POLITECNICO DI TORINO
Corso di Dottorato in Matematica per le Scienze dell'Ingegneria - XXI Ciclo

Tesi di dottorato - MAT/05

Distributed optimization with communication constraints

Paolo FRASCA

DISSERTATION COMMITTEE:
prof. Fabio Fagnani, PoliTo, Advisor
prof. Sandro Zampieri, Università di Padova, Advisor
prof. Nicola Bellomo, PoliTo, PhD program director
prof. Andrea Bacciotti, PoliTo
prof. Alessandro Chiuso, Università di Padova
prof. Lane Hughston, Imperial College, London

March 2009

To my late grandfather.

Abstract

This thesis addresses problems of optimization to be solved collectively (distributely, we say) by groups of networked *agents* having limited communication and computation capabilities. The focus is on two paradigmatic problems, average consensus and coverage control. The former consists in computing the average of values which are priory known to the agents, while the latter consists in optimally deploying robotic agents in a given environment. Their solution is possible thanks to iterative algorithms which exploit the available communications among agents, described by a graph.

Research has been very intense in latest years on these problems, due to strong applicative motivations: these algorithms are crucial for the design of networks of sensors, and for the coordinated motion of groups of unmanned vehicles. Namely, the ultimate goal of these algorithms is to make such artificial networks able to self-organize and perform specific tasks without a centralized supervision.

In the thesis networks have been considered, which allow only communication in pairs, or limited to messages belonging to a discrete alphabet. Several novel algorithms for the solution of the above optimization problems are proposed: they are proved to lead the agents to the desired solution, and their performance is studied. The mathematical tools required for the analysis are novel, or need to be used in a non-standard way. They are drawn from very different branches of pure and applied mathematics: spectral graph theory, Markov chains, combinatorics, linear algebra, set-valued analysis, convex and discrete geometry, topology of hyper-spaces of sets.

Contents

Abstract i

Table of contents iii

Acknowledgements vi

General introduction 1
- Problem overview 1
- Mathematical tools 3
- Statement of contributions 3
- Perspectives 5

I Digital consensus 7

1 Average consensus algorithms and network models 9
- 1.1 Notations 9
- 1.2 Graph and matrix theory 10
  - 1.2.1 Examples 11
- 1.3 Average consensus problems and algorithms 13
  - 1.3.1 How to construct a Perron matrix 16
  - 1.3.2 Scaling of performance 17
  - 1.3.3 Gossip algorithms 17
- 1.4 Networks of digital links 18
  - 1.4.1 Quantizers 19
  - 1.4.2 Consensus on non ideal networks 20
  - 1.4.3 Related works 22

2 Static quantization and fixed topology 23
- 2.1 The algorithm 23
  - 2.1.1 Deterministic analysis 25
- 2.2 Bounded error model 26
- 2.3 Probabilistic model 34
- 2.4 Simulations and conclusion 41

2.4.1 Speed of convergence 41
2.4.2 Dependence on γ 41
2.4.3 Dependence on N 44
2.4.4 Conclusions 46
2.A Analysis of the Example 2.2 47
2.B Proof of Proposition 2.3 50

3 Quantized gossip **53**
3.1 Problem statement 54
3.2 Gossip algorithms with deterministic quantizers 55
3.2.1 Partially quantized strategy 55
3.2.2 Globally quantized strategy 60
3.2.3 Speed of convergence 62
3.3 Gossip algorithms with probabilistic quantizers 63
3.3.1 Partially quantized strategy 64
3.3.2 Globally quantized strategy 70
3.4 Conclusion and open questions 72

4 Efficient quantization **75**
4.1 Encoding-decoding schemes 75
4.2 Zooming-in zooming-out strategy 76
4.3 Simulation results 80
4.4 Conclusions 81

5 Perspectives on the consensus problem **85**

II Partitions optimization **87**

6 Gossip coverage optimization **89**
Introduction 89
6.1 Coverage optimization via multicenter functions 90
6.1.1 Partitions and multicenter optimization 90
6.1.2 Distributed coverage control and its limitations 93
6.2 Partitions-based gossip coverage algorithm 94
6.2.1 The partitions-based gossip coverage algorithm 94
6.2.2 Designing a continuous algorithm 97
6.2.3 Analysis and convergence results 99
6.2.4 Simulation results and implementation remarks 101
6.3 LaSalle invariance principle for set-valued maps 102
6.4 The set of partitions 107
6.5 Continuity of relevant functions and maps 109
6.5.1 Convergence proofs 114

7 Gossip in one dimensional partitions 115

7.1 Special properties 115
7.2 Uniform density: equitable partitioning and average consensus 118
7.3 Robots motion protocols 124

8 Perspective on partitions optimization 125

Bibliography 133

Acknowledgements

Primum pro nummata vini: The three years work on which this thesis is based would not have been possible without the support of the following organizations and people.

- Fondazione CRT - Istituto Lagrange, which paid my salary and covered part of my travel expenses, through the grant *Coordinated control and identification of dynamical systems with communication constraints*;

- Istituto Nazionale di Alta Matematica, which covered some of my travel expenses, through various grants;

- the Dipartimento di Matematica del Politecnico di Torino, which hosted my activities for most of the time;

- prof. Francesco Bullo, which allowed me to spend a significant time abroad, at the University of California at Santa Barbara.

From the scientific side, I am deeply grateful to my mentors during these three years: my advisors Fabio Fagnani and Sandro Zampieri, and Francesco Bullo, for their patience and their efforts to teach me how to do the job, and the need for clarity above all. Special thanks go to Ruggero Carli, my most constant coauthor. I am also indebted to Francesco Bullo for many LaTeX suggestions, and for allowing me to take advantage of its bibliographical files.

Finally, many people influenced, encouraged, helped, or taught me something during these years. I'll try to list here some of them, with no claim of completeness, in alphabetic order: Andrea Bacciotti, Nicola Bellomo, Giacomo Como, Federica Garin, Fabio Pasqualetti, Jacob Pejsachowicz, Paolo Tilli.

General introduction

Problems overview

In the latest years, the applied mathematics community has witnessed an increasing interest for studying control, estimation and algorithmic problems over networks. A common feature of these problems is the fact that there is a fundamental constraint on the information flow: data and capabilities are *distributed* across a possibly large number of *agents*, which communicate among each other through some communication network. A quite large variety of different engineering problems do fit in this general setting. Concrete specifications can be obtained by choosing the common goal, the agents internal model, the communication network. Such problems are mostly new to the engineering community and give to the applied mathematician the way of posing and solving interesting mathematical problems. This thesis aims at contributing to this emerging field, focusing on two prototypical problems: the *consensus* and *coverage* problems. Both of them are special cases of optimization problems in a distributed system of agents.

Hence, we shall introduce distributed optimization, consensus, and coverage problems. In general, an optimization problem can be written as

$$\min_{\theta \in \Theta} f(\theta). \tag{1}$$

A distributed optimization problem (for N agents) is an optimization problem in which $f(\theta)$ depends on a vector of functions $(f_i(\theta))_{i=1}^N$. Each function f_i is *local* to agent i, in the sense that it can be computed by the agent i using the available information. Since communication among the agents is limited to a communication graph, these optimization problems can not be solved by usual algorithms, but need *iterative* procedures involving some sort of *information spreading*. Typical theoretical results include proving the convergence of the proposed algorithms and estimating its speed of convergence.

Let us now introduce an important family of problems. The function f is said [66] to be *separable* if

$$f(\theta) = \sum_{i=1}^{N} f_i(\theta). \tag{2}$$

These functions are attracting a special attention [79, 66, 71, 70] because, in spite of their simple structure, include important examples. The problems studied in this thesis fall in this framework.

The average consensus problem

Suppose we have a directed graph $\mathcal{G} = (V, E)$, with nodes set $V = \{1, \dots, N\}$ and edges set $E \subset V \times V$, and a real quantity x_i for every node $i \in V$. The *average consensus problem* consists of computing the average $x_{\text{ave}} = N^{-1} \sum_{i=1}^{N} x_i$ in an iterative and distributed way, exchanging information among nodes exclusively along the available edges in $\mathcal{E}$. The average consensus problem can be seen as the quadratic minimization problem relative to the function $f(\theta)$ in (2), obtained by choosing $f_i(\theta) = (x_i - \theta)^2$. Then the minimum (1) is achieved for $\theta = x_{\text{ave}}$. This problem appears in the control community since the early 80's (decentralized computation [85], load balancing [32]) and, more recently, has attracted much attention [73, 24] for a variety of possible applications:

(i) data fusion [89, 13] in sensors networks;

(ii) clocks synchronization [58, 16];

(iii) coordinated control of mobile autonomous agents.

In the latter case, we want a group of self-propelled agents to complete a task involving coordinated motion, without any external supervision. Examples of tasks include going to a *rendezvous* point, aligning velocities, and reaching or keeping a spacial disposition (formation). Early works on the subject include [86, 53], and more recent works include [61, 62, 28, 50], while an accessible introduction can be found in [63]. Two recent books [81, 11] are available on the subject.

Moreover, related problems have been studied in the probability [26] and computer science [75] communities since the 70's, as well as by game theorists in the context of coordination games[91, 64].

The average consensus problem is especially attractive for an applied mathematician because, while it is suitable to a variety of applications, its simplicity allows a rather deep theoretical understanding.

The coverage problem

Coverage control has the goal of optimizing the deployment, that is the spatial configuration, of a group of robots in a given environment [11]. The optimization problem can be written in terms of a cost function, depending on the agents' positions. In this form, the problem has been recently solved in a distributed way [29].

In this thesis, a novel approach to the problem is taken. We aim at the optimization of a *partition* of the given environment, each portion of which is assigned to one agent. Optimality is then defined with respect to how well the agents can control their own regions. To be more formal, let the environment to apportion be Q, a compact subset of $\mathbb{R}^d$, and let the variable $\theta = \{\theta_i\}_{i=1}^N$ belong to a space of partitions of Q into subregions, assigned to each agent. Let ϕ be a measure on Q, and $\mathrm{Cd}(\theta_i)$ denote the centroid of the region θ_i. Then if we define

$$f_i(\theta) = \int_{\theta_i} \|y - \mathrm{Cd}\,(\theta_i)\|^2 d\phi,$$

we have written the coverage control problem as a distributed optimization problem.

To solve the problem, we design iterative algorithms which, starting from an initial partition, at each time step update the partition according to some local rule. The goal is to make the sequence of partitions converge to an optimal one. Among possible algorithms, we shall prefer those which require the least communication effort.

Mathematical tools

The spectrum of mathematics which are involved in these problems is extremely wide, depending on the variety of problems. Example giving, the specific case of the coverage problem studied in Chapter 6 has required techniques from convex geometry and analysis. Since the emerging field of *distributed control and optimization* lies at the crossing of control, optimization, computer science and information theory, many mathematical tools come from these fields. Classical optimization methods are obviously a reference. Besides, we will happen to use results from stability and feedback control of dynamical systems, both in classical settings and in the context of switching systems.

Other mathematical tools, like *graph theory*, are specifically relevant for distributed system. Indeed, since the potential communication are effectively depicted using a graph, many aspects of graph theory are relevant. Most interest has been devoted to spectral graph theory: interesting connections have been found between the eigenvalues of the adjacency matrix of the graph and the performance of various algorithms. An important role is played by *discrete mathematics and combinatorics*, for various reasons. First, as a background for graph theory, and second, because of the application to digital processors, which inherently deal with finite sets of inputs and outputs. An example is the family of symbolic dynamics induced by quantization, that we are going to study along Chapters 2 and 3. Probabilistic tools, mostly about Markov chains, are of the greatest importance. This happens because in most cases the performance of the algorithms is evaluated on random instances, and several proposed algorithms are themselves randomized. Moreover, consensus problems have a natural connection with the classical problem of the random walk on a graph.

Finally, we want to stress the interplay between computer simulations, useful to provide experimental evidences about the algorithms, and the rigorous analysis which is offered by mathematics. The role of mathematician in dealing with such delicate issues is crucial, and appreciated by the engineering community.

Statement of contributions

Our contribution in this emerging field has been focused on the design and analysis of distributed algorithm to solve the problems of consensus and coverage under severe communication limitations. As we said, available links between agents are usually described as a graph. Most precedent literature has assumed the communication channels between the agents to be analog, permanent, reliable, and able to transmit information instantaneously and correctly. Our contribution has been in designing and analyzing algorithmic solutions for cases in which the above assumptions do not hold.

The results of the thesis have been mostly published, or submitted for publication, in conferences proceedings [19, 18, 20, 21, 40, 38] or journals [41, 23]: note however that the content of the thesis is wider and contains further results and insights. Note moreover that some work of mine, done in preparation to my master's thesis, has appeared in [30]. The latter paper deals with multi-agent systems from the perspective of mathematical physics and modeling, rather than of control, and its content is not included in this thesis.

Part I: Digital consensus

This research has been done in collaboration with Ruggero Carli, Fabio Fagnani, Thomas Taylor (Arizona State University), and Sandro Zampieri. My stay at the University of Padova in fall 2006 has been the occasion to start working on this topic.

All literature on the consensus problem, until year 2006, assumed the agents to be able to communicate with their neighbors in the graph through an ideal communication channel: instantaneous, reliable, and able to transmit real numbers. Relaxing these assumption to consider digital channels, able to send quantized information, have given rise to the research exposed in Part I of the thesis. The material is organized as follows. In Chapter 1 we survey the state of the art about average consensus algorithms, with no claim of completeness, but rather in view of presenting our results. Moreover, we introduce the quantization constraint: the messages the agents can send belong to a discrete space, and are obtained as the image of the agent state through a map, called the quantizer. Several quantizers are introduced, and we preliminarily discuss their properties and impact on consensus algorithms.

In the following chapters, three specific problems are considered.

(i) Chapter 2. *Deterministic algorithms for average consensus over networks with time-invariant topology and digital channels.* Deterministic time-invariant consensus algorithms have been studied in many papers, starting from the pioneering work [85]. In this chapter we consider the constraint of having a digital communication with messages belonging to a countable set, and we propose a novel update rule for the agents' states. Such update rule has the advantages of simplicity and of preserving the average through iterations, but due to the quantization, consensus can not be approached with arbitrary precision. We study the committed steady state error, as well as the speed of convergence of the algorithm, as functions of the communication graph. Understanding how the properties scale with N has been a major concern. Due to the non-linearity of the system, most results are obtained by a bounded error model, which is studied with both a worst-case and a probabilistic approach. Most results appeared in [18, 20, 21, 41].

(ii) Chapter 3. *Randomized "Gossip" algorithms for consensus, using pairwise digital communication.* At least since [9], people have studied algorithms in which only random subsets of the edges are used, or available, at each time step. Namely, the attention has been focused on communication in pairs. Indeed, pairwise communication can be easier to establish in an hazardous environment, and algorithm using pairwise communication have shown to be effective in reducing the amount of communication required for convergence. In the chapter we study average consensus problems in which communication, besides being quantized, is constrained to happen between pairs of agents. We consider two different update rules, and the advantage of using a randomized quantizer instead of a deterministic one. This is done by a detailed probabilistic analysis of the convergence of such algorithms, and of their performance in terms of achievable precision. Roughly the same set of results is exposed in [40, 23].

(iii) Chapter 4. *Efficient quantization techniques.* Assuming a static topology and the use of a saturated quantizer with finite range, we propose a more elaborated coding-decoding scheme, which is able to drive the systems to converge with arbitrary precision to average consensus at exponential speed, at the price of using minimal additional memory and computational resources for coding and decoding. The result appeared in [19], and is presented here and in the report [22] in a slightly refined version, together with encouraging simulations.

Finally, Chapter 5 is devoted to comment on the obtained results, and point out perspective research in this field.

Part II: partitions optimization with pairwise communication.

This research has been developed in collaboration with Francesco Bullo and Ruggero Carli, during my stay as a visiting scholar *at the University of California at Santa Barbara in spring 2008.*

Part II of the thesis undertakes the solution of coverage problems for convex environments under the constraint of pairwise communication, thanks to a novel approach to coverage problems. In previous literature [29, 11], a coverage algorithm is described as a dynamical system of the robots positions: in our approach, the state space consists of the *partitions of the given environment into subregions*. The convergence properties of the gossip algorithm are established through two results of independent interest. First, describing the properties of a suitable space of partitions, and the continuity of certain geometric maps. Second, proving a LaSalle invariance principle for switching maps on metric spaces. While the framework is set, and convergence is proven, for an environment of any dimension (Chapter 6), a special attention is devoted to the one dimensional case (Chapter 7). Indeed the latter enjoys a fruitful connection with the average consensus problem. A brief account of these results has appeared in [38], and a more extended report is available in [39].

Perspectives

The research presented above is rich in potential developments, as the whole field of control with limited communication: see Chapter 5 and Chapter 8 for a concise account of open problems and research perspectives on consensus and coverage problems, respectively. We wish to present here two open problems, which we consider of strategic interest, and which we are going to address in the next future. First, in the consensus problem, the constraint of quantized communication has been applied, in the current literature, with an assumption of reliability of the channel. It is thus apparent the interest of relaxing this assumption, and providing algorithms able to assure convergence to consensus in presence of digital noisy channels. Second, the novel approach to coverage problems, by means of dynamical systems of the partitions, is promising of wider application: in more general environments (non-convex or discrete), and using more general cost functions.

Part I

Digital consensus

Chapter 1

Average consensus algorithms and network models

This first chapter, besides giving some technical preliminaries, introduces the consensus problems on networks of digital links, which are studied in Chapters 2, 3 and 4.

After introducing some standard notations in Section 1.1, in Section 1.2 we lay down the graph theoretical framework which takes into account the network structure, together with some relevant result. Then, in Section 1.3 we introduce the consensus problem, and overview the state-of-the-art algorithms for solving it on networks allowing perfect transmission of information. Section 1.4 discusses the way to model a network of digital or digital noisy links, introducing some relevant quantizers, and proposing a framework to adapt the "ideal case" algorithms to more general networks model. Such adaptation has been introduced by the author in [18]: from Section 1.4 on, the contents of the thesis are novel.

1.1 Notations

The symbols $\mathbb{N}, \mathbb{Z}, \mathbb{R}, \mathbb{C}$ denote the natural, integer, real and complex numbers, respectively. The symbols $\mathbb{Z}_{\geq 0}, \mathbb{R}_{>0}, \mathbb{R}_{\geq 0}$ denote the nonnegative integers, and the positive and nonnegative reals, respectively. If $x \in \mathbb{C}$, then $|x|$ denotes its absolute value, whereas if A is a set, $|A|$ denotes its cardinality. The notations $\lfloor \cdot \rfloor$ and $\lceil \cdot \rceil$ are used to denote the floor and ceiling maps from $\mathbb{R}$ to $\mathbb{Z}$.

Given $N \in \mathbb{N}$, $\mathbf{1}_N \in \mathbb{R}^N$ is a vector having all entries equal to 1, and I_N is the identity matrix of dimension N. Given a matrix $M \in \mathbb{R}^{m \times n}$, M^* denotes its conjugate transpose. Given a subspace V of $\mathbb{R}^N$, $V^\perp$ denotes the subspace orthogonal to V. The matrix $\Omega_N = I - N^{-1}\mathbf{1}_N\mathbf{1}_N^*$ represents the operator projecting on the space $\mathrm{span}\{\mathbf{1}_N\}^\perp$, where $\mathrm{span}\{v\} = \{x \in X \mid x = \lambda v, \lambda \in \mathbb{R}\}$ is the subspace spanned by a vector v. When sizes can be inferred from the context, the subscripts are dropped. Given a matrix $M \in \mathbb{R}^{N \times N}$, $\mathrm{diag}(M)$ means a diagonal matrix with the same diagonal elements of the matrix M. Given a vector $m \in \mathbb{R}^N$, $\mathrm{diag}(m)$ means a diagonal matrix having the components of m as diagonal elements. Given a vector $x \in \mathbb{R}^N$ with $\|x\|$ and $\|x\|_\infty$ we denote respectively the Euclidean norm and the sup-norm. Accordingly, given a matrix $M \in \mathbb{R}^{N\times N}$, with $\|M\|$ and $\|M\|_\infty$ we denote the induced operator norms.

$\mathbb{P}$ and $\mathbb{E}$ denote the probability of a given event and its mathematical expectation.

To compare the asymptotic behavior of two functions $f, g : \mathbb{N} \to \mathbb{R}$, we use the notations:

- $f = o(g)$ if $\lim_{n\to\infty} \frac{f(n)}{g(n)} = 0$;
- $f \asymp g$ if $\lim_{n\to\infty} \frac{f(n)}{g(n)} = 1$;
- $f = \Theta(g)$ if there exist $N_0 \in \mathbb{N}$ and k_1, k_2 positive real numbers, such that $k_1|g(n)| \le f(n) \le k_2|g(n)|$ for all $n \ge N_0$;
- $f = O(g)$ if there exist $N_0 \in \mathbb{N}$ and a positive real number k, such that $f(n) \le k|g(n)|$ for all $n \ge N_0$;
- $f = \Omega(g)$ if there exist $N_0 \in \mathbb{N}$ and a positive real number k, such that $f(n) \ge kg(n)$ for all $n \ge N_0$.

1.2 Graph and matrix theory

Matrices

A matrix $M \in \mathbb{R}^{N\times N}$ is said to be *normal* if $A^* A = AA^*$. A normal matrix is always diagonalizable, namely a matrix is normal if and only if it can be diagonalized by a unitary matrix. A matrix $M \in \mathbb{R}^{N\times N}$ is said to be *nonnegative* if $M_{ij} \ge 0$ for all i and j. A matrix M is said to be *stochastic* if it is nonnegative and the sums along each row are equal to 1. Moreover a matrix M is said to be *doubly stochastic* if it is stochastic and also the sums along each column are equal to 1. Given a matrix $M \in \mathbb{R}^{N\times N}$, let $\sigma(M)$ denote the set of eigenvalues of M and $\rho(M)$ the spectral radius of M: $\rho(M) = \max\{|\lambda| : \lambda \in \sigma(M)\}$. When the matrix M is stochastic, $\rho(M) = 1$, and thus it is worth to define the *essential* spectral radius as

$$\rho_{\text{ess}}(M) = \begin{cases} 1 & \text{if } \dim \ker(M - I) > 1 \\ \max\{|\lambda| : \lambda \in \sigma(M) \setminus \{1\}\} & \text{if } \dim \ker(M - I) = 1\,, \end{cases} \tag{1.1}$$

Graphs

Communication links among agents are modeled by a *directed graph* $\mathcal{G} = (V, E)$ where $V = \{1, \dots, N\}$ is the set of vertices and E is a subset of $V \times V$. The cardinality of V is said to be the order of the graph. An element of E is said to be a (directed) edge. Any $(i, i) \in E$ is called a *self loop*. A path in $\mathcal{G}$ consists of a sequence of vertices $i_1 i_2 \ldots\ldots i_r$ such that $(i_\ell, i_{\ell+1}) \in E$ for every $\ell = 1, \dots, r-1$; i_1 (resp. i_r) is said to be the initial (resp. terminal) vertex of the path. A vertex i is said to be connected to a vertex j if there exists a path with initial vertex i and terminal vertex j. A graph is said to be *undirected* (or *symmetric*) if $(i, j) \in E$ implies that $(j, i) \in E$. A directed graph is said to be (weakly) connected if, given any pair of vertices i and j, either i is connected to j or j is connected to i. A directed graph is said to be *strongly connected* if, given any pair of vertices i and j, i is connected to j. On an undirected graph, these two notions coincide, and we just talk about connectivity. A graph is said to be *fully connected* or *complete* if $E = V \times V$. Given $\mathcal{G}_1 = (V, E_1)$ and $\mathcal{G}_2 = (V, E_2)$, we say that $\mathcal{G}_1 \subset \mathcal{G}_2$ if $E_1 \subset E_2$.

The *adjacency matrix* of $\mathcal{G}$, denoted by $A_{\mathcal{G}}$ or just A, is a $\{0, 1\}$-valued matrix indexed by the elements in V defined by letting $A_{ij} = 1$ if and only if $(j, i) \in E$ and $j \neq i$. Define the *in-degree* of a vertex i as $\sum_j A_{ij}$ and the *out-degree* of a vertex j as $\sum_i A_{ij}$. A graph is said to be d-regular if the in-degree and the out-degree of all its nodes is equal to d. In an undirected graph the in-degree and the out-degree are equal, and we just talk of *degree* of a node.

Given a graph $\mathcal{G}$, denote the in-degree of node $i \in V$ by d_i, and $D_{\mathcal{G}} = \mathrm{diag}\,(d_1, \ldots, d_N)$. Then, the Laplacian matrix of the graph $\mathcal{G}$ is defined as $L_{\mathcal{G}} = D_{\mathcal{G}} - A_{\mathcal{G}}$. Such matrix, and its eigenvalues, are of the greatest importance to understand the properties of graph $\mathcal{G}$ and of algorithms based on it. For a general treatment, we refer the reader to [25].

Given a nonnegative matrix $M \in \mathbb{R}^{N \times N}$, we can define an induced graph $\mathcal{G}_M = (V, E_M)$ by putting an edge (j, i) in $E_M \subset (V \times V)$ if $M_{ij} > 0$. Given a graph $\mathcal{G}$ on V, a matrix M is said to be *adapted* or *compatible* with $\mathcal{G}$ if $\mathcal{G}_M \subset \mathcal{G}$.

1.2.1 Examples

We describe here some relevant families of graphs, which are used later as examples.

Example 1.1 (Abelian Cayley graphs and matrices) An interesting class of graphs are the *Cayley graphs* defined on Abelian groups [4]. Let G be any finite Abelian group of order $|G| = N$, whose operation is denoted by $+$, and let S be a subset of G. The Cayley graph $\mathcal{G}(G, S)$ is the directed graph with vertex set G and arc set

$$E = \{(g, h) \in G \times G : h - g \in S\}\,.$$

A notion of Cayley structure can also be introduced for matrices. A matrix $M \in \mathbb{R}^{G \times G}$ is said to be a *Cayley* matrix over the group G if

$$M_{i,j} = M_{i+h,j+h} \qquad \forall\; i, j, h \in G\,.$$

This means that a generating vector $\pi \in \mathbb{R}^N$ exists such that $M_{i,j} = \pi(i - j)$.

Note that the adjacency matrix of a Cayley graph is a Cayley matrix, defined on the same group. Cayley matrices enjoy many properties: among others, they are *normal*. Moreover, it is possible to write their spectrum explicitly [24].

We now list some important examples of Cayley graphs

Example 1.2 (Circulant graphs) If $G = \mathbb{Z}_N$, the cyclic group of order N, then the Cayley graph $\mathcal{G}(G, S)$ is said to be *circulant*, and the Cayley matrix M is said to be a *circulant* matrix [33].

Example 1.3 (Tori and ring graphs) Let $G = \mathbb{Z}_n^d$, and let e_i be the vector of length d having the i-th entry equal to 1 and the other entries equal to 0. The Cayley graphs $\mathcal{G}(G, \{e_i\}_{i=1,\dots,d})$ and $\mathcal{G}(G, \{e_i\}_{i=1,\dots,d} \cup \{-e_i\}_{i=1,\dots,d})$ are called the *directed and undirected d-dimensional torus graphs*, respectively. The important special cases when $d = 1$ are circulant graphs, and we refer to them as *rings*.

The graph which is obtained removing one edge from an undirected ring is called a *line* graph.

Example 1.4 (Hypercube graphs) The n-dimensional hypercube graph, the Cayley graph $\mathcal{G}(\mathbb{Z}_2^n, \{e_i\}_{i=1,\dots,n})$, is the graph obtained drawing the edges of a n-dimensional hypercube. It has $N = 2^n$ nodes which can be identified with the binary words of length n. Two nodes are neighbors if the corresponding binary words differ in only one component. Thus the graph is undirected and every node has degree n.

Example 1.5 (Random geometric graphs) A random geometric graph of order N is obtained through the following probabilistic construction. Consider N random points in the unit square, independently drawn and uniformly distributed. Then, two vertices have an edge between them, if their Euclidean distance is below a threshold R. Two examples are in Figure 1.1.

The random geometric graph is commonly used to model wireless networks [49]. Of special interest are families of random geometric graphs of increasing order, in which R depends on N. Among them, the choice $R = \Theta(\sqrt{\log N/N})$, for $N \to \infty$, has attracted special attention, as it is able to assure connectivity of the graph with probability going to 1 as N goes to infinity, while keeping low the number of edges. For a complete treatment of random geometric graphs, see [76].

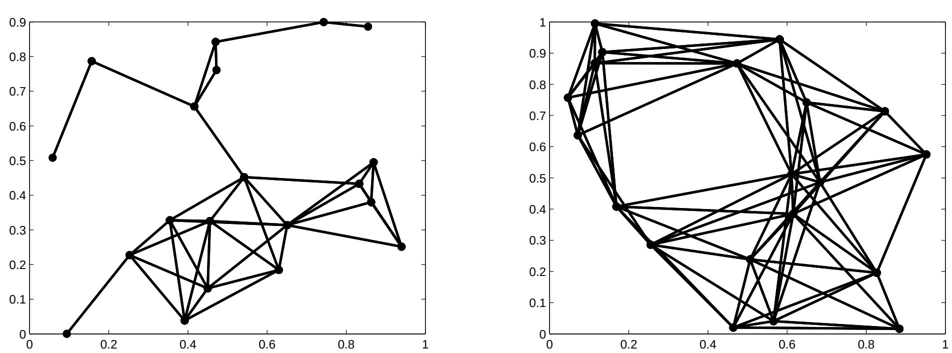

Figure 1.1: Two sample random geometric graphs, generated with $N = 20$ and $R = 0.3$ and $R = 0.5$, respectively.

1.3 Average consensus problems and algorithms

The *average consensus problem* is the problem of driving the states of a set of dynamical systems to a final common state which corresponds to the average of initial states of each system. In a broader sense, a *consensus* problem requires to drive the states to *any* common state.

This mathematical problem can be seen as the simplest example of coordination task. In facts, it can be used to model both the control of multiple autonomous vehicles which all have to be driven to the centroid of the initial positions (rendezvous), and to model the decentralized estimation of a quantity from multiple measures coming from distributed sensors. For this combination of simplicity and significance it is of great interest to the applied mathematician. In the latest years, algorithms to solve consensus problems have attracted a lot of interest. Several algorithms for average consensus have been proposed in the literature: among the vast literature, we refer the reader to [73] [24] and references therein. The difficulty of the problem resides in the communication constraints which are given to the agents, represented by a graph: nodes are agents and edges are available communication links. Let us now be more precise in stating the problem. First, with the aim on giving more emphasis to the role of communication, than to the dynamics of single agents, we restrict ourselves to sets of identical discrete-time linear systems, whose coupling is described by a graph.

Let there be a graph of agents $\mathcal{G} = (V, E)$, with $V = \{1, \ldots, N\}$. We assume that only if $(j, i) \in E$, j can transmit information about its state to i. Let there also be N systems, one for each agent, whose dynamics are described, for $t \in \mathbb{N}_0$, by the following discrete-time state equations

$$x_i(t+1) = x_i(t) + u_i(t) \qquad i = 1, \ldots, N$$

where $x_i(t) \in \mathbb{R}$ is the state of the i-th system/agent, and $u_i(t) \in \mathbb{R}$ is the control input. More compactly we can write

$$x(t+1) = x(t) + u(t) \tag{1.2}$$

where $x(t), u(t) \in \mathbb{R}^N$.

Definition 1.1 (The consensus problem) *Given the system* (1.2)*, the* consensus problem *consists in designing a sequence of control inputs $u(t)$ yielding the consensus of the states, namely*

(i) if $x(0) = \lambda\mathbf{1}$ with $\lambda \in \mathbb{R}$, then $x(t) = x(0)$ for all $t \in \mathbb{N}$.

(ii) it exists $\alpha \in \mathbb{R}$ such that

$$\lim_{t\to+\infty} x(t) = \alpha\mathbf{1}. \tag{1.3}$$

If we require $\alpha = x_{\mathrm{ave}}(0)$, we have the average consensus problem.

The important case we restrict to, and which has been mostly studied in literature, is a static linear feedback

$$u(t) = Kx(t) \qquad K \in \mathbb{R}^{N\times N} \tag{1.4}$$

In such case the system (1.2) is given by the following closed loop system

$$x(t+1) = (I+K)x(t). \tag{1.5}$$

The matrix $I+K$ is commonly called a Perron matrix, and denoted by P. The problem of designing the controller thus reduces to find a matrix K *adapted to* $\mathcal{G}$.

Methods for design are already clear in the literature [24].

Proposition 1.1 *The consensus problem for system (1.5) is solved if and only if the following condition holds.*

(i) The only eigenvalue of $I+K$ on the unit circle is 1;

(ii) the eigenvalue 1 has algebraic multiplicity one (namely it is a simple root of the characteristic polynomial of $I+K$), and $\mathbf{1}$ is its eigenvector;

(iii) all the other eigenvalues are strictly inside the unit circle.

In the sequel we will restrict to matrices K such that $I+K$ is a nonnegative matrix, that is a matrix with all elements nonnegative. Proposition 1.1 then implies that $I+K$ has to be a stochastic matrix. Moreover,

$$\lim_{t\to\infty}(I+K)^t = \mathbf{1}v^*$$

where $v \in \mathbb{R}^N$ is the unique probability vector such that $v^*(I+K) = v^*$. This implies that

$$\lim_{t\to\infty} x(t) = v^* x(0)\mathbf{1}.$$

In the special case when $v = N^{-1}\mathbf{1}$ we obtain that the consensus is achieved at the average of the initial conditions. In this case, $P = I+K$ is a doubly stochastic matrix and K is said to be an *average consensus controller*.

The following assumption on P guarantees that K is an average consensus controller. We are going to use it during Chapters 2 and 4.

Assumption 1.1 *P is a doubly stochastic matrix such that $P_{ii} > 0$, for all $i \in \{1,\dots,N\}$, and $\mathcal{G}_P$ is strongly connected.*

We collect here some facts about such matrix P, which are easy to check or consequences of the Perron-Frobenius Theorem [44].

Lemma 1.2 *Let P satisfy Assumption 1.1 and let $\Omega = I - N^{-1}\mathbf{1}\mathbf{1}^*$. Then the following facts hold.*

- *The eigenvalues of P are bounded in modulus by 1.*

- *There is only one eigenvalue of modulus 1. It is simple, and equal to 1. The corresponding eigenvector is* $\mathbf{1}$.
- $\lim_{t\to\infty} P^t = N^{-1}\mathbf{1}\mathbf{1}^*$.
- ΩP *has the same eigenvalues and eigenvectors system as* P*, except that the eigenvalue 1 is replaced by 0. This implies that*
$$\lim_{t\to\infty} \Omega P^t = 0.$$
- $P\Omega = \Omega P$ *and* $\Omega^2 = \Omega$.
- $\|P\Omega\| < 1$.

The lemma implies convergence to the average consensus.

Proposition 1.3 (Average consensus) *Let* $x(t)$ *evolve following* (1.5)*, and* P *satisfy Assumption 1.1. Then* $\lim_{t\to\infty} x(t) = x_{\text{ave}}(0)\mathbf{1}$.

Conversely, the following result [24] characterizes those graphs $\mathcal{G}$ for which the average consensus problem is solvable, in the sense that an average consensus controller compatible with $\mathcal{G}$ can be found.

Proposition 1.4 (Solvability of average consensus) *Let* $\mathcal{G}$ *be a directed graph containing all self-loops. The following conditions are equivalent:*

(i) The average consensus problem is solvable on $\mathcal{G}$.

(ii) $\mathcal{G}$ *is strongly connected.*

Furthermore, if the above conditions are satisfied, any K *such that* $I+K$ *is doubly stochastic and* $\mathcal{G}_{I+K} = \mathcal{G}$*, solves the average consensus problem.*

For the analysis of consensus algorithms, it is useful to introduce some quantities, depending on the state $x(t)$.

Definition 1.2 (Consensus figures) *Let* $x(t)$ *be the state of a system evolving through a consensus algorithm. Then we define the* disagreement *vector*

$$y(t) = x(t) - x_{\text{ave}}(t),$$

the scalar bias

$$z(t) = |x_{\text{ave}}(t) - x_{\text{ave}}(0)|,$$

and the total disagreement

$$\bar{y}(t) = x(t) - x_{\text{ave}}(0).$$

Remark that

$$y(t) = \Omega x(t).$$

Moreover, let

$$d(t) = \frac{1}{\sqrt{N}}\|y(t)\|.$$

Thus the latter quantity represents the normalized distance of the state $x(t)$ from the average of the states.

Whenever an algorithm drives a system to consensus, $\lim_{t\to\infty} y(t) = 0$. It is thus interesting to evaluate the rate of convergence of the algorithm, that we define as

$$R = \sup_{x(0)} \limsup_{t\to\infty} d(t)^{1/t}. \tag{1.6}$$

For the linear algorithm (1.5), we have the following result [24].

Proposition 1.5 (Rate of convergence) *Let $x(t)$ follow* (1.5)*. Then the rate of convergence is a function of P. Namely,*

$$R(P) = \rho_{\mathrm{ess}}(P).$$

1.3.1 How to construct a Perron matrix

Restricting to Perron matrices satisfying Assumption 1.1, we are left with the problem of, given a strongly connected graph, designing one such matrix, adapted to it. Proposition 1.4 assures that this can be done: the interest is then in making an optimal choice with respect to some performance parameter. The more natural and investigated index is the rate of convergence $R(P)$, which reduces to studying $\rho_{\mathrm{ess}}(P)$. This problem is well known in the literature [8], since it correspond to designing the fastest mixing Markov chain on the given graph. An account of this optimization problem, from the point of view of consensus, can be found in [24].

The problem, though well known, is not simple: example giving, there are many degrees of freedom (one for each edge). For this reason, we list here some of the most natural restrictions which can be imposed to the matrix P to obtain a simpler problem. This exposition is also preliminary to discuss specific optimization problems, later in the exposition.

(i) Given a graph $\mathcal{G}$, for any *gain* $\gamma \in (0,1)$ we can use a *weighted Laplacian* Perron matrix

$$P = I - \gamma D^{-1} L_{\mathcal{G}}.$$

This choice leaves the gain parameter free for optimization. This construction does not give a doubly stochastic matrix unless

$$\sum_{i\neq j} d_i^{-1} A_{ij} = 1 \quad \forall\, j \in V,$$

which is true if $\mathcal{G}$ is symmetric or regular.

(ii) if $\mathcal{G}$ is d-regular, we can set $\gamma = \frac{d}{d+1}$ in (i), obtaining a *uniform* Perron matrix $P = \frac{1}{d+1}(I + A_{\mathcal{G}})$, which is doubly stochastic.

(iii) Posing $d_M = \max_{i\in V} d_i$, we can define a *Maximum degree* Perron matrix as

$$P = \left(1 - \frac{1}{d_M+1} D\right) I + \frac{1}{d_M+1} A,$$

which is doubly stochastic when the out-degree of each node equals its in-degree.

1.3.2 Scaling of performance

We introduce here an important issue for the analysis of consensus algorithm, that of how performance scales with increasing the number of agents. To see this, we consider sequences of graphs of increasing size, and the speed of convergence as the performance index. E.g., for $N \in \mathbb{N}$, consider ring graphs on N nodes, and uniform weights matrices P_N adapted to them. A natural question is: does $\rho_{\text{ess}}(P_N)$ depend on N? From the properties of circulant matrices [33] it comes out that

$$\begin{aligned}\rho_{\text{ess}}(P_N) &= \left(\frac{1}{2} + \frac{1}{2}\cos\frac{2\pi}{N}\right)^{\frac{1}{2}} \\ &\asymp 1 - \frac{\pi^2}{2}\frac{1}{N} \qquad \text{for } N \to \infty.\end{aligned}$$

Then the rate of convergence goes to 1 as N goes to infinity, at polynomial speed. Such increase of the spectral radius is much more general, as we can see from the following result regarding Cayley graphs.

Lemma 1.6 ([24]) *Let G be a finite Abelian group of order N and $S \subset G$. Let $\mathcal{G}(G, S)$ be the Cayley graph associated with G and S, and P any matrix adapted to it. If $|S| = \nu$, then $\rho_{\text{ess}}(P) = 1 - CN^{-2/\nu}$, where $C > 0$ is a constant independent of G and S.*

From the point of view of consensus, it means that the convergence rate degrades as N grows, at a polynomial speed. Note that there are families of graphs which do not show this dependence. Indeed, the dependence on N can be logarithmical, as for the hypercube graphs (2.21), or there can be no such dependence. Indeed, there are sequences of graphs such that the matrices adapted to them have the essential spectral radius bounded away from 1, uniformly in N. The easiest example are the complete graphs, while other examples are the so called *expander* graphs. A fact similar to Lemma 1.6 is conjectured to hold for many classes of graphs, like (random) geometric graphs and planar grids or tessellations. A complete understanding of which classes of graphs show such dependence of N is still an open problem, which is not going to be further addressed here.

1.3.3 Gossip algorithms

Besides the static linear controller considered in the previous sections, more refined algorithms have been proposed in the literature. Restricting to linear algorithms, interesting cases are obtained using at each time step only a subset of the links. Sufficient and necessary conditions for convergence are known, both if the choices of the subsets are deterministic [65], and if they are stochastic [36].

Particularly interesting is the so called (symmetric) gossip algorithm: at every time instant a randomly chosen pair of agents communicates, and they average their states. Such algorithm, brought to wide audience by [9], has many appealing features: it reduces the number of communications with respect to deterministic algorithms and avoids data collision. Let us describe such algorithm in more detail. Assume we are given an undirected graph $\mathcal{G} = (V, \mathcal{E})$, $\mathcal{E} \subset \{(i,j) : i,j \in V\}$. At each time step, one edge (i,j) is randomly selected in $\mathcal{E}$ with probability $W^{(i,j)}$ such that $\sum_{(i,j)\in\mathcal{E}} W^{(i,j)} = 1$. Let W be the matrix with entries $W_{ij} = W^{(i,j)}$. The two agents connected by the selected edge average their states according to

$$\begin{aligned} x_i(t+1) &= \frac{1}{2}x_i(t) + \frac{1}{2}x_j(t) \\ x_j(t+1) &= \frac{1}{2}x_j(t) + \frac{1}{2}x_i(t) \end{aligned} \tag{1.7}$$

while

$$x_h(t+1) = x_h(t) \qquad \text{if } h \neq i,j. \tag{1.8}$$

Let $E_{ij} = (e_i - e_j)(e_i - e_j)^*$ and

$$P(t) = I - \frac{1}{2}E_{ij} \tag{1.9}$$

where $e_i = [0,\dots,0,1,0,\dots,0]^*$ is a $N \times 1$ unit vector with the i-th component equal to 1, then (1.7) and (1.8) can be written in a vector form as

$$x(t+1) = P(t)x(t). \tag{1.10}$$

Note that $P(t)$ is a symmetric doubly stochastic matrix, and then (1.10) preserves the average of states. It is known [36] that, if the graph $\mathcal{G}$ is connected and each edge $(i,j) \in \mathcal{E}$ can be selected with a strictly positive probability $W^{(i,j)}$, then (1.10) reaches, almost surely, the *average consensus*, that is

$$\lim_{t\to\infty} x(t) = x_{\mathrm{ave}}\mathbf{1},$$

for every initial condition $x(0)$.

1.4 Networks of digital links

Most of the literature on the consensus problem assumes that the communication channels between the nodes allows to communicate real numbers with no errors. In practical applications this can be an unrealistic assumption: the communication is in most cases digital and the channel is noisy. There can be strict bandwidth constraints, as well as communication delays, and packet losses due to interferences and erasures. These limitations are mostly likely to be significant for a network communicating in a wireless fashion, which is nowadays a typical choice.

In the thesis, we are going to concentrate on modeling the links as instantaneous and lossless digital channels. This choice forces a quantization on the real numbers that agents have to transmit, but keeps aside the issues of delays and packet losses.

1.4.1 Quantizers

The first design issue is how to quantize the states, that is how to map the continuous space of states into a discrete alphabet of messages. Given $X \subseteq \mathbb{R}^\nu$, we call quantizer a map $q : X \to S$, where S is a finite or countable set. If we have a vector $x \in X^N$, with a slight abuse of notation, we will use the notation $q(x) \in S^N$ to denote the vector such that $q(x)_i = q(x_i)$. Many quantizers have been proposed in the vast literature on the subject [68]. In this thesis, we concentrate on *uniform* quantizers, that is quantizers such that the measure of $q^{-1}(s)$ is equal for all $s \in S$.

Deterministic quantizer

The first quantizer we consider is the *deterministic quantizer* with countable alphabet. Let $q_d : \mathbb{R} \to \mathbb{Z}$ be the mapping sending $z \in \mathbb{R}$ into its nearest integer, namely,

$$q_d(z) = n \in \mathbb{Z} \quad \Leftrightarrow \quad \begin{array}{ll} z \in [n - 1/2, n + 1/2[, & \text{if } z \geq 0 \\ z \in]n - 1/2, n + 1/2], & \text{if } z < 0. \end{array} \tag{1.11}$$

It enjoys the property that, for all $z \in \mathbb{R}$,

$$|z - q_d(z)| \leq \frac{1}{2}.$$

Most results in Chapters 2 and 3 are about systems of agents which communicate using q_d.

Truncation quantizer

Similar to the latter is the *truncation quantizer*

$$q_T(z) = \lfloor z \rfloor. \tag{1.12}$$

The analysis of systems based on q_T is often similar to q_d, taking into account that the quantization error can be bounded as

$$|z - q_d(z)| \leq 1.$$

For this reason, the quantizer q_T will be treated briefly in the sequel.

Probabilistic quantizer

The *probabilistic quantizer* $q_p : \mathbb{R} \to \mathbb{Z}$ is defined as follows. For any $x \in \mathbb{R}$, $q_p(x)$ is a random variable on $\mathbb{Z}$ defined by

$$q_p(x) = \begin{cases} \lfloor x \rfloor & \text{with probability} \quad \lceil x \rceil - x \\ \lceil x \rceil & \text{with probability} \quad x - \lfloor x \rfloor. \end{cases} \tag{1.13}$$

The following lemma discusses two important properties of the probabilistic quantizer.

Lemma 1.7 *For every $x \in \mathbb{R}$, it holds that*

$$\mathbb{E}\left[q_p(x)\right] = x. \tag{1.14}$$

Moreover

$$\mathbb{E}\left[(x - q_p(x))^2\right] \leq \frac{1}{4}. \tag{1.15}$$

Proof: Equation (1.14) is immediate and (1.15) follows from computing

$$\mathbb{E}\left[(x - q_p(x))^2\right] = x\lfloor x \rfloor + x\lceil x \rceil - \lfloor x \rfloor \lceil x \rceil - x^2 \leq \frac{1}{4}.$$

□

The quantizers defined above map $\mathbb{R}$ into $\mathbb{Z}$ and have quantization bins of length 1. More general quantizers, having as quantization step a positive real number ε, can be obtained from $q : \mathbb{R} \to \mathbb{Z}$ by defining $q^{(\varepsilon)}(x) = \varepsilon q(x/\varepsilon)$. Hence, the general case can be simply recovered by a suitable scaling, and our choice is not restrictive.

1.4.2 Consensus on non ideal networks

We can make the framework (1.4) more general if we assume that the agents can not directly access their neighbors' states, but can obtain estimates or approximations of them. Let $\hat{x}_j^i(t)$ be the estimate that agent i has of the state of agent j, at time t. The control input $u_i(t)$ assumes the following form

$$u_i(t) = \sum_{j=1}^{N} K_{ij} \hat{x}_j^i(t). \tag{1.16}$$

Due to the possibility that $\hat{x}_j^i(t) \neq x_j(t)$, such a control is not in general guaranteed to yield convergence to consensus, nor to preserve the average of states. This general framework is suitable for several specialization.

(i) If the agents communicate via a reliable instantaneous analog channel, we can just assume that $\hat{x}_j^i(t) = x_j(t)$, and thus the algorithm reduces to (1.4).

(ii) If the agents communicate via a digital channel, we can set $\hat{x}^i_j(t) = q(x_j(t))$. Then

$$x(t+1) = x(t) + Kq(x(t)). \tag{1.17}$$

This update rule appeared for the first time in [18] and is studied deeply in Chapter 2. It shares with the linear ideal case (1.4) the important property of preserving the average of states, in spite of the additional constraint of quantization.

Proposition 1.8 (Average preservation) *Let $x(t)$ evolve following* (1.17) *and $P = I + K$ satisfy Assumption 1.1. Then $x_{\mathrm{ave}}(t) = x_{\mathrm{ave}}(0)$ for all $t \in \mathbb{N}$.*

Proof: Since P is doubly stochastic, we have that

$$\begin{aligned} x_{\mathrm{ave}}(t+1) &= N^{-1}\mathbf{1}^* x(t+1) \\ &= N^{-1}\mathbf{1}^* x(t) \; + N^{-1}\mathbf{1}^* \left(P - I\right) q(x(t)) \\ &= N^{-1}\mathbf{1}^* x(t) = x_{\mathrm{ave}}(t). \end{aligned}$$

□

(iii) An important case is when the states are communicated with an additive noise, only depending on the source agent j: thus $\hat{x}^i_j(t) = x_j(t) + n_j(t)$ for all i neighbor to j. This case of *broadcasting noise* has been studied in [88, 18, 80]. The main application of such additive noise is modeling the effects of quantization, namely, the quantization error $x - q(x)$. If $q = q_p$, then the error $x - q_p(x)$ naturally enjoys statistical properties, expressed in Lemma 1.7. If instead $q = q_d$, no statistical property holds for the error, and the bound $|x - q_d(x)| \leq \frac{1}{2}$ is the only a priori information on it. However, for the analysis purposes, it can be sensible to *assume* such error to be a random variable which has zero mean, and the interval $[-1/2, 1/2]$ as support: then its variance is bounded by $1/4$. Provided suitable independence assumptions, the *same probabilistic model*, which (roughly speaking) consists in assuming $q(x) = x+n$, can be used for the analysis of both the probabilistic and the deterministic quantizer. However, the significance of the model is different in the two cases. For the probabilistic quantizer, the model is a priori correct: it forgets the exact definition of the quantizer, focusing on its statistics. Instead, for the deterministic quantizer there is no a priori guarantee of correctness, since the model assumes a deterministic quantity to be random: the validation of the model can only come from an a posteriori comparison with experimental results. Such a twofold analysis is going to be carried on in Chapters 2 and 3, and has appeared elsewhere in the literature, with various levels of consciousness [2, 54, 80].

(iv) If the agents communicate via an analog channel, subject to an *additive channel noise*, we can set $\hat{x}^i_j(t) = x_j(t) + n^i_j(t)$, where $n^i_j(t)$ represents the added noise. $n^i_j(t)$ are assumed to be independent random variables with zero mean. This setting is most natural whenever the neighbors states are *sensed*, rather than communicated. For this reason, it has been studied in [80], and in [18]. It is also attracting attention for the analysis of digital communication over noisy channels [17].

Remark 1.1 (Naive quantization) In a naive approach, one which is not guided by the control theoretical framework (1.16), could suggest that the agents apply to received data the same linear algorithm as in the ideal case, so that the states evolve following

$$x_i(t+1) = P_{ii}\, x_i(t) + \sum_{j \neq i} P_{ij}\, q(x_j(t)). \tag{1.18}$$

Differently from (1.16), the system (1.18) does not guarantee that $x_{\text{ave}}(t+1) = x_{\text{ave}}(t)$, that is the average of the states is not preserved. This drawback, which has been remarked for the first time in [89], is due to the loss of symmetry between neighbors in the use of information. Moreover, the dynamical system (1.18), besides not preserving the average of states, shows in simulations poor convergence properties [18]. For these reasons, we do not discuss it any longer.

1.4.3 Related works

The effects of quantization in feedback control problems have been widely studied in the past mainly in the stabilization problem. We refer to the survey [68] for an introduction.

Granularity effects different from quantized communication have been tackled in few papers on the consensus problems [35], [55]. In [55], the authors study systems having (and transmitting) integer-valued states and propose a class of gossip algorithms which preserve the average of states and are guaranteed to converge up to one quantization bin. Similar problems have been studied by computer scientists, in the context of load balancing. Our setting is different: we study the real-valued consensus problem under the assumption that communications are quantized, and we assume an ideal noiseless digital channel. This problem has first been posed in the final section of [89]. While the work of this thesis was in progress, some papers have considered consensus with quantized communication [69], [3, 2], [54]. The recently appeared technical report [69] studies a deterministic quantization scheme, not preserving the average, in which the agents can only store quantized data. Several results are obtained on the error between the convergence value and the initial average, and on the speed of convergence. The papers [3, 2] actually deal with quantized communication, using the probabilistic quantizer. Their scheme achieves almost surely a consensus at a quantization level, but the average is preserved in expectation only. The probabilistic quantizer also appears in [87], and, with an equivalent definition, in [54], where it is used with time dependent consensus gains.

Chapter 2

Static quantization and fixed topology

This chapter contains an extended analysis of the average consensus algorithm appeared in [18, 20, 21] and in [41], which features a uniform quantizer to encode the messages among agents, and a time invariant communication graph. Thanks to a suitably designed update, this algorithm is able to overcome the drawbacks of a naive treatment of quantization, and to guarantee that the agents states get eventually close to the desired average, at a speed which is roughly the same as in the ideal case. Since, as we shall see, the algorithm preserves the average of the initial condition, its performance is defined in terms of the distance of the states from the average. A special attention is given to the scalability in the number of agents N of such performance, and this crucial issue is studied in the context special sequences of graphs.

2.1 The algorithm

We have seen in Chapter 1 that, assuming ideal communication between agents connected by a directed graph, the average consensus is solved by the algorithm

$$x_i(t+1) = \sum_{j=1}^{N} P_{ij} x_j(t),$$

where $x_i(t) \in \mathbb{R}$ is the state of the i-th agent at the time t, provided that Assumption 1.1 holds. More compactly we can write

$$x(t+1) = Px(t), \tag{2.1}$$

where $x(t)$ is the column vector whose entries $x_i(t)$ represent the agents states.

In presence of quantized communication, following the framework (1.16), we propose the algorithm

$$x(t+1) = x(t) \; + \; (P - I)\, q_d(x(t)). \tag{2.2}$$

or a corresponding algorithm with the probabilistic quantizer

$$x(t+1) = x(t) + (P - I)\, q_p(x(t)). \tag{2.3}$$

As stated in Proposition 1.8, algorithms (2.2) and (2.3) preserve the average of the initial conditions, that is $x_{\text{ave}}(t) = x_{\text{ave}}(0)$ for all $t \geq 0$.

The algorithms (2.2) and (2.3), because of the quantization effects, are not expected to converge in the sense (1.3). What we can hope is for the agents to reach states which are close to the average $x_{\text{ave}}(0)$. To measure this asymptotic disagreement, we introduce the performance index

$$d_\infty(P, x(0)) = \limsup_{t\to\infty} \frac{1}{\sqrt{N}} ||y(t)||,$$

where $y(t) = \Omega x(t)$, as in Definition 1.2. We can get rid of the initial condition by considering

$$d_\infty(P) = \sup_{x(0)} d(P, x(0)). \tag{2.4}$$

Clearly $d_\infty(P)$ depends on the matrix P and, via P, on the communication topology. The *analysis* problems we would like to address are:

(i) given a matrix P, evaluate how big $d_\infty(P)$ is;

(ii) given a sequence of matrices P_N, adapted to graphs of increasing size N, estimate how $d_\infty(P_N)$ depends on N.

Moreover, we are interested in the following *design* problem:

> given a communication topology, find the matrix P, adapted to the topology, which minimizes $d_\infty(P)$.

With respect to the linear iteration (2.1), the quantization effects introduce nonlinearities which make the states evolution hard to analyze. Indeed, we are able to carry on such an analysis only in very specific examples, in Section 2.1.1. To overcome these limitations, instead of studying the exact evolution of the algorithm, we model the effect of quantization as disturbances of the linear system (2.1). This leads to a worst-case analysis and a probabilistic model. The worst-case analysis is obtained by a bounded error model, maximizing $d_\infty(P)$ over all possible realizations of the communication errors. This is done in Section 2.2 where we prove in general the *convergence to a neighborhood of the average*, obtaining upper bounds on its size. These bounds are *independent of the initial condition* but depend on the diffusion matrix. The worst-case analysis is significant for both q_d and q_p, but its results are intrinsically conservative: we are going to show that, on a sequence of matrices P_N adapted to hypercube graphs of order N, $d_\infty(P_N)$ grows as the logarithm of N. Such behavior is in disagreement with the experimental evidence displayed in Section 2.4.

For this reason, in Section 2.3 we propose an alternative method and develop a probabilistic analysis, modeling the quantization error as additive random noise affecting the received data. As discussed in Section 1.4.2, the significance of the probabilistic analysis is a priori guaranteed when the quantizer map is q_p. Interestingly, simulations show that the probabilistic analysis is very close to the experimental evidence for the deterministic map as well.

A classical mean squared analysis for the asymptotic error can be carried on in the probabilistic model. It comes out that, under mild assumptions, the expected behavior depends only on the assumed distribution of the errors and on the spectrum of P, and does not depend on N. We shall comment on this discrepancy later in the chapter.

2.1.1 Deterministic analysis

We start with a remark about the best performance achievable by (2.2). It is clear that, when all states lie in the same quantization interval, namely $q_d(x_i(T)) = Q$ for all i, differences are not perceivable and states do not evolve. Therefore the best the algorithm can assure is that the system reaches such an equilibrium in which $q_d(x_i(t)) = Q$ for all i and for all $t \geq T$. This implies both that $|y_i(t)| \leq 1$ for all i and for all $t \geq T$, and that $d_\infty \leq 1/2$. Obtaining such a performance is not trivial: simulations show that the error from the agreement can actually be bigger. However, it is worth noting that the states of the agents subject to (2.2) are bounded. In particular one can see, by convexity arguments, that for any node $i \in V$ and for all $t \geq 0$,

$$x_i(t) \in \left[\min_{j\in V}\{q_d(x_j(0))\} - \frac{1}{2},\ \max_{j\in V}\{q_d(x_j(0))\} + \frac{1}{2}\right]. \tag{2.5}$$

This simple result is of some applicative interest from the point of view of quantizer's design. Indeed it implies that, although the quantizer has in principle a countably infinite range space, only a *finite number of symbols* is going to be actually used by the algorithm, and such number has an upper bound which depends on the initial condition only.

On the other hand, (2.5) is a very weak result from the point of view of convergence, since we can hope that in general the disagreement decreases as time goes on and, moreover, that the asymptotic disagreement does not depend on the initial conditions, but just possibly on P. We consider now two examples in which the evolution of the system can be studied explicitly, the complete graph and the directed ring graph.

Example 2.1 (Complete graph) If the communication graph is complete and the communication is not quantized, the average consensus problem can be solved in one step taking $P = \frac{1}{N}\mathbf{1}\mathbf{1}^*$. We now compute the performance degradation due to quantization. The system is in this case

$$x(t+1) = x(t) - q_d(x(t)) + N^{-1}\mathbf{1}\mathbf{1}^* q_d(x(t)). \tag{2.6}$$

We have that, for $t \geq 1$,

$$\begin{aligned}
|y_i(t)| &= \left| x_i(t) - \frac{1}{N}\sum_{j=1}^{N} x_j(0) \right| = \left| x_i(t) - \frac{1}{N}\sum_{j=1}^{N} x_j(t-1) \right| \\
&= \left| x_i(t-1) - q\left(x_i(t-1)\right) + \frac{1}{N}\sum_{j=1}^{N} q\left(x_j(t-1) - \frac{1}{N}\sum_{j=1}^{N} x_j(t-1) \right| \\
&\leq \left| x_i(t-1) - q\left(x_i(t-1)\right)\right| + \frac{1}{N}\sum_{j=1}^{N} \left| q\left(x_j(t-1) - x_j(t-1)\right| \\
&= \frac{1}{2} + \frac{1}{N}\frac{N}{2} = 1.
\end{aligned}$$

This implies that

$$d_\infty(N^{-1}\mathbf{1}\mathbf{1}^*) \leq 1. \tag{2.7}$$

Example 2.2 (Directed ring) Now we consider a more interesting example, the directed ring graph of Example 1.3 with a uniform Perron matrix: each agent communicates with only one neighbor and evolves following

$$x_i(t+1) = x_i(t) + \frac{1}{2}[q_d(x_{i+1}(t)) - q_d(x_i(t))] \quad i = 1, \ldots, N, \tag{2.8}$$

where summation of the indexes is to be intended $\mod N$. The evolution of (2.8) can be studied exactly by means of a *symbolic dynamics* approach[1]. This analysis is definitively not trivial, but permits us to characterize precisely the evolution of (2.8) and to obtain a strong result. Indeed, we show in Appendix 2.A that there exists $T \in \mathbb{N}$ such that

$$|x_i(t) - x_j(t)| \leq 1 \quad \forall\, i, j \quad \forall\, t > T.$$

It follows that

$$d_\infty(P) \leq 1/2.$$

2.2 Bounded error model

An exact analysis of the dynamics of system (2.2), as we did in the previous two examples, is not feasible for general graphs. In this section we undertake a worst-case analysis which can instead be applied to the general case. We start by observing that (2.2) can be rewritten in the following way

$$x(t+1) = Px(t) + (P - I)(q_d(x(t)) - x(t)). \tag{2.9}$$

[1] A similar analysis can be done for q_T, leading to analogous results, whereas a similar symbolic dynamics can not be defined for q_p.

Notice that $\|q_d(x(t)) - x(t)\|_\infty \leq 1/2$. In order to carry out a worst-case analysis of (2.9), we introduce the following bounded error model [2]

$$\begin{cases} x_w(t+1) = Px_w(t) + (P-I)e(t), & x_w(0) = x(0) \\ y_w(t) = \Omega x_w(t) \end{cases} \tag{2.10}$$

where $e(t) \in \mathbb{R}^N$ is such that $\|e(t)\|_\infty \leq 1/2$ for all $t \geq 0$. Notice that in this case $e(t)$ is no more a quantization error, but instead represents an unknown bounded disturbance. Clearly, when $e(t) = q_d(x(t)) - x(t)$ it turns out that $x_w(t) = x(t)$ and $y_w(t) = y(t)$ for all $t \geq 0$.

We define now a performance index for (2.10), considering the worst asymptotic disagreement, worst with respect to all the possible choices of the time sequence of the vectors $e(t)$. To be more precise, let us introduce

$$\mathcal{E}^\infty = \{\{e(\cdot)\}_{t=0}^\infty \mid \|e(t)\|_\infty \leq \frac{1}{2},\ \forall\, t \geq 0\}, \tag{2.11}$$

namely the set of all the sequences of N-dimensional vectors having sup norm less than 1/2. Then, for the system (2.10), we define

$$d_\infty^w(P, x_w(0)) = \sup_{\mathcal{E}^\infty} \limsup_{t\to\infty} \frac{1}{\sqrt{N}} ||y_w(t)||. \tag{2.12}$$

Note that $\lim_{t\to\infty} \Omega P^t = 0$. This implies that the asymptotic behavior of $y_w(t)$ is independent of the initial condition $x_w(0)$ and hence this is the case also for the quantity $d_\infty^w(P, x_w(0))$. Thus, from now on we denote $d_\infty^w(P, x_w(0))$ simply by $d_\infty^w(P)$. As a preliminary remark, note that

$$d_\infty(P) \leq d_\infty^w(P).$$

The following result provides a general bound for d_∞^w.

Proposition 2.1 *Let P be a matrix satisfying Assumption 1.1. Then*

$$d_\infty^w(P) \leq \frac{1}{1 - \|P\Omega\|}. \tag{2.13}$$

Proof: Consider now $y_w(t)$. Standard algebraic tools, and Lemma 1.2, yield

$$\begin{aligned} y_w(t) &= \Omega P^t x(0) + \Omega \sum_{s=0}^{t-1} P^s (I-P) e(t-s-1) \\ &= (P\Omega)^t y(0) + \sum_{s=0}^{t-1} (P\Omega)^s (I-P) e(t-s-1). \end{aligned}$$

[2] A similar bounded error model can be set for the system (2.3) as well, taking $\|e(t)\|_\infty \leq 1$. All results in the sequel follows with obvious changes. However, since for (2.3) the probabilistic model (2.18) has a stronger significance, we do not further detail its worst-case analysis.

Now we have that

$$
\begin{aligned}
\|y_w(t)\| &= \|(P\Omega)^t y_w(0) + \sum_{s=0}^{t-1} (P\Omega)^s (I-P) e(t-s-1)\| \\
&\leq \|(P\Omega)^t\| \, \|y_w(0)\| + \|I-P\| \sum_{s=0}^{t-1} \|(P\Omega)\|^s \|e(t-s-1)\| \\
&= \|(P\Omega)^t\| \, \|y_w(0)\| + \sqrt{N} \frac{1-\|P\Omega\|^t}{1-\|P\Omega\|}
\end{aligned}
$$

where in the last inequality we used the facts that $\|I-P\| \leq 2$ and $\|e(t)\| \leq \sqrt{N}/2$ for all $t \geq 0$. By letting $t \to \infty$ we obtain (2.13). □

Note that, if P is normal, we have that $\|P\Omega\| = \rho_{\text{ess}}(P)$ and hence (2.13) becomes

$$
d^w(P) \leq \frac{1}{1-\rho_{\text{ess}}(P)}.
$$

However, when P is a normal matrix the bound on $d^w_\infty(P)$ can be improved as stated in the next proposition.

Proposition 2.2 *If P is a normal matrix satisfying Assumption 1.1, then*

$$
d^w_\infty(P) \leq \frac{1}{2} \sum_{s=0}^{\infty} \rho(P^s(I-P)). \tag{2.14}
$$

Proof: Starting from the expression of $y_w(t)$ provided along the proof of Proposition 2.1 we can write that

$$
\begin{aligned}
\|y_w(t)\| &\leq \|(P\Omega)^t y_w(0)\| + \|\sum_{s=0}^{t-1} (P\Omega)^s (I-P) e(t-s-1)\| \\
&\leq \|(P\Omega)^t y_w(0)\| + \frac{\sqrt{N}}{2} \|(P\Omega)^s (I-P)\|
\end{aligned}
$$

Since P is normal we have that $\|(P\Omega)^s(I-P)\| = \rho((P\Omega)^s(I-P)) = \rho(P^s(I-P))$. By letting $t \to \infty$ in the last inequality, we obtain (2.14). □
Note that, from the sub-multiplicative inequality $\|(P\Omega)^s(I-P)\| \leq \|P\Omega\|^s \, \|I-P\|$, it follows immediately that $\frac{1}{2}\sum_{s=0}^{\infty} \rho(P^s(I-P)) \leq \frac{1}{1-\rho_{\text{ess}}(P)}$. Then the bound (2.14) is indeed an improvement of the bound (2.13).

Example 2.3 (Complete graph) Let $P = \frac{1}{N}\mathbf{1}\mathbf{1}^*$. Hence $\frac{1}{1-\|P\Omega\|} = 1$. This is an alternative way to prove (2.7). However, since $\left(\frac{1}{N}\mathbf{1}\mathbf{1}^*\right)^k = \frac{1}{N}\mathbf{1}\mathbf{1}^*$ for all $k > 0$, it follows immediately that $\frac{1}{2}\sum_{s=0}^{\infty} \rho(P^s(I-P)) = \frac{1}{2}$ and this represents a refinement of (2.7). We conclude that $d_\infty(\frac{1}{N}\mathbf{1}\mathbf{1}^*) \leq \frac{1}{2}$.

To further improve the bound (2.14) requires some more work. Given $c \in \mathbb{C}$ and $r \in \mathbb{R}$ such that $r \geq 0$, we denote by

$$B_{c,r} = \{z \in \mathbb{C} \,|\, \|z - c\| \leq r\}$$

the closed ball of complex numbers of radius r and centered in c. The proof of the following result is deferred to the Appendix 2.B.

Proposition 2.3 *Let P be a normal matrix satisfying the Assumption 1.1. Let R be such that $0 < R < 1$ and $\sigma(P) \subseteq B_{1-R,R}$ and let $\bar{\rho} = \rho_{\mathrm{ess}}(P)$ denote the essential spectral radius of P. Then*

$$\sum_{s=0}^{\infty} \rho(P^s(I-P)) \leq \frac{1}{1-R} + \sqrt{\frac{8R}{(1-R)(1-\bar{\rho})}}. \tag{2.15}$$

Remark that the given bound depends on R and $\bar{\rho}$ only, which are functions of the spectrum of P. Let us illustrate it with an example.

Example 2.4 (Directed ring) Consider the directed ring graph introduced in Example 2.2 and the evolution law given by (2.8). In this case P is the circulant matrix

$$P = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & \frac{1}{2} & \frac{1}{2} & 0 & \cdots & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{2} & \frac{1}{2} & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \cdots & \vdots & \vdots & \vdots \\ \frac{1}{2} & 0 & 0 & 0 & \cdots & 0 & 0 & \frac{1}{2} \end{pmatrix}.$$

We know from Corollary 2.15 that $d_\infty(P) \leq \frac{1}{2}$. Let us compare this fact with the bound in Proposition 2.1. Since $\rho_{\mathrm{ess}}(P) = 1 - \frac{\pi^2}{2}\frac{1}{N^2} + o\left(\frac{1}{N^2}\right)$ and since each circulant matrix is a normal matrix we have that $\frac{1}{1-\|P\Omega\|} = \frac{1}{1-\rho_{\mathrm{ess}}(P)} = \Theta\left(N^2\right)$. Observe now that all the eigenvalues of P are inside the ball $B_{\frac{1}{2},\frac{1}{2}}$. Hence we obtain that $\frac{1}{1-R}+\sqrt{\frac{8R}{(1-R)(1-\bar{\rho})}} = \Theta(N)$. This means that the bound (2.15) improves the bound proposed in (2.13). Moreover, by numerical experiments one can check that $\frac{1}{2}\sum_{s=0}^{\infty} \rho(P^s(I-P)) = \Theta(N)$: the bound (2.15) is tight.

If P is symmetric we can provide a stronger result.

Proposition 2.4 *Let P be a symmetric stochastic matrix satisfying Assumption 1.1. Let R be such that $0 < R < 1$ and $\sigma(P) \subseteq B_{1-R,R}$ and let $\bar{\rho} = \rho_{\mathrm{ess}}(P)$ denote the essential spectral radius of P. Then,*

$$\sum_{s=0}^{+\infty} \rho\left(P^s(I-P)\right) \leq \frac{3}{2} + \frac{1}{1-R} + \frac{1}{2}\log\left(\frac{1}{1-\bar{\rho}}\right). \tag{2.16}$$

Proof: Assume that $\sigma(P) = \{\lambda_0 = 1, \lambda_1, \dots, \lambda_{N-1}\}$. Note that $(\sigma(P) \setminus \{1\}) \subseteq [1-2R, \bar\rho]$. We want to upper bound $\rho(P^s(I-P)) = \max_{k=1}^{N-1} |\lambda_k^s(1-\lambda_k)|$. To do this, consider the function $f(x) = |x^s(1-x)|$. It is continuous in $[-1,1]$, positive and decreasing in $[-1,0]$, it vanishes in $x=0$ and in $x=1$ and has a local maximum in $x = x_M = \frac{s}{1+s}$, with $f(x_M) = (\frac{s}{1+s})^s(\frac{1}{1+s})$. We need to evaluate $\max_{x\in[1-2R,\bar\rho]} f(x)$.

First observe that there exists $\bar s$, only depending on the value $1-2R$, such that for all $s > \bar s$ we have $f(1-2R) < (\frac{s}{1+s})^s(\frac{1}{1+s})$ and then the global maximum of $f(x)$ is assumed at $x = x_M$ if $x_M \le \bar\rho$. Since x_M tends to 1 as s goes to infinity, it happens that, when $s > \frac{\bar\rho}{1-\bar\rho}$, $x_M(s) > \bar\rho$. In conclusion we have that

$$\max_{x\in[1-2R,\bar\rho]} f(x) = \begin{cases} 2R|1-2R|^s, & \text{if } 0 \le s \le \bar s; \\ (\frac{s}{1+s})^s(\frac{1}{1+s}), & \text{if } \bar s < s \le s^*; \\ \bar\rho^s(1-\bar\rho), & \text{if } s^* < s < \infty. \end{cases}$$

where $s^* = \lfloor \frac{\bar\rho}{1-\bar\rho} \rfloor$. Hence we can write

$$\sum_{s=0}^{t-1} \rho(P^s(I-P))$$
$$\le \sum_{s=0}^{\bar s} 2R|1-2R|^s + \sum_{s=\bar s+1}^{s^*} \left(\frac{s}{1+s}\right)^s \left(\frac{1}{1+s}\right) + \sum_{s=s^*+1}^{t-1} \bar\rho^s(1-\bar\rho).$$

Notice now that

$$\sum_{s=0}^{\bar s} 2R|1-2R|^s \le \frac{1}{1-R},$$

and that

$$\sum_{s=s^*+1}^{t-1} \bar\rho^s(1-\bar\rho) = \bar\rho^{s^*+1} - \bar\rho^t \le 1.$$

Notice finally that, since $\sum_{i=1}^m 1/i \le 1 + \ln m$ we can argue that

$$\begin{aligned} \sum_{s=\bar s+1}^{s^*} \left(\frac{s}{1+s}\right)^s \left(\frac{1}{1+s}\right) &\le \sum_{s=0}^{s^*} \left(\frac{1}{2}\right)\left(\frac{1}{1+s}\right) \le \frac{1}{2}(1+\log(s^*+1)) \\ &\le \frac{1}{2} + \frac{1}{2}\log\left(\frac{\bar\rho}{1-\bar\rho}+1\right) = \frac{1}{2} + \frac{1}{2}\log\left(\frac{1}{1-\bar\rho}\right). \end{aligned}$$

Putting together these three inequalities we obtain (2.16). □

The significance of the result can be seen with an example.

Example 2.5 (Undirected ring) In this example we consider the circulant matrix

$$P = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & 0 & 0 & \cdots & 0 & 0 & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & 0 & \cdots & 0 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \cdots & \vdots & \vdots & \vdots \\ \frac{1}{3} & 0 & 0 & 0 & \cdots & 0 & \frac{1}{3} & \frac{1}{3} \end{pmatrix}.$$

Simulations, shown in Figure 2.3, suggest that $d_\infty(P)$ be uniformly bounded on N. Consider the bound (2.13): since $\rho_{\text{ess}}(P) \asymp 1 - \frac{4}{3}\frac{\pi^2}{N^2}$ we obtain that $\frac{1}{1-\|P\Omega\|} = \frac{1}{1-\rho_{\text{ess}}(P)} = \Theta(N^2)$. On the other hand, observe that all the eigenvalues of P are greater than $-\frac{1}{3}$. Hence $R = 2/3$, and then $\frac{3}{2} + \frac{1}{1-R} + \frac{1}{2}\log\left(\frac{1}{1-\rho_{\text{ess}}(P)}\right) = \Theta(\log N)$. Moreover, numerically it is possible to observe that also $\frac{1}{2}\sum_{s=0}^{\infty}\rho(P^s(I-P))$ grows logarithmically, meaning that asymptotically in N, (2.16) is tight.

The bound (2.14) is in many cases logarithmical in N. Let us consider a sequence of Cayley Perron matrices P_N with their diagonal elements bounded away from zero, uniformly in N, and adapted to a Cayley graphs of constant degree. Then, from Lemma 1.6, we argue that

$$\frac{1}{1-R} + \log\sqrt{\frac{1}{1-\rho_{\text{ess}}(P_N)}} = \Omega(\log N).$$

We can not argue from this fact the stronger claim that $d_\infty^w(P_N) = \Omega(\log N)$, but we can provide another sequence enjoying such property.

Theorem 2.5 *Let there be a sequence of hypercube graphs of size $N = 2^n$, and a sequence of uniform Perron matrices P_N adapted to them. Then*

$$d_\infty^w(P_N) = \frac{n}{2} = \frac{\log_2 N}{2}.$$

Consider an hypercube graph and the uniform Perron matrix adapted to it. First we give the following preliminary result.

Lemma 2.6 *Let P be a uniform Perron matrix adapted to an hypercube graph. Then*

$$\sum_{s=0}^{\infty}\rho(P^s(I-P)) = n = \log_2 N. \tag{2.17}$$

Proof: From [31], the eigenvalues of P are, for $k \in \{0,\dots,n\}$, $\lambda_k = 1 - \frac{2k}{n+1}$, with multiplicity $p_k = \binom{n}{k}$. Then,

$$\begin{aligned}\sum_{s=0}^{\infty}\rho(P^s(I-P)) &= \sum_{s=0}^{\infty}\rho_{\text{ess}}(P^s)\rho(I-P)\\ &= \sum_{s=0}^{\infty}\left(1-\frac{2}{n+1}\right)^s\left(2-\frac{2}{n+1}\right) = n.\end{aligned}$$

□

We are able now to provide the proof of Theorem 2.5.
Proof: First we rewrite the expression of $d_\infty^w(P)$. Since P is symmetric, it is diagonalizable by an orthogonal matrix, and we can write that $P = \sum_{h=0}^{N-1}\lambda_h q_h q_h^*$ where

q_h are orthonormal. These facts are true also for $P^s(I-P)$. Moreover we have that $\rho\left(P^s(I-P)\right) = ||P^s(I-P)||$. We define

$$y_w^{(f)}(t) = \sum_{s=0}^{t-1} P^s(I-P)e(t-s-1),$$

and we know from the proof of Proposition 2.1 that

$$\limsup_{t\to\infty} \|y_w(t)\| = \limsup_{t\to\infty} \|y_w^{(f)}(t)\|.$$

Then,

$$\begin{aligned}
\|y_w^{(f)}(t)\|^2 &= \|\sum_{s=0}^{t-1} P^s(I-P)e(t-s-1)\|^2 \\
&= \|\sum_{s=0}^{t-1}\sum_{h=0}^{N-1} \lambda_h^s(1-\lambda_h)q_h q_h^* e(t-s-1)\|^2 \\
&= \|\sum_{h=0}^{N-1} q_h(1-\lambda_h)\sum_{s=0}^{t-1}\lambda_h^s q_h^* e(t-s-1)\|^2 \\
&= \sum_{h=0}^{N-1}\left[(1-\lambda_h)\sum_{s=0}^{t-1}\lambda_h^s q_h^* e(t-s-1)\right]^2.
\end{aligned}$$

From the proof of Proposition 2.1, $[d_\infty^w(P)]^2 = \max_{\mathcal{E}^\infty} \limsup_{t\to\infty} \frac{1}{N}\|y_w^{(f)}(t)\|^2$.

Now we start using combinatorial tools. Indeed the vertices of the hypercube, as well as the eigenvalues and eigenvectors of P, can be indexed by the subsets of $\{1\ldots,n\}$. subsets of $\{1\ldots,n\}$.
With this indexing, for each $I\subseteq\{1\ldots n\}$ the corresponding eigenvalue is $\lambda_I = 1-\frac{2|I|}{n+1}$ and the eigenvector is the 2^n-dimensional vector $q^{(I)}$, such that its $J-$th component is equal to $q_J^{(I)} = 2^{-n/2}(-1)^{|I\cap J|}$.
Let $T\in\mathbb{N}$ and consider the sequence of vectors $\{e(t)\}_{t\in\mathbb{Z}_{\geq 0}}$ such that the J-th component of the vector $e(t)$ is equal to $\frac{1}{2}(-1)^{T-1-r}(-1)^{|J|}$, where r is the remainder in the Euclidean division of t over T. Observe that $e(t+T)=e(t)$ for all $t\geq 0$. Observe, moreover, that $e(t)$ is an eigenvector of P corresponding to the eigenvalue $\frac{1-n}{1+n}$ for all $t\geq 0$. Hence we

have that

$$
\begin{aligned}
\frac{1}{N}\|y_w^{(f)}(T)\|^2 &= \frac{1}{N}\sum_{h=0}^{N-1}\left[(1-\lambda_h)\sum_{s=0}^{T-1}\lambda_h^s q_h^* e(T-s-1)\right]^2 \\
&= \frac{1}{2^n}\left[\left(1-\frac{1-n}{n+1}\right)\sum_{s=0}^{T-1}\left(\frac{1-n}{n+1}\right)^s 2^{-\frac{n}{2}}\sum_{J\subseteq\{1,\dots,n\}}(-1)^{|J|}\frac{1}{2}(-1)^s(-1)^{|J|}\right]^2 \\
&= \frac{1}{4^n}\left[\frac{n}{n+1}\sum_{s=0}^{T-1}\left(\frac{n-1}{n+1}\right)^s 2^n\right]^2 \\
&= \frac{n^2}{(n+1)^2}\left[\frac{1-\left(\frac{n-1}{n+1}\right)^T}{1-\frac{n-1}{n+1}}\right]^2 \\
&= \frac{n^2}{4}\left[1-\left(\frac{n-1}{n+1}\right)^T\right]^2.
\end{aligned}
$$

Assume now that T is even. By recalling that $e(t+T)=e(t)$ for all $t\geq 0$, for $t=kT$ where $k\in\mathbb{N}$, it turns out that

$$
\begin{aligned}
\frac{1}{N}\|y_w^{(f)}(kT)\|^2 &= \frac{n^2}{4}\left[1-\left(\frac{n-1}{n+1}\right)^T\right]^2\sum_{u=0}^{k-1}\left(\frac{1-n}{1+n}\right)^{uT} = \\
&= \frac{n^2}{4}\left[1-\left(\frac{n-1}{n+1}\right)^T\right]^2\left[\frac{1-\left(\frac{n-1}{n+1}\right)^{kT}}{1-\left(\frac{n-1}{n+1}\right)^T}\right]^2 \\
&= \frac{n^2}{4}\left[1-\left(\frac{n-1}{n+1}\right)^{kT}\right]^2.
\end{aligned}
$$

Letting $k\to\infty$ we obtain that, for the particular sequence considered,

$$\lim_{k\to\infty}\frac{1}{N}\|y_w^{(f)}(kT)\|^2 = \frac{n^2}{4}$$

Therefore we have proved that

$$\limsup_{t\to\infty}\frac{1}{N}\|y_w^{(f)}(t)\|^2 \geq \frac{n^2}{4}$$

and hence $[d_\infty^w(P)]^2 \geq \frac{n^2}{4}$. Now, Lemma 2.6 implies that $[d_\infty^w(P)]^2 \leq \frac{n^2}{4}$, and then the claim follows. □

Theorem 2.5 is not in accordance with the evidence of simulations in Section 2.4, which do not show any significant dependence of d_∞ on N. This discrepancy confirms the interest for a probabilistic model of the quantization error.

2.3 Probabilistic model

In the previous section we have shown that the bounded error model does not seem to really capture the behavior of the quantized system (2.2) nor of (2.3). In particular the upper bound to the performance we have found seems to be quite conservative. In this section we undertake a probabilistic approach, modeling the quantization error as a random variable. We carry on a classical mean square analysis and we show that it is quite close to simulations of the real quantized system.

For all $i \in V$ and $t \in \mathbb{Z}_{\geq 0}$, let $n_i(t)$ be random variables such that

- $n_i(t)$ are uncorrelated, that is $\mathbb{E}\left[n_i(t)n_j(\tau)\right] = 0$ if $i \neq j$ or $t \neq \tau$
- $n_i(t) \in [-1/2, 1/2]$
- $\mathbb{E}[n_i(t)] = 0$

Remark that $\mathbb{E}[n_i(t)^2] \leq \frac{1}{4}$. Let $n(t)$ be the random vector whose components are $n_i(t)$ and consider the probabilistic model

$$\begin{cases} x_r(t+1) = Px_r(t) + (P-I)n(t), \\ x_r(0) = x(0) \end{cases} \tag{2.18}$$

and the variable $y_r(t) = \Omega x_r(t)$. We define

$$d_\infty^r(P, x_r(0)) = \limsup_{t\to\infty} \sqrt{\frac{1}{N}\mathbb{E}[\|y_r(t)\|^2]}.$$

Like for the model in Section 2.2, we have that $d_\infty^r(P, x_r(0))$ is independent of the initial condition $x_r(0)$. Hence, in the sequel we denote $d_\infty^r(P, x_r(0))$ with the symbol $d_\infty^r(P)$.

The following result is the main one of this section.

Theorem 2.7 *Let P be a matrix satisfying Assumption 1.1. Then*

$$[d_\infty^r(P)]^2 \leq \frac{1}{4N} \operatorname{tr}\left[(I-P)(I-\tilde{P}\tilde{P}^*)^{-1}(I-P)^*\right].$$

where $\tilde{P} = P\Omega$. In particular, if P is normal, and $\sigma(P) = \{1, \lambda_1, \dots, \lambda_{N-1}\}$ denotes the spectrum of P, we have that

$$[d_\infty^r(P)]^2 \leq \frac{1}{4N}\sum_{i=1}^{N-1} \frac{|1-\lambda_i|^2}{1-|\lambda_i|^2}. \tag{2.19}$$

Proof: Define $\Sigma_{yy}(t) = \mathbb{E}[y_r(t)y_r(t)^*]$, and remark that $\frac{1}{N}\mathbb{E}[||y_r(t)||^2] = \frac{1}{N}\operatorname{tr}\Sigma_{yy}(t)$. Using the facts that $\Omega^k = \Omega$ for all positive integer k and $\Omega(P-I) = P-I$, it is easy to see that y_r satisfies the following recursive equation

$$y_r(t+1) = \tilde{P}y_r(t) + (P-I)n(t).$$

Now, thanks to the hypotheses on $n_i(t)$,

$$\begin{aligned}\Sigma_{yy}(t+1) &= \mathbb{E}[y_r(t+1)y_r(t+1)^*] \\ &= \mathbb{E}[\tilde{P}y_r(t)y_r(t)^*\tilde{P}^*] + (I-P)\mathbb{E}[n(t)n(t)^*](I-P)^* \\ &\leq \tilde{P}\,\Sigma_{yy}(t)\,\tilde{P}^* + \frac{1}{4}(I-P)(I-P)^*,\end{aligned}$$

where we used the bound on the variance. Then a simple recursion gives

$$\Sigma_{yy}(t) \leq \tilde{P}^t\,\Sigma_{yy}(0)\,(\tilde{P}^*)^t + \frac{1}{4}\sum_{s=0}^{t-1}\tilde{P}^s(I-P)(I-P)^*(\tilde{P}^*)^s.$$

Since P satisfies Assumption 1.1, from Lemma 1.2, $\rho_{\text{ess}}(P^*P) < 1$. Moreover we have that $\rho(\tilde{P}) = \rho_{\text{ess}}(P) < 1$ and $\rho(\tilde{P}^*\tilde{P}) = \rho_{\text{ess}}(P^*P) < 1$. Using the linearity and cyclic property of the trace,

$$\begin{aligned}\operatorname{tr}\Sigma_{yy}(t) &= \operatorname{tr}\left[\tilde{P}^t\,\Sigma_{yy}(0)\,(\tilde{P}^*)^t\right] + \operatorname{tr}\left[\frac{1}{4}\sum_{s=0}^{t-1}(PP^* - (P+P^*) + I)(\tilde{P}\tilde{P}^*)^s\right] \\ &= \operatorname{tr}\left[\tilde{P}^t\,\Sigma_{yy}(0)\,(\tilde{P}^*)^t\right] + \frac{1}{4}\operatorname{tr}\left[(PP^* - (P+P^*) + I)(I - (\tilde{P}\tilde{P}^*)^t)(I - \tilde{P}\tilde{P}^*)^{-1}\right].\end{aligned}$$

and hence

$$\lim_{t\to\infty}\operatorname{tr}\Sigma_{yy}(t) \leq \frac{1}{4}\operatorname{tr}\left[(PP^* - (P+P^*) + I)(I - \tilde{P}\tilde{P}^*)^{-1}\right].$$

If moreover P is normal, we can find a unitary matrix O of eigenvectors and a diagonal matrix of eigenvalues Λ, such that $P = O\Lambda O^*$. This implies

$$\operatorname{tr}\left[(PP^* - (P+P^*) + I)(1 - \tilde{P}\tilde{P}^*)^{-1}\right] = \sum_{i=1}^{N-1}\frac{|1-\lambda_i|^2}{1-|\lambda_i|^2}.$$

□

Remark that if we could assume the random variable $n_i(t)$ to have a common variance (or a common bound on it) σ^2, the above result would become

$$[d_\infty^r(P)]^2 = \frac{\sigma^2}{N}\operatorname{tr}\left[(I-P)(I-\tilde{P}\tilde{P}^*)^{-1}(I-P)^*\right].$$

We shall not comment further on this more restrictive setting, which has been treated in [41].

From now on we restrict to the case in which P is normal. Note that the expression for the mean square error of formula (2.19) is the product of two terms,

$$d_\infty^r(P) = \frac{1}{4}\Phi(P)$$

where

$$\Phi(P) = \frac{1}{N}\sum_{i=1}^{N-1}\frac{|1-\lambda_i|^2}{1-|\lambda_i|^2},$$

is a functional[3] of the matrix P, depending only on its spectral structure.

In the sequel of this section, we focus on the functional Φ: we compute it for some topologies, and we discuss how it depends on N for sequences of matrices of increasing size.

We start by observing that, if $\rho_{\text{ess}}(P) \leq B < 1$ then $\Phi(P) \leq \frac{N-1}{N}\frac{4}{1-B^2}$. This implies that, given a sequence of matrices of increasing size, if their essential spectral radiuses are uniformly bounded away from 1, then the functional cost $\Phi(P)$ is uniformly bounded in N. This fact is true also for the worst case bound d_∞^w as from Proposition 2.1. The interesting fact is that the performance index $\Phi(P)$ can exhibit the same behavior *even when the essential spectral radiuses are not bounded away from* 1.

Let us start with some examples.

Example 2.6 (Complete graph) In this case it is easy to compute that

$$\Phi(N^{-1}\mathbf{1}\mathbf{1}^*) = \frac{N-1}{N}.$$

Example 2.7 (Directed ring) Consider the matrix P defines in Example 2.4. In this case $\Phi(P)$ can be exactly computed. We have

$$\begin{aligned}
\Phi(P_N) &= \frac{1}{N}\sum_{h=1}^{N-1}\frac{(1-\lambda_h)(1-\lambda_h^*)}{1-\lambda_h\lambda_h^*}\\
&= \frac{1}{N}\sum_{h=1}^{N-1}\frac{(1-(1/2+1/2\,e^{i\frac{2\pi}{N}h}))(1-(1/2+1/2\,e^{-i\frac{2\pi}{N}h}))}{1-(1/2+1/2\,e^{i\frac{2\pi}{N}h})(1/2+1/2\,e^{-i\frac{2\pi}{N}h})}\\
&= \frac{1}{N}\sum_{h=1}^{N-1}\frac{1/2(1-\cos(\frac{2\pi}{N}h))}{1/2(1-\cos(\frac{2\pi}{N}h))} = \frac{N-1}{N}.
\end{aligned}$$

Example 2.8 (Undirected ring) Consider the undirected ring graph and the matrix P introduced in Example 2.5. The eigenvalues of P are

$$\lambda_h = \frac{1}{3} + \frac{2}{3}\cos\left(\frac{2\pi}{N}h\right) \quad h = 0,\ldots,N-1,$$

and we have

$$\Phi(P_N) = \frac{1}{N}\sum_{h=1}^{N-1}\frac{1-\lambda_h}{1+\lambda_h} = \frac{1}{N}\sum_{h=1}^{N-1}\frac{1-\cos\left(\frac{2\pi}{N}h\right)}{2+\cos\left(\frac{2\pi}{N}h\right)}.$$

[3] Remarkably, the functional $\Phi(P)$ also arises, with a rather different meaning, in [34], as a cost functional depending on the transient of the algorithm (2.1) over graphs with ideal communication.

In this case it is difficult to work out the computation explicitly. However, it is possible to compute the limit for $N \to \infty$, since the summation can be interpreted as Riemann sum relative to the function $f(x) = \dfrac{1-\cos(x)}{2+\cos(x)}$. We thus obtain

$$\lim_{N\to\infty} \Phi(P_N) = \frac{1}{2\pi}\int_0^{2\pi} \frac{1-\cos(x)}{2+\cos(x)}dx = \sqrt{3}-1. \tag{2.20}$$

Example 2.9 (Hypercube) Consider the hypercube graph and the uniform Perron matrix $P = \frac{1}{n+1}(I+A)$. From [31], P has $n+1$ distinct eigenvalues, which are, for $k \in \{0,\dots,n\}$,

$$\lambda_k = 1 - \frac{2k}{n+1}, \tag{2.21}$$

with multiplicity $p_k = \binom{n}{k}$. We have that

$$\begin{aligned}
\Phi(P) &= \frac{1}{N}\sum_{i=1}^{N-1} \frac{|1-\lambda_i|^2}{1-|\lambda_i|^2} \\
&= \frac{1}{2^n}\sum_{k=1}^{n} \frac{\left(\frac{2k}{n+1}\right)^2}{1-\left(\frac{n+1-2k}{n+1}\right)^2}\binom{n}{k} \\
&= \frac{1}{2^n}\sum_{k=1}^{n} \frac{k}{n+1-k}\binom{n}{k} \\
&= \frac{1}{2^n}\sum_{k=1}^{n} \binom{n}{k-1} \\
&= \frac{2^n-1}{2^n} = \frac{N-1}{N}.
\end{aligned}$$

Hence,

$$\Phi(P) = \frac{N-1}{N}.$$

Remark 2.1 While in the previous section the hypercube provided the negative example for the worst-case behavior, the probabilistic analysis is in agreement with the evidence showed in the simulations. This highlights the differences between the bounded error model and the probabilistic model: with the same assumptions on P the two worst-case analysis and the mean-square analysis give different results.

The above cases are encompassed by the following general result, which was conjectured in [41], and has been recently proved by Federica Garin and Enrico Lovisari [45].

Theorem 2.8 *Let $N = n^d$, and P be the uniform Perron matrix adapted to the $d-$dimensional directed torus. Then,* $\Phi(P) = \dfrac{n^d-1}{n^d}$.

Proof: We can see P as the (linear) map $P : \mathbb{R}^{\mathbb{Z}_n^d} \to \mathbb{R}^{\mathbb{Z}_n^d}$ defined by

$$(Px)_{a_1,\dots,a_d} = \frac{1}{d+1} x_{a_1,\dots,a_d} + \sum_{j=1}^{d} \frac{1}{d+1} x_{a_1,\dots,a_j-1,\dots,a_d}$$

Since $\Phi(P) = \frac{1}{n^d} \sum_{\substack{\lambda\in\sigma(P)\\ \lambda\neq 1}} \frac{|1-\lambda|^2}{1-|\lambda|^2}$, where $\sigma(P)$ is the spectrum of P, we start from the knowledge of the eigenvalues of P,

$$\sigma(P) = \{\lambda_a : a \in \mathbb{Z}_n^d\} \qquad \text{where } \lambda_a = \frac{1}{d+1}(1 + \omega^{a_1} + \cdots + \omega^{a_d}),\ \omega = e^{i2\pi/n}$$

(considered with their multiplicities).

The key idea of the proof is to partition $\sigma(P)$ in disjoint subsets such that within any subset $|\lambda|$ is constant and moreover, for any such subset O, $\sum_{\lambda\in O} \frac{|1-\lambda|^2}{1-|\lambda|^2} = |O|$. The partition is going to be the orbits under the action of a suitably defined group.

Define the (linear) map $g : \mathbb{Z}_n^d \to \mathbb{Z}_n^d$ given by

$$g(a_1, \dots, a_d) = (a_2 - a_1, a_3 - a_1, \dots, a_d - a_1, -a_1)$$

Remark that g is, you can see it as the composition of the embedding of $\mathbb{Z}_n^d$ in $\mathbb{Z}_n^{d+1}$ (inserting a 0 in the last position), the substraction of the first entry to all entries, and the projection of $\mathbb{Z}_n^{d+1}$ onto $\mathbb{Z}_n^d$ by removing the first entry, as follows:

$$\begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_d \end{bmatrix} \longmapsto \begin{bmatrix} a_1 \\ a_2 \\ \vdots \\ a_d \\ 0 \end{bmatrix} \longmapsto \begin{bmatrix} 0 \\ a_2 - a_1 \\ \vdots \\ a_d - a_1 \\ -a_1 \end{bmatrix} \longmapsto \begin{bmatrix} a_2 - a_1 \\ \vdots \\ a_d - a_1 \\ -a_1 \end{bmatrix}$$

Computing the successive powers of g, we obtain two useful facts:

(i) g^{d+1} is the identity: indeed

$$\begin{aligned}
g^2(a_1,\dots,a_d) &= (a_3 - a_2, a_4 - a_2, \dots, a_{d-1} - a_2, a_d - a_2, -a_2, a_1 - a_2)\\
g^3(a_1,\dots,a_d) &= (a_4 - a_3, a_5 - a_3, \dots, a_d - a_3, -a_3, a_2 - a_3, a_1 - a_3)\\
&\dots\\
g^{d-1}(a_1,\dots,a_d) &= (a_d - a_{d-1}, -a_{d-1}, a_1 - a_{d-1}, \dots, a_{d-2} - a_{d-1})\\
g^{d}(a_1,\dots,a_d) &= (-a_d, a_1 - a_d, \dots, a_{d-1} - a_d)\\
g^{d+1}(a_1,\dots,a_d) &= (a_1,\dots,a_d)
\end{aligned}$$

(ii) $\forall a = (a_1,\dots,a_d) \in \mathbb{Z}_n^d$ and $\forall k \in \{1,\dots,d\}$, $\lambda_{g^k(a)} = \omega^{-a_k}\lambda_a$. Then, $|\lambda_{g^k(a)}| = \lambda_a$.

These two facts allow to prove that

$$\begin{aligned}\sum_{k=0}^{d} \lambda_{g^k(a)} &= \sum_{k=0}^{d} \omega^{-a_k}\lambda_a = \lambda_a(1+\omega^{-a_1}+\cdots+\omega^{-a_d})\\ &= \lambda_a(d+1)\lambda_a^* = (d+1)|\lambda_a|^2,\end{aligned}$$

and deduce that

$$\begin{aligned}\sum_{k=0}^{d} \frac{|1-\lambda_{g^k(a)}|^2}{1-|\lambda_{g^k(a)}|^2} &= \sum_{k=0}^{d} \frac{1+|\lambda_{g^k(a)}|^2 - 2\,\mathrm{Re}(\lambda_{g^k(a)})}{1-|\lambda_{g^k(a)}|^2}\\ &= \frac{d+1+(d+1)|\lambda_a|^2 - 2\sum_{k=0}^{d}\mathrm{Re}(\lambda_{g^k(a)})}{1-|\lambda_a|^2}\\ &= \frac{d+1+(d+1)|\lambda_a|^2 - 2(d+1)\lambda_a^2}{1-|\lambda_a|^2}\\ &= d+1 \qquad\qquad \forall a\in\mathbb{Z}_n^d.\end{aligned}$$

Let G be the cyclic group $G=\{1,g,g^2,\dots,g^d\}$, and denote by O_a the orbit of $a\in\mathbb{Z}_n^d$ under the action of G, i.e. $O_a=\{a,g(a),g^2(a),\dots,g^d(a)\}$. It is well known that the orbits are all disjoint. Moreover, notice that for each orbit O, if you choose any element $a\in O$, then

$$\begin{aligned}\sum_{\lambda\in O}\frac{|1-\lambda|^2}{1-|\lambda|^2} &= \sum_{a\in O}\frac{1}{|G|}\sum_{k=0}^{d}\frac{|1-\lambda_{g^k(a)}|^2}{1-|\lambda_{g^k(a)}|^2}\\ &= \frac{|O|}{|G|}\sum_{k=0}^{d}\frac{|1-\lambda_{g^k(a)}|^2}{1-|\lambda_{g^k(a)}|^2}\\ &= \frac{|O|}{|G|}(d+1) = |O|\end{aligned}$$

Finally notice that the eigenvalue $1=\lambda_0$ and $O_0=\{0\}$. So we can conclude:

$$\begin{aligned}\Phi(P) &= \frac{1}{n^d}\sum_{\substack{\lambda\in\sigma(P)\\ \lambda\neq 1}}\frac{|1-\lambda|^2}{1-|\lambda|^2}\\ &= \frac{1}{n^d}\sum_{\substack{a\in\mathbb{Z}_n^d\\ \lambda\neq 0}}\frac{|1-\lambda_a|^2}{1-|\lambda_a|^2}\\ &= \frac{1}{n^d}\sum_{\substack{O\in\mathbb{Z}_n^d/G\\ O\neq O_0}}|O|\\ &= \frac{1}{n^d}(|Z_n^d|-1).\end{aligned}$$

□

We thus have a class of graphs for which $\Phi(P)$ does not diverge when N goes to infinity, while $\rho_{\mathrm{ess}}(P)$ does. Moreover, numerical computations demonstrate that the same is true for sequences of random geometric graphs of increasing size. A theoretical proof of this fact is still missing.

After these examples, we look for some more general condition.

Proposition 2.9 *Let $B_{c,r}$ be the closed ball of complex numbers with center c and radius r, and let P satisfy Assumption 1.1. Then it exists $0 < R < 1$ such that $\sigma(P) \subseteq B_{1-R,R}$, and*

$$\Phi(P) \leq \frac{R}{1-R}. \tag{2.22}$$

Proof: The existence of such R is a direct consequence of Lemma 1.2. It means that the spectrum is contained in a disc of radius R internally tangent in 1 to the unit disc of the complex plane.

Then, we need to prove (2.22). Every eigenvalue $\lambda \in B_{1-R,R}$, is such that

$$\lambda = (1-r) + r\, e^{i\theta}$$

with $\theta \in [0, 2\pi[$ and $0 \leq r \leq R$. Moreover, if $\lambda \neq 0$, then $\theta > 0$, and

$$\begin{aligned}
\frac{|1-\lambda|^2}{1-|\lambda|^2} &= \frac{|r - re^{i\theta}|^2}{1-|1-r+r\,e^{i\theta}|^2} \\
&= \frac{r^2|1-e^{i\theta}|^2}{1-(1-r)^2-2r(1-r)\cos\theta - r^2} \\
&= \frac{r^2 2(1-\cos\theta)}{2r(1-r)(1-\cos\theta)} \\
&= \frac{r}{1-r} \leq \frac{R}{1-R} \quad \forall\, \lambda \neq 0.
\end{aligned}$$

This yields the result. □

Note that the bound (2.22) depends only on R and *does not depend on the essential spectral radius of the matrix P*, while the worst-case bounds provided in Proposition 2.3 and in Proposition 2.4 did. The following result takes advantage of this remark.

Corollary 2.10 *Let P satisfy Assumption 1.1, and let $p = \min_i P_{ii}$. Then*

$$\Phi(P) \leq \frac{1-p}{p}. \tag{2.23}$$

Proof: Lemma 1.2 and Gershgorin Theorem imply that

$$\sigma(P) \subseteq \bigcup_i B_{P_{ii},1-P_{ii}} \subseteq B_{p,1-p}.$$

□

The above result has the following interpretation. If in a family of matrices P_N we have that $\min_i (P_N)_{ii}$ is lower bounded uniformly in N, then (2.23) gives a uniform bound to the asymptotic displacement. This is a useful hint to construct sequences of matrices whose performance scales well with N. It would be enough to prescribe that the agents assign a minimum weight to their own values. Such a condition has appeared elsewhere in the literature to guarantee robustness of consensus algorithms [69].

2.4 Simulations and conclusion

In this section we compare the evolution in time and the performance of the deterministic quantization algorithm (2.2), of the probabilistic quantization algorithm (2.3), and of the probabilistic model (2.18) (with noises sampled uniformly in $(-1/2, 1/2)$), which has been defined for analysis purposes. Indeed, one of the goals of this section is to validate the models we introduced in the theoretical analysis.

2.4.1 Speed of convergence

Figure 2.1 shows an example of the time evolution of the three algorithms above, comparing them with the ideal communication algorithm (1.5). It is evident that, as long as the disagreement is far from the size of the quantization step, there is no significant difference between the algorithms, and they converge exponentially with the same rate. Instead, when the quantization step and the disagreement value become comparable, the quantization effects become prevalent.

2.4.2 Dependence on γ

We have seen in Section 1.3.1 that the matrix P can be chosen in a way to minimize the convergence rate of the ideal communication algorithm. On the other hand, from the results above in this chapter, it appears the possibility of optimizing the achievable performance $d_\infty(P)$, or at least $\Phi(P)$ which is a good description of $d_\infty(P)$. A natural question is then to ask whether these two optimization problems are independent, or rather there is some trade-off between speed and achievable precision. It comes out that the latter is the case.

To give an example of this, we focus on the special case of weights dependent on a gain parameter γ, representing the confidence given to received information. Let us then consider weighted Laplacian Perron matrices

$$P_\gamma = I + \gamma D^{-1} L_{\mathcal{G}},$$

and undirected ring graphs. Simulations are shown in Figure 2.2, and moreover we can readily extend the computations in Example 2.8 to prove the following fact.

Proposition 2.11 *Let $\mathcal{G}$ be an undirected ring, and $P_\gamma = I + \gamma D^{-1} L_{\mathcal{G}}$. Then*

$$\lim_{N\to\infty} \Phi(P_\gamma) = \frac{1}{\sqrt{1-\gamma}} - 1.$$

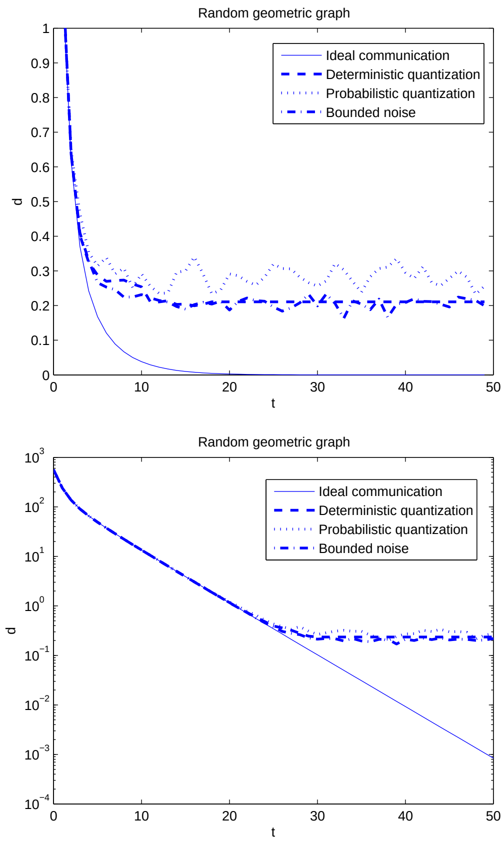

Figure 2.1: Disagreement $d(t)$ for algorithms (2.2), (2.3), (2.18) and (1.5). Random geometric graph on 20 nodes, maximum degree weights.

Let us comment on these figures. If $\gamma > 1/2$, that is, agents are very confident in the incoming information, then the performance worsen unboundedly as γ tends to 1. If instead $\gamma < 1/2$, that is, agents are less confident in the incoming information, then we have good convergence results, similar to the case $\gamma = 1/2$. However, we do not see in simulations the error going to zero as γ goes to zero, as suggested by the probabilistic analysis. The reason is clear: with quantized communication we can have no warranty, in general, of having the asymptotic states closer than 1 in the infinite-norm. This highlights that the probabilistic approach is not significant when we are already so near to the consensus.

The experimental results are according to what can be proven in a parametric study of the simplest communication configuration, when the system is made of only two communicating agents.

Example 2.10 Suppose $N = 2$, $x = (x_1, x_2)$ and

$$P = \begin{pmatrix} 1-\gamma & \gamma \\ \gamma & 1-\gamma \end{pmatrix}.$$

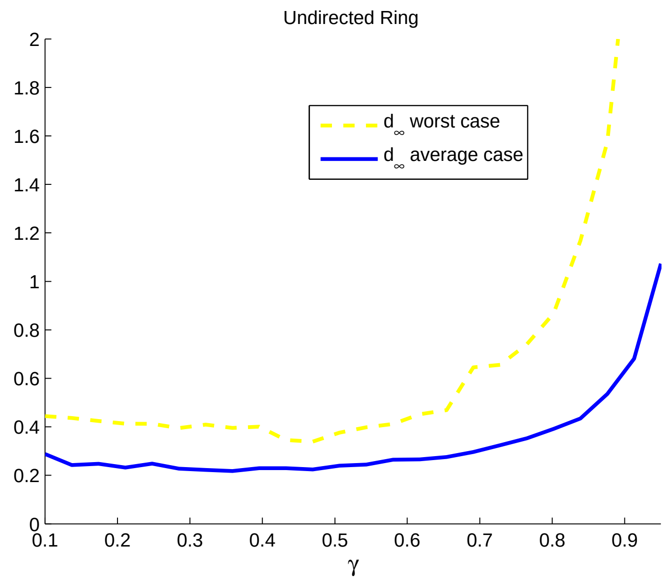

Figure 2.2: Dependence of the performance d_∞ on γ for a directed ring. Plot shows average and worst case of 25 simulations with random initial conditions.

Then using (2.14) it's easy to compute that

$$d_\infty(P) \leq \begin{cases} \frac{1}{2}\frac{\gamma}{1-\gamma} & \text{if } \gamma \geq 1/2 \\ \frac{1}{2} & \text{if } \gamma \leq 1/2. \end{cases}$$

It is interesting to remark that the obtained bound is *tight* for $\gamma \to 1$ and for $\gamma \leq 1/2$. Indeed, we can check that the system has a periodic point of period 2 such that $x = ((1-\gamma)A, A)^T$ with $A = \frac{1}{2-\gamma}\left\lfloor \frac{1}{2}\frac{2-\gamma}{1-\gamma}\right\rfloor$. The amplitude of this cycle is $x_2 - x_1 = \left\lfloor \frac{1}{2}\frac{2-\gamma}{1-\gamma}\right\rfloor \frac{\gamma}{2-\gamma}$ and this is asymptotically equivalent to $\frac{1}{1-\gamma}$ for $\gamma \to 1$. Moreover, as already remarked in general, the states can differ by one at equilibrium points, and this gives the tightness for small γ.

Now we can show the optimization trade-off between $\rho_{\text{ess}}(P)$ and $d_\infty(P)$. Indeed, it is shown in [24] that, if P is as in Proposition 2.11, then the optimal $\rho_{\text{ess}}(P)$ is achieved for

$$\gamma = \frac{2}{3 - \cos\frac{2\pi}{N}}.$$

Then, for the undirected ring the optimal γ goes to 1 for N going to infinity. But Proposition 2.11 implies that such a choice is not compatible with achieving a prescribed precision. As a consequence, a non trivial trade-offs between speed of convergence and limit behavior is apparent.

To conclude, we have observed that, at least on some families of graphs, one can not use the gains to optimize, given the graph, the achievable precision without trading-off the speed of convergence. If instead one could design the network topology, then choosing a graph with a large spectral gap would imply a low essential spectral radius of P, and then both a fast convergence and good precision.

2.4.3 Dependence on N

Simulations are performed to investigate the dependence on N of d_∞, for algorithms (2.2), (2.3) and (2.18). All plots are obtained averaging the outcome of the algorithms, which have been run from 25 random initial conditions, sampled from a uniform distribution in $[-100, 100]$, for each value of N. To have a range of significant examples, we consider four different topologies, which we have already shown to be significant:

(i) the undirected ring with uniform weights (Figure 2.3);

(ii) the undirected torus with uniform weights (Figure 2.4);

(iii) the hypercube with uniform weights (Figure 2.5);

(iv) the *random geometric* graph, where the threshold has been chosen as $R = \sqrt{5 \log N/(N)}$, with maximum degree weights (Figure 2.6, in this case we averaged also with respect to different realizations of the graph).

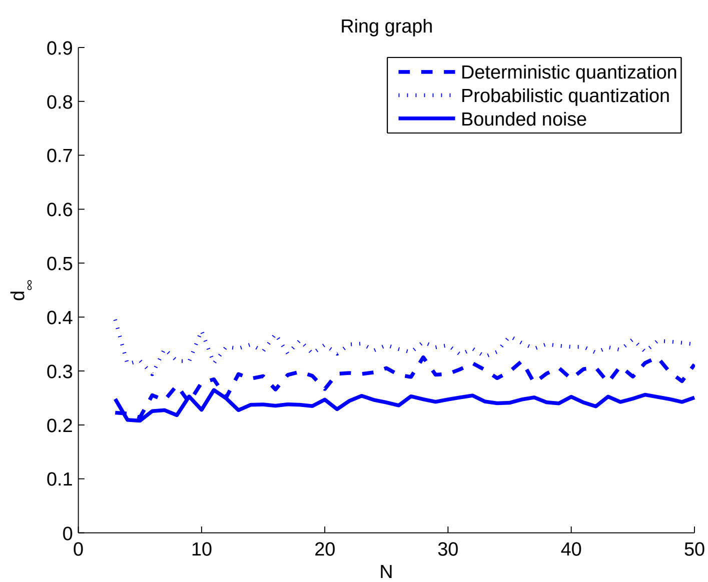

Figure 2.3: Performance of the undirected rings.

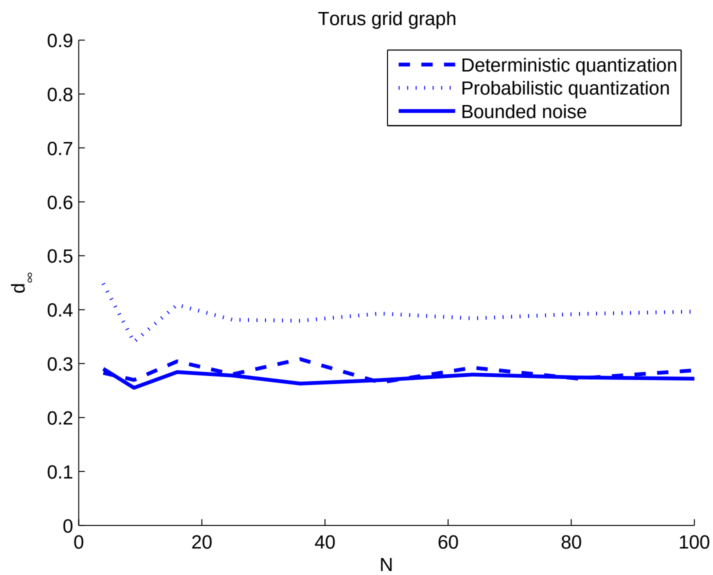

Figure 2.4: Scaling of performance of the undirected torus graphs.

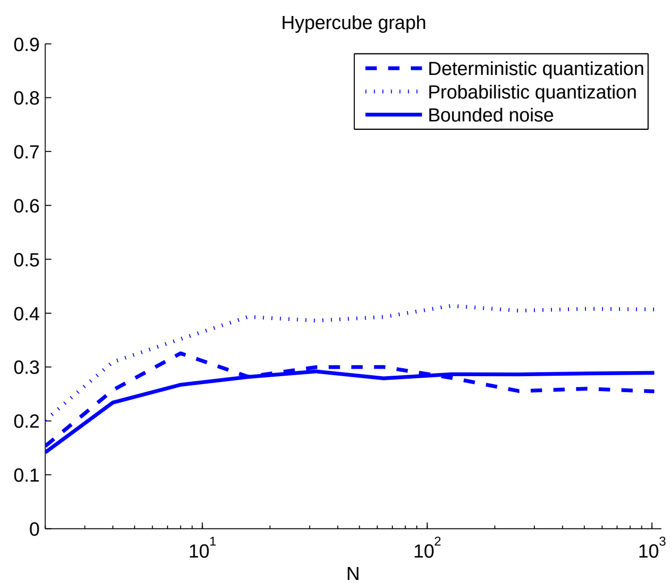

Figure 2.5: Scaling of performance of the hypercube graphs.

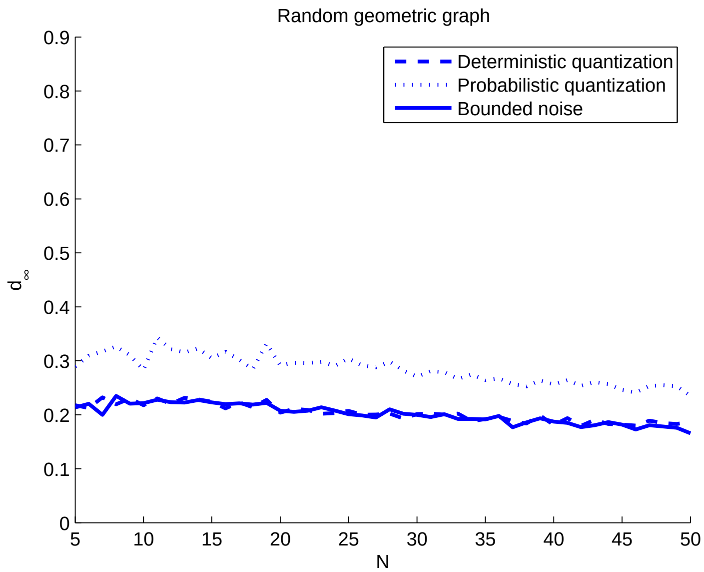

Figure 2.6: Scaling of performance of random geometric graphs.

From our simulations we can draw several remarks. Even though the theoretical results suggest that the performance could worsen with increasing N, simulations show a nicer behavior. In all the considered cases, the errors $d_\infty(P)$ are small, compared to the quantization step, and not increasing with N. Hence, we argue that the worst-case analysis, which predicts $d^w_\infty(P)$ to be logarithmical in N for the hypercube, is too pessimistic. This can be understood looking at the proof of this fact: the sequence of quantization errors which gives the worst effect is one extreme event. Then, it is not likely to be seen in simulations, which show typical behaviors. Instead, the probabilistic approach finds an *a posteriori* justification in its better agreement with simulations.

2.4.4 Conclusions

In this chapter we studied the effects of a uniform quantization on the average consensus problem, and, starting from the well-known linear diffusion algorithm, we proposed a simple and effective adaptation which is able to preserve the average of states and to drive the system reasonably near to the consensus. A special attention has been given to the scalability in N of the performance, which is a crucial issue for applications in which the number of agents is huge. In this direction we obtained several favorable results, which we applied to sequences of graphs with regularity. In particular, we showed how the convergence properties of the algorithm depend on the spectrum of a matrix defining the system: this gives deep theoretical insights, and also leads to useful design suggestion, as in Corollary 2.10.

An exact analysis of the nonlinear dynamics of the whole system has been performed in one special case, and general convergence results have been obtained by a worst-case analysis and by a probabilistic model. The relationship between worst-case and probabilistic models deserves some more comment. Indeed, the worst-case model points out a dependence of the performance from the number of agents, which is a bad news for the applications. The probabilistic model, instead, assuming a probabilistic structure on the quantization errors, does not show this dependence. However, such dependence is not apparent from simulations. We argue from this difference that the poor worst-case performance is due to few bad realizations of the errors, which have a small probability in the probabilistic model, and do not happen in the actual quantized system.

Appendices

Here we include some long proofs which we have skipped during the chapter.

2.A Analysis of the Example 2.2

Here we construct and study the symbolic dynamics underlying the system (2.8). To start, we need the following technical lemma. Let $\lfloor\cdot\rfloor$ and $\lceil\cdot\rceil$ denote the floor and ceiling operators from $\mathbb{R}$ to $\mathbb{Z}$.

Lemma 2.12 *Given $\alpha, \beta \in \mathbb{N}$ and $x \in \mathbb{R}$, it holds*

$$\lfloor x\rfloor = \left\lfloor \frac{\lfloor \alpha x\rfloor}{\alpha}\right\rfloor \tag{2.24}$$

$$q_d(x) = \lfloor x + 1/2\rfloor = \left\lceil \frac{1}{2}\left\lfloor \frac{\lfloor 2\beta x\rfloor}{\beta}\right\rfloor\right\rceil . \tag{2.25}$$

Proof: We first prove (2.24). Let $m = \lfloor x\rfloor$. So

$$\begin{aligned} m &\leq x < m+1 \\ \alpha m &\leq \alpha x < \alpha m + \alpha. \end{aligned}$$

Hence, we can find $s \in \mathbb{N}, 0 \leq s \leq \alpha - 1$ such that $\alpha m + s \leq \alpha x < \alpha m + s + 1$. This yields $\lfloor \alpha x\rfloor = \alpha m + s$ and $\left\lfloor \frac{\lfloor \alpha x\rfloor}{\alpha}\right\rfloor = m$.

Then we prove equation (2.25). The equality $q_d(x) = \lfloor x + 1/2\rfloor$ is clear from the definition of $q_d(x)$. To prove the second equality, let $h = \lfloor 2x\rfloor$. Then $h \leq 2x < h+1$, from which follows that

$$\frac{h}{2} + \frac{1}{2} = \frac{h+1}{2} \leq x + 1/2 < \frac{h+2}{2} = \frac{h}{2} + 1.$$

From this inequality it follows that $\lfloor x + 1/2\rfloor = \lceil \frac{h}{2}\rceil$. This, with (2.24), implies (2.25). $\square$

We define $n_i(t) = \lfloor 2x_i(t) \rfloor$. Simple properties of floor and ceiling operators, together with the above lemma, allow us to remark that $q_d(x_i(t)) = \left\lceil \frac{n_i(t)}{2} \right\rceil$ and to derive from (2.8) that

$$\begin{aligned}
x_i(t+1) &= x_i(t) - \frac{1}{2}q_d(x_i(t)) + \frac{1}{2}q_d(x_{i+1}(t)) \\
\lfloor 2x_i(t+1) \rfloor &= \lfloor 2x_i(t) \rfloor + q_d(x_{i+1}(t)) - q_d(x_i(t)) \\
n_i(t+1) &= n_i(t) + \left\lceil \frac{n_{i+1}(t)}{2} \right\rceil - \left\lceil \frac{n_i(t)}{2} \right\rceil \\
n_i(t+1) &= \left\lfloor \frac{n_i(t)}{2} \right\rfloor + \left\lceil \frac{n_{i+1}(t)}{2} \right\rceil .
\end{aligned}$$

We have thus found an iterative system involving only the symbolic signals $n_i(t)$:

$$\begin{aligned}
&n_i(t+1) = g_1(n_i(t), n_{i+1}(t)) \qquad \forall\, i \in \{1, \dots, N\} \\
&\text{where } g_1(h,k) = \left\lfloor \frac{h}{2} \right\rfloor + \left\lceil \frac{k}{2} \right\rceil .
\end{aligned} \tag{2.26}$$

Since $n_i(t) = \lfloor 2x_i(t) \rfloor$, the asymptotic analysis of (2.26) allows us to obtain information about the asymptotic of $x_i(t)$ up to quantization errors equal to 1.

We now start the analysis of system (2.26). Define the following quantities: $m(t) = \min_{1 \le i \le N} n_i(t)$, $M(t) = \max_{1 \le i \le N} n_i(t)$, $D(t) = M(t) - m(t)$. From the form of (2.26) one can easily remark that $m(t)$ can not decrease and $M(t)$ can not increase. Hence $D(t)$ is not increasing.

A much stronger result about the monotonicity of $D(t)$ is the content of the following lemma.

Lemma 2.13 *If $D(t_0) \ge 2$, there exists $T \in \mathbb{N}$ such that $D(t_0 + T) < D(t_0)$.*

Proof: Let $I_m(t) = \{j \in \mathbb{Z}_N \text{ s.t. } n_j(t) = m(t)\}$. The idea of the proof is to show that the cardinality of the set $I_m(t)$ eventually decreases if $D(t) \ge 2$.

Notice first that, for $h, k \in \mathbb{Z}$, $g_1(h+2, k+2) = g(h,k) + 2$. Hence, by an appropriate translation of the initial condition, we can always restrict ourselves to the case $m(t_0) \in \{0, 1\}$.

Case $m(t_0) = 0$. Notice that

$$g_1(h,k) > 0 \;\forall h \ge 0,\, k > 0, \quad g_1(h,0) > 0 \;\forall h \ge 2 .$$

This easily implies that $I_m(t)$ is not increasing. Now, since $D(t_0) \ge 2$, we can find $j_0 \in I_m(t_0)$ and two integers $U > 0$ and $W \ge 0$ such that

$$\begin{aligned}
&n_{j_0-W-1}(t_0) > 1 \\
&n_{j_0-v}(t_0) = 1 \quad 0 < v \le W \\
&n_{j_0+u}(t_0) = 0 \quad 0 \le u < U \\
&n_{j_0+U}(t_0) > 0.
\end{aligned}$$

After W time steps,

$$\begin{aligned} &n_{j_0-W-1}(t_0+W) > 1 \\ &n_{j_0-u}(t_0+W) = 0 \quad W-U+1 \le u \le W \\ &n_{j_0+W-U}(t_0+W) > 0. \end{aligned}$$

At the following time step, one 0 will then disappear

$$\begin{aligned} &n_{j_0-W-1}(t_0+W+1) > 0 \\ &n_{j_0-u}(t_0+W+1) = 0 \quad W-U+2 \le u \le W \\ &n_{j_0+W-U-1}(t_0+W+1) > 0. \end{aligned}$$

This implies that $|I_m(t_0+W+1)| < |I_m(t_0)|$.

Case $m(t_0)=1$. Notice that

$$g_1(h,k) > 1 \;\forall h \ge 2,\ k \ge 1, \quad g_1(1,k) > 1 \;\forall k \ge 3\,.$$

This easily implies that $|I_m(t)|$ is again not increasing. Now, since $D(t_0) \ge 2$, we can find $j_0 \in I_m(t_0)$ and an integer $W \ge 0$ such that

$$\begin{aligned} &n_{j_0}(t_0) = 1 \\ &n_{j_0+w}(t_0) = 2 \quad 1 \le w \le W \\ &n_{j_0+W+1}(t_0) > 2. \end{aligned}$$

The evolution of the above configuration yields, after W instant steps

$$\begin{aligned} &n_{j_0}(t_0+W) = 1 \\ &n_{j_0+1}(t_0+W) > 2. \end{aligned}$$

The next step, we obtain $n_{j_0}(t_0+W+1) > 1$. Therefore, $|I_m(t_0+W+1)| < |I_m(t_0)|$.

In both cases we have proven that $|I_m(t)|$ strictly decreases in finite number of steps. A straightforward induction principle then implies that a finite $T \in \mathbb{N}$ exists such that $m(t_0+T) > m(t_0)$. This proves the result. □

The interesting consequence of this lemma is the following result.

Theorem 2.14 *Let $n(t)$ evolve following* (2.26)*. Then, there exists $T \in \mathbb{N}$ and $h \in \mathbb{Z}$ such that, for all $t > T$, it holds that $D(t) < 2$ and one of the following condition is met*

(i) $n_i(t) = h$, $\forall i$*;*

(ii) $\{n_i(t) \,:\, i = 0, \dots N\} = \{2h-1, 2h\}$ *and each $n_i(t)$ is constant in time;*

(iii) $\{n_i(t) \,:\, i = 0, \dots N\} = \{2h, 2h+1\}$ *and each $n_i(t)$ is periodic in time of period N.*

Moreover, for all $t \in \mathbb{N}$, $\sum_{i=1}^N n_i(t) = \sum_{i=1}^N n_i(0)$.

Proof: From Lemma 2.13, it follows that a finite $T \in \mathbb{N}$ can be found, such that $D(t) < 2$ for all $t > T$. Once we reach his condition, there are two possibilities: either the $n_i(t)$ are all equal or they differ by 1. In the first case, the system remains constant (case (i)). In the second case, it follows from the way g is defined that if the lowest state is odd, the evolution is constant (case (ii)), while if the lowest one is even, the state evolution is a leftward shift (case (iii)). This is periodic of period N (and possibly also of some divisor of N). □

We can now go back to the original system

Corollary 2.15 *For system (2.8), there exists $T \in \mathbb{N}$ such that*

$$|x_i(t) - x_j(t)| \leq 1 \quad \forall\, i, j \quad \forall\, t > T, \tag{2.27}$$

and hence $d_\infty(P) \leq 1/2$.

Proof: Immediate consequence of Theorem 2.14, considering the relation $n_i(t) = \lfloor 2x_i(t) \rfloor$. □

Remark that (2.27) is satisfied while the system has reached a non consensus fixed point or a periodic point.

Remark 2.2 Using the quantizer q_T instead of q_d would induce a symbolic dynamics evolving via

$$g_5(h, k) = \left\lceil \frac{h}{2} \right\rceil + \left\lfloor \frac{k}{2} \right\rfloor,$$

which can be studied in an analogous way, leading to a result similar to Theorem 2.14

2.B Proof of Proposition 2.3

Proof: To upper bound $\sum_{s=0}^{\infty} \rho(P^s(I-P))$, we estimate $\rho(P^s(I-P)) = \max_{k=1}^{N-1} |\lambda_k^s(1-\lambda_k)|$ for all s. In order to do so we consider the function $f : \mathbb{C} \to \mathbb{R}$ defined as $f(z) = z^s(1-z)$. Let us consider the closed balls $B_{1-R,R}$ and $B_{0,\bar{\rho}}$. By Gershgorin's Theorem we have that $\sigma(P) \subseteq B_{1-R,R}$. By the definition of essential spectral radius it holds that $\sigma(P) \setminus \{1\} \subseteq B_{0,\bar{\rho}}$. Hence $\sigma(P) \setminus \{1\} \subseteq B_{0,\bar{\rho}} \cap B_{1-R,R}$. Let $A = B_{1-R,R} \cap B_{0,\bar{\rho}}$. Clearly

$$\max_{k=1}^{N-1} |\lambda_k^s(1-\lambda_k)| \leq \max_{z \in A} |f(z)|.$$

Since f is an analytic function and A is a compact set, from the Maximum Modulus Principle it follows that

$$\max_{k=1}^{N-1} |\lambda_k^s(1-\lambda_k)| \leq \max_{z \in \partial A} |f(z)|,$$

where ∂A denotes the boundary of A.
Consider now the curves $\gamma : [0, 2\pi] \to \mathbb{C}$,

$$\gamma(t) = 1 - R + Re^{jt},$$

and $\theta : [0, 2\pi] \to \mathbb{C}$,

$$\theta(t) = \bar{\rho} e^{jt},$$

which represent, respectively, the boundaries of $B_{1-R,R}$ and of $B_{0,\bar{\rho}}$. In the following, since $|f(z)| = |f(z^*)|$, we consider γ and θ only on the interval $[0, \pi]$.
By calculating the intersection between γ and θ one can see that $\partial A = \tilde{\gamma} \cup \tilde{\theta}$ where

$$\tilde{\gamma} = \left\{ z = z_x + i z_y \, \in \, \gamma \, : \, z_x \leq \frac{1 - 2R + \bar{\rho}^2}{2(1-R)} \right\}$$

and

$$\tilde{\theta} = \left\{ z = z_x + i z_y \, \in \, \theta \, : \, z_x \geq \frac{1 - 2R + \bar{\rho}^2}{2(1-R)} \right\}.$$

We then perform the optimization on these two sets separately. Consider $|f(z)|$ on $\tilde{\gamma}$. By straightforward calculations one can show that

$$|f(\gamma(t))|^2 = 2R^2(1 - \cos t) \left[1 - 2R(1-R)(1 - \cos t)\right]^s.$$

Now let $x = R\cos t + 1 - R$. In order to analyze the behavior of $|f(\gamma(t))|^2$ we introduce the following auxiliary function

$$F(x) = 2R(1 - x) \left[1 - 2(1-R)(1-x)\right]^s.$$

A straightforward calculation shows that studying $|f(z)|^2$ on $\tilde{\gamma}$ is equivalent to study F on the interval $I = \left[1 - 2R, \frac{1-2R+\bar{\rho}^2}{2(1-R)}\right]$. Since

$$F'(x) = 2R \left[1 - 2(1-R)(1-x)\right]^{s-1} \left[-1 + 2(1-R)(s+1)(1-x)\right],$$

we have that $F'(x) = 0$ for x equal to $x_1 = 1 - \frac{1}{2(1-R)}$ and $x_2 = 1 - \frac{1}{2(1-R)(s+1)}$, and that F is monotone increasing in $[x_1, x_2]$ and monotone decreasing for $[x_2, +\infty)$. Note that $1 - \frac{1}{2(1-R)} \leq 1 - 2R$ for all $R > 0$. Hence F, restricted to I, has a maximum in $1 - 2R$ if $x_2 \leq 1 - 2R$, in x_2 if $1 - 2R \leq x_2 \leq \frac{1-2R+\bar{\rho}^2}{2(1-R)}$, in $\frac{1-\bar{\rho}^2}{2R(1-R)}$ if $x_2 \geq \frac{1-\bar{\rho}^2}{2R(1-R)}$. We have that $x_2 \leq 1 - 2R \Leftrightarrow s \leq \frac{(1-2R)^2}{4R(1-R)}$, $1 - 2R \leq x_2 \leq \frac{1-2R+\bar{\rho}^2}{2(1-R)} \Leftrightarrow \frac{(1-2R)^2}{4R(1-R)} < s < \frac{\bar{\rho}^2}{1-\bar{\rho}^2}$, $x_2 \geq \frac{1-\bar{\rho}^2}{2R(1-R)} \Leftrightarrow s \geq \frac{\bar{\rho}^2}{1-\bar{\rho}^2}$. Let $\bar{s} = \lfloor \frac{(1-2R)^2}{4R(1-R)} \rfloor$ and $s^* = \lfloor \frac{\bar{\rho}^2}{1-\bar{\rho}^2} \rfloor$. Then

$$\max_{x \in I} F(x) = \begin{cases} 4R^2 \left(1 - 2R\right)^{2s} & \text{if} \quad s \leq \bar{s} \\ \frac{R}{1-R} \frac{s^s}{(s+1)^{s+1}} & \text{if} \quad \bar{s} + 1 \leq s \leq s^* \\ \frac{R}{1-R} \bar{\rho}^{2s} \left(1 - \bar{\rho}^2\right) & \text{if} \quad s \geq s^* + 1. \end{cases}$$

Let us now consider $|f(z)|^2$ on $\tilde{\theta}$. By simple algebraic manipulations one can see that

$$|f(\theta(t))|^2 = \bar{\rho}^{2s} \left(1 + \bar{\rho}^2 - 2\bar{\rho}\cos t\right)$$

Note that $|f(\theta(t))|^2$ is monotone increasing for $t \in [0, \pi]$ and hence it reaches its maximum on $\tilde{\theta}$ when $\cos t = \frac{1-2R+\bar{\rho}^2}{2(1-R)}$, that is at the intersection with $\tilde{\gamma}$.

Therefore we can conclude that

$$\max_{k=1}^{N-1} |\lambda_k^s(1-\lambda_k)| \leq \begin{cases} 2R\,|1-2R|^s & \text{if} \quad s \leq \bar{s} \\ \sqrt{\frac{R}{1-R}\frac{s^s}{(s+1)^{s+1}}} & \text{if} \quad \bar{s}+1 \leq s \leq s^* \\ \sqrt{\frac{R}{1-R}\bar{\rho}^{2s}\,(1-\bar{\rho}^2)} & \text{if} \quad s \geq s^*+1. \end{cases}$$

Hence we can write

$$\begin{aligned} &\sum_{s=0}^{t-1} \rho(P^s(I-P)) \\ &\quad \leq \sum_{s=0}^{\bar{s}} 2R\,|1-2R|^s + \sum_{s=\bar{s}+1}^{s^*} \sqrt{\frac{R}{1-R}\frac{s^s}{(s+1)^{s+1}}} + \sum_{s=s^*+1}^{t-1} \sqrt{\frac{R}{1-R}\bar{\rho}^{2s}\,(1-\bar{\rho}^2)}. \end{aligned}$$

Notice now that

$$\sum_{s=0}^{\bar{s}} 2R\,|1-2R|^s \leq \frac{2R}{1-|1-2R|} \leq \frac{1}{1-R}.$$

and that

$$\begin{aligned} \sum_{s=s^*+1}^{t-1} \sqrt{\frac{R}{1-R}\bar{\rho}^{2s}\,(1-\bar{\rho}^2)} &\leq \sqrt{\frac{R}{1-R}}\sqrt{1-\bar{\rho}^2}\sum_{s=0}^{\infty}\bar{\rho}^s \\ &= \sqrt{\frac{R(1-\bar{\rho}^2)}{(1-R)(1-\bar{\rho})^2}} \\ &\leq \sqrt{\frac{2R}{(1-R)(1-\bar{\rho})}}. \end{aligned}$$

Notice finally that, since $\sum_{i=1}^{m}\sqrt{\frac{1}{i+1}} \leq 2\sqrt{m+1}$, we can argue that

$$\begin{aligned} \sum_{s=\bar{s}+1}^{s^*} \sqrt{\frac{R}{1-R}\frac{s^s}{(s+1)^{s+1}}} &\leq \sqrt{\frac{R}{1-R}}\sum_{s=1}^{s^*}\sqrt{\frac{1}{2}}\sqrt{\frac{1}{1+s}} \\ &\leq \sqrt{\frac{4R}{2(1-R)(1-\bar{\rho}^2)}} \\ &\leq \sqrt{\frac{2R}{(1-R)(1-\bar{\rho})}}. \end{aligned}$$

Putting together these three inequalities we obtain (2.15). □

Chapter 3

Quantized gossip

The goal of this chapter is to analyze the effects of quantization on the (symmetric) gossip algorithm (see Section 1.3.3 and [9]), or, equivalently, the opportunity of using a gossip communication when quantization of the messages is imposed. The agents' states are assumed to be real numbers, while the sent messages are integer numbers: we consider both a *deterministic* uniform quantizer and a *probabilistic* uniform quantizer, which are defined rigorously in Section 3.1. To perform the states update, we introduce two alternative strategies, the *partially quantized* strategy and the *globally quantized* strategy, depending on whether the systems use exact information regarding their own state, or not, to update their states. We analyze these strategies both with the deterministic quantizer and with the probabilistic quantizer.

Through both analytical results and simulations, we investigate two design questions: wether the agents should use the deterministic or the probabilistic quantizer, and wether they should use, or not, exact information regarding their own states in the update.

Our results are obtained by two different techniques. For the partially quantized strategy with probabilistic quantization, we give a mean squared error analysis and convergence is proved for time going to infinity. In the three other cases, we study a Markov chain symbolic dynamics, obtaining results of convergence in finite time. Such a fact is remarkable, since it underlines the discrete nature of the problem, in spite of the state space being continuous. We show that the globally quantized strategy, both using the deterministic quantizer and the probabilistic quantizer, ensures that, almost surely, the consensus is reached in a finite time. The drawback of this strategy is that it does not preserve the average of the initial conditions. On the other hand, the partially quantized strategy preserves the initial average at each iteration of the algorithm, but does not guarantee that an exact consensus is reached. However, we prove that the states get as close to consensus as the size of quantization steps.

The chapter is organized as follows. In Section 3.1 we formulate the problem. In particular we introduce the *partially* quantized strategy and the *globally* quantized strategy. In Section 3.2 and in Section 3.3, we analyze these two strategies assuming, respectively, that the systems quantize the information by means of deterministic quantizers and by means of probabilistic quantizers. Finally in Section 3.4 we gather out our conclusions.

A former version of this work appeared as [40] and a complete version, containing most of the material, as [23].

Recall the definition of the standard gossip algorithm (1.10). In this chapter, we make the following assumption.

Assumption 3.1 *The graph $\mathcal{G} = (V, \mathcal{E})$ is a undirected connected graph and, at every time instant $t \geq 0$, each edge $(i,j) \in \mathcal{E}$ can be selected with a strictly positive probability $W^{(i,j)}$.*

3.1 Problem statement

Note that the gossip algorithm (1.10) relies upon a crucial assumption: each agent transmits to its neighboring agents the precise value of its state. Here we assume that the communication network is constituted of digital links. This prevents the agents from having a precise knowledge about the state of the other agents. Indeed, through a digital channel, the i-th agent can only send to its neighbors symbolic data: using this data, the neighbors of the i-th agent can build an estimate of the i-th agent's state. We denote this estimate by $\hat{x}_i(t)$, and let $\hat{x}(t) = [\hat{x}_1(t), \dots, \hat{x}_N(t)]^*$. Here, the estimate is simply the received symbol, computed via the *deterministic quantizer* (1.11) or the *probabilistic quantizer* (1.13).

We introduce *two updating rules* for the states, which use quantized information. In the first strategy, if (i,j) is the edge selected at the t-th iteration, i and j, in order to update its state, use only the estimates of their states, as follows,

$$\begin{aligned} x_i(t+1) &= \frac{1}{2}\hat{x}_i(t) + \frac{1}{2}\hat{x}_j(t) \\ x_j(t+1) &= \frac{1}{2}\hat{x}_j(t) + \frac{1}{2}\hat{x}_i(t), \end{aligned} \tag{3.1}$$

or, equivalently in vector form, by recalling the definition of $P(t)$ in (1.9),

$$x(t+1) = P(t)\hat{x}(t). \tag{3.2}$$

To define the second strategy, we remark that (1.7) can be written as

$$\begin{aligned} x_i(t+1) &= x_i(t) - \frac{1}{2}x_i(t) + \frac{1}{2}x_j(t) \\ x_j(t+1) &= x_j(t) - \frac{1}{2}x_j(t) + \frac{1}{2}x_i(t). \end{aligned}$$

We then propose the following updating rule, where the agents use also perfect information regarding their own states,

$$\begin{aligned} x_i(t+1) &= x_i(t) - \frac{1}{2}\hat{x}_i(t) + \frac{1}{2}\hat{x}_j(t) \\ x_j(t+1) &= x_j(t) - \frac{1}{2}\hat{x}_j(t) + \frac{1}{2}\hat{x}_i(t), \end{aligned} \tag{3.3}$$

or, equivalently in vector form,

$$x(t+1) = x(t) + (P(t) - I)\hat{x}(t). \tag{3.4}$$

We call the law (3.2) *globally quantized* and the law (3.4) *partially quantized*.

It is important to remark that, since $P(t)$ is a doubly stochastic matrix for all $t \geq 0$, Proposition 1.8 can be invoked to assure that the *partially quantized* law (3.3) preserves the initial state average.

We proceed with our analysis of these two rules by assuming first that $\hat{x}_i(t) = q_d(x_i(t))$, i.e., the information transmitted is quantized by means of deterministic quantizer, and then by assuming that $\hat{x}_i(t) = q_p(x_i(t))$, i.e., the information transmitted is quantized by means of probabilistic quantizer.

3.2 Gossip algorithms with deterministic quantizers

In this section we assume that the information exchanged between the agents is quantized by means of the deterministic quantizer q_d described in (1.11), namely $\hat{x}_i(t) = q_d(x_i(t))$. In this context, we separately analyze the partially and globally quantized strategies, starting from the first one.

3.2.1 Partially quantized strategy

Consider the partially quantized strategy

$$\begin{aligned} x_i(t+1) &= x_i(t) - \frac{1}{2}q_d(x_i(t)) + \frac{1}{2}q_d(x_j(t)) \\ x_j(t+1) &= x_j(t) - \frac{1}{2}q_d(x_j(t)) + \frac{1}{2}q_d(x_i(t)), \end{aligned} \tag{3.5}$$

Recalling Definition 1.2, we let

$$y(t) = \left(I - \frac{1}{N}\mathbf{1}\mathbf{1}^*\right)x(t), \tag{3.6}$$

and remark that $y(t) = x(t) - \frac{1}{N}\mathbf{1}\mathbf{1}^*x(0)$. As well, we let

$$d(t) = \frac{1}{\sqrt{N}}\|y(t)\|_2, \tag{3.7}$$

Such quantity represents the normalized distance of the state $x(t)$ from the average of the states.

As an example we report in Figure 3.1 the result of simulations relative to a connected random geometric graph. Such graph has been drawn placing $N = 50$ nodes uniformly at random inside the unit square and connecting two nodes whenever the distance between them is less that $R = 0.3$. The initial condition $x_i(0)$ is randomly chosen inside the interval $[-100, 100]$ for all $1 \leq i \leq N$. Note that $d(t)$ does not converge to 0, meaning that the average consensus is not reached, but the values get very close to the initial average.

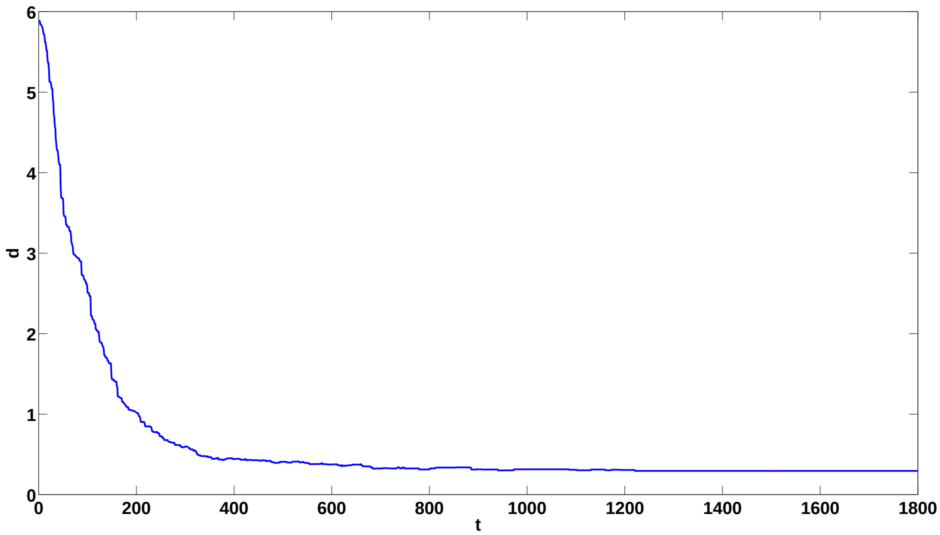

Figure 3.1: Behavior of d for a connected random graph with $N = 50$ in case of deterministic quantizers and of partially quantized strategy.

In the following we give a proof of this fact, quantifying the distance from consensus that the states of the agents asymptotically achieve. This will be done exploiting a natural *symbolic dynamics* interpretation of the states dynamics, similarly to what we have done in Appendix 2.A, and extending the results in [55]. We define $n_i(t) = \lfloor 2x_i(t) \rfloor$ for all $i \in V$, and let $n(t) = [n_1(t), \dots, n_N(t)]^*$. Simple properties of floor and ceiling operators, together with Lemma 2.12, allow us to remark that $q_d(x_i(t)) = \left\lceil \frac{n_i(t)}{2} \right\rceil$ and that

$$\begin{aligned} x_i(t+1) &= x_i(t) - \frac{1}{2} q_d(x_i(t)) + \frac{1}{2} q_d(x_j(t)) \\ \lfloor 2x_i(t+1) \rfloor &= \lfloor 2x_i(t) \rfloor - q_d(x_i(t)) + q_d(x_j(t)), \end{aligned}$$

from which we can obtain that

$$\begin{aligned} n_i(t+1) = n_i(t) - \left\lceil \frac{n_i(t)}{2} \right\rceil + \left\lceil \frac{n_j(t)}{2} \right\rceil \\ = \left\lfloor \frac{n_i(t)}{2} \right\rfloor + \left\lceil \frac{n_j(t)}{2} \right\rceil . \end{aligned}$$

We have thus found an iterative system involving only the symbolic signals $n_i(t)$. When the edge (i, j) is selected, i and j adjourn their states following the pair dynamics

$$(n_i(t+1), n_j(t+1)) = g_2(n_i(t), n_j(t)) \tag{3.8}$$

where $g : \mathbb{Z} \times Z \to \mathbb{Z} \times \mathbb{Z}$ is

$$g_2(h,k) = \left(\left\lfloor \frac{h}{2} \right\rfloor + \left\lceil \frac{k}{2} \right\rceil, \left\lfloor \frac{k}{2} \right\rfloor + \left\lceil \frac{h}{2} \right\rceil \right).$$

Notice that g is symmetric in the arguments, in the sense that if $g_2(h,k) = (\eta, \chi)$, then $g_2(k,h) = (\chi, \eta)$. The analysis of the evolution of (3.8) then allows us to obtain information about the asymptotics of $x_i(t)$, since $n_i(t) = \lfloor 2x_i(t) \rfloor$. Before stating the main result regarding the convergence properties of (3.8), we define the following set

$$\mathcal{R} = \left\{ r \in \mathbb{Z}^N : \exists \alpha \in \mathbb{Z} \text{ s. t. } r - \alpha \mathbf{1} \in \{0,1\}^N \right\}. \tag{3.9}$$

We have the following result.

Theorem 3.1 *Let $n(t)$ evolve according to (3.8). For every fixed initial condition $n(0) \in \mathbb{Z}$, almost surely there exists $T_{con} \in \mathbb{N}$ such that $n(t) \in \mathcal{R}$ for all $t \geq T_{con}$.*

Proof: The proof is based [72] on verifying the following three facts:

(i) the set $\mathcal{R}$, defined in (3.9), is an invariant subset for the evolution described by (3.8);

(ii) $n(t)$ is a Markov process on a finite number of states;

(iii) there is a positive probability for $n(t)$ to reach a state in $\mathcal{R}$ in a finite number of steps.

Let us now check them in order.

(i) Let $h \in \mathbb{Z}$. Observe that

$$g_2(h, h+1) = \begin{cases} (h+1, h) & \text{if } h \text{ is even} \\ (h, h+1) & \text{if } h \text{ is odd} \end{cases}$$

This implies that $\mathcal{R}$ is an invariant subset for the dynamics described by (3.8).

(ii) Markovianity immediately follows from the fact that subsequent random choices of the edges are independent. We prove now that the states are finite. To this aim let $(h', k') = g_2(h,k)$. By the structure of g, it is easy to see that

$$\max\{h', k'\} \leq \max\{h, k\} \qquad \min\{h, k\} \leq \min\{h', k'\}.$$

Therefore we have that $m(n(t)) \geq m(n(0))$ and $M(n(t)) \leq M(n(0))$. This implies that the cardinality of the set of the states is upper bounded by $(M(n(0)) - m(n(0)) + 1)^N$.

(iii) First define

$$m(t) = \min_{1 \leq i \leq N} n_i(t) \tag{3.10}$$

$$M(t) = \max_{1 \leq i \leq N} n_i(t), \tag{3.11}$$

and,

$$D(t) = M(t) - m(t).$$

The proof of (iii) is based on the following strong result about the monotonicity of $D(t)$: if $D(t) \geq 2$, then there exists $\tau \in \mathbb{N}$ such that

$$\mathbb{P}[D(t+\tau) < D(t)] > 0. \tag{3.12}$$

Now we prove (3.12).

Let $\mathcal{I}(t) = \{j \in V \text{ s.t. } n_j(t) = m(t)\}$. We start by proving that $|\mathcal{I}(t)|$, i.e., the cardinality of $\mathcal{I}(t)$, does not increase and that, if $D(t) \geq 2$, then there is a positive probability that it decreases within a finite number of time steps. Notice first that, for $h, k \in \mathbb{Z}$, $g_2(h+2, k+2) = g_2(h,k) + 2$. Hence, by an appropriate translation of the initial condition, we can always restrict ourselves to the case $m(t) \in \{0, 1\}$, which of course is easier to handle.
Case $m(t) = 0$. In this case it is possible for a nonzero state to decrease to 0, but only in the case of a swap between 0 and 1. This assures that $|\mathcal{I}(t)|$ is nonincreasing. Let $\mathcal{S}(t)$ denote the set of nodes which have value $m(t)+2$ or larger. Since $D(t) \geq 2$ then $\mathcal{S}(t)$ is non empty at time t. Now let $(v_1, v_2, \ldots, v_{p-1}, v_p)$ be a shortest path between $\mathcal{I}(t)$ and $\mathcal{S}(t)$. Such a path exists since $\mathcal{G}$ is connected. Note that $v_1 \in \mathcal{I}(t)$ and $v_p \in \mathcal{S}(t)$ and that $\{v_2, \ldots, v_{p-1}\}$ could be an empty set; in this case a shortest path between $\mathcal{I}(t)$ and $\mathcal{S}(t)$ has length 1. Moreover note also that all the nodes in the path except v_1 and v_p have value 1 at time t, otherwise $(v_1, v_2, \ldots, v_{p-1}, v_p)$ would not be a shortest path. Since each edge of the communication graph has a positive probability of being selected in any time, there is also a positive probability that in the $p-1$ time units following t the edges of this path are selected sequentially, starting with the edge (v_1, v_2). At the last step of this sequence we have that the values of v_{p-1} and v_p are updated. By observing again, that the pair of value $(0, 1)$ is transformed by (3.8) into the pair $(1, 0)$ we have that the value of v_{p-1}, when the edge (v_{p-1}, v_p) is selected, is equal to 0. This update, for the form of (3.8), causes the value of both nodes to be strictly greater than 0. Therefore, this proves that $|\mathcal{I}(t+p-1)| < |\mathcal{I}(t)|$ with positive probability. Clearly, if $|\mathcal{I}(t)| = 1$ then we have also that $D(t+p-1) < D(t)$ with positive probability.
Case $m(t) = 1$. In this case no state can decrease to 1, and thus $|\mathcal{I}(t)|$ is not increasing. Let $\mathcal{I}(t)$, $\mathcal{S}(t)$ and $(v_1, v_2, \ldots, v_{p-1}, v_p)$ be defined as in the previous case. Obviously in this case all the nodes $v_2, \ldots, v_{p-1}$ in the path have value equal to 2. Moreover observe that also the sequence of edges (v_{p-1}, v_p), (v_{p-2}, v_{p-1}), $\ldots$,(v_2, v_3), (v_1, v_2) has positive probability of being selected in the $p-1$ time units following t. At the last step of this sequence of edges, the values of v_1 and v_2 are updated. Clearly the value of v_1 is equal to 1. Since the value of v_p at time t is greater or equal to 3, and since the pair $(2, 3)$ is transformed by (3.8) into $(3, 2)$, we have that the value of v_2 when the edge (v_1, v_2) is selected, is greater or equal to 3. This update, for (3.8), causes the value of both nodes to be strictly greater than 1. Hence $|\mathcal{I}(t+p-1)| < |\mathcal{I}(t)|$ with positive probability. Again, if $|\mathcal{I}(t)| = 1$ then we have also that $D(t+p-1) < D(t)$ with positive probability.

Consider now the following sequence of times $t_0 = t, t_1, t_2, \ldots$. For each $i \geq 0$, if $|\mathcal{I}(t)| > 1$, then we let t_{i+1} to be the first time for which there is a positive probability that $|\mathcal{I}(t_{i+1})| < |\mathcal{I}(t_i)|$. Let now $k \in \mathbb{N}$ be such that $|\mathcal{I}(t_k)| = 1$. Then we have that $D(t_{k+1}) < D(t_k)$. This ensures the validity of (3.12).

The proof of the fact (iii) follows directly from (3.12). Indeed, let $\bar{n} \notin \mathcal{R}$, then, from a repeated application of (3.12) it follows that, there exists a path connecting $\bar{n}$ to a state $\bar{n}' = [\bar{n}'_1, \ldots, \bar{n}'_N]$, such that $\max\{\bar{n}'_1, \ldots, \bar{n}'_N\} - \min\{\bar{n}'_1, \ldots, \bar{n}'_N\} < 2$, that is, $\bar{n}' \in \mathcal{R}$.

This proves the thesis. $\square$

We can go back to the original system, and prove the following result.

Corollary 3.2 *Consider the algorithm (3.5). Then, almost surely, there exists $T_{con} \in \mathbb{N}$ such that*

$$|x_i(t) - x_j(t)| \leq 1 \quad \forall\, i, j \quad \forall\, t \geq T_{con}. \tag{3.13}$$

and hence,

$$\|x(t) - x_{\text{ave}}\mathbf{1}\|_\infty \leq 1.$$

Proof: The proof is an immediate consequence of Theorem 3.1 and of the relation $n_i(t) = \lfloor 2x_i(t) \rfloor$, which assure that the states belong to two consecutive quantization bins. $\square$

Remark 3.1 Note that the quantizer q_T, instead of q_d, would induce a dynamics

$$g_6(h,k) = \left(\left\lfloor \frac{k}{2} \right\rfloor + \left\lceil \frac{h}{2} \right\rceil, \left\lfloor \frac{h}{2} \right\rfloor + \left\lceil \frac{k}{2} \right\rceil \right),$$

which can be studied similarly.

Remark 3.2 It is worth noting that Theorem 3.1 is an extension of Lemma 3 and Theorem 1 in [55]. In [55] the authors introduced a class of *quantized gossip algorithms*, satisfying the following assumptions. Let (i, j) be the edge selected at time t and let $n_i(t)$ and $n_j(t)$ the values at time t of node i and of node j respectively. If $n_i(t) = n_j(t)$ then $n_i(t+1) = n_i(t)$ and $n_j(t+1) = n_j(t)$. Otherwise, defined $D_{ij} = |n_i(t) - n_j(t)|$, the method used to update the values has to satisfy the following three properties:

(P1) $n_i(t+1) + n_j(t+1) = n_i(t) + n_j(t)$,

(P2) if $D_{ij}(t) > 1$ then $D_{ij}(t+1) < D_{ij}(t)$, and

(P3) if $D_{ij}(t) = 1$ and (without loss of generality) $n_i(t) < n_j(t)$, then $n_i(t+1) = n_j(t)$ and $n_j(t+1) = n_i(t)$. Such update is called *swap*.

Now we substitute the property $(P3)$ either with the property

(P3') if $D_{ij}(t) = 1$ and (without loss of generality) $n_i(t) < n_j(t)$, then, if $n_i(t)$ is odd, then $n_i(t+1) = n_j(t)$ and $n_j(t+1) = n_i(t)$, otherwise if $n_i(t)$ is even then $n_i(t+1) = n_i(t)$ and $n_j(t+1) = n_j(t)$

or with the property

(P3") if $D_{ij}(t) = 1$ and (without loss of generality) $n_i(t) < n_j(t)$, then, if $n_i(t)$ is even then $n_i(t+1) = n_j(t)$ and $n_j(t+1) = n_i(t)$, otherwise if $n_i(t)$ is odd then $n_i(t+1) = n_i(t)$ and $n_j(t+1) = n_j(t)$.

If we consider the class of algorithms satisfying (P1), (P2), (P3') or satisfying (P1), (P2), (P3"), it is possible to prove that Lemma 3 and Theorem 1 stated in [55] hold true also for this class. The proofs are analogous to that of Theorem 3.1 provided above. Moreover it is easy to see that the algorithm (3.8) satisfies the properties (P1), (P2), (P3'). This represents an alternative way to prove Theorem 3.1.

3.2.2 Globally quantized strategy

In this subsection we consider the globally quantized strategy

$$\begin{aligned} x_i(t+1) &= \frac{1}{2}q_d(x_i(t)) + \frac{1}{2}q_d(x_j(t)) \\ x_j(t+1) &= \frac{1}{2}q_d(x_j(t)) + \frac{1}{2}q_d(x_i(t)), \end{aligned} \tag{3.14}$$

We underline immediately that the fact that (3.14) uses only quantized information and not perfect information combined with quantized information as in (3.5) makes the analysis of (3.5) slightly easier than the analysis of (3.14).
Remarkably, we show in this subsection that the law (3.5) drives, almost surely, the systems to exact consensus at an integer value. Unfortunately, the initial average of states is not preserved in general. Again, the analysis of this algorithm can be performed efficiently by means of the *symbolic dynamics* previously introduced.
Let again $n_i(t) = \lfloor 2x_i(t) \rfloor$ for all $i \in V$. From (3.14) and the fact that $q_d(x_i(t)) = \left\lceil \frac{n_i(t)}{2} \right\rceil$ we obtain

$$(n_i(t+1), n_j(t+1)) = (g_3(n_i(t), n_j(t)), g_3(n_i(t), n_j(t))) \tag{3.15}$$

where $g_3 : \mathbb{Z} \times \mathbb{Z} \to \mathbb{Z}$ is defined as

$$g_3(h,k) = \left\lceil \frac{h}{2} \right\rceil + \left\lceil \frac{k}{2} \right\rceil.$$

Define

$$\mathcal{A} = \{y \in \mathbb{Z}^N : \exists \alpha \in \mathbb{Z} \text{ such that } y = 2\alpha\mathbf{1}\}. \tag{3.16}$$

We have the following result.

Theorem 3.3 *Let $n(t)$ evolve according to (3.15). For every fixed initial condition $n(0)$, almost surely there exists $T_{con} \in \mathbb{N}$ such that $n(t) \in \mathcal{A}$ for all $t \geq T_{con}$.*

Proof: As for the proof of Theorem 3.1, it will be sufficient to verify the following three facts:

(i) each element in the set $\mathcal{A}$ is invariant for the evolution described by (3.15);

(ii) $n(t)$ is a Markov process on a finite number of states;

(iii) there is a positive probability for $n(t)$ to reach a state in $\mathcal{A}$ in a finite number of steps.

Let us now check them in order.

(i) is trivial.

(ii) Markovianity immediately follows from the fact that subsequent random choices of the edges are independent. To prove that the states are finite, define $m(t)$ and $M(t)$ as in (3.10) and (3.11). Let $p, q \in \mathbb{Z}$ with $p \leq q$. Then, from the structure of g_3 we have that $p \leq g_3(p, q) \leq q + r_q$ where r_q denotes the remainder in the division of q by 2. It follows that

$$m(0) \leq n_i(t) \leq M(0) + r_{M(0)} \;\forall i \in V \;\; \forall t \geq 0\,. \tag{3.17}$$

This yields (ii).

(iii) Let us fix $t = t_0$, and assume that $n(t_0) \notin \mathcal{A}$. We prove that there exists $\tau \in \mathbb{N}$ such that $\mathbb{P}\left[n(t_0 + \tau) \in \mathcal{A}\right] > 0$. We start by observing that, from the assumption of having a connected graph, there exists $(h, k) \in \mathcal{E}$ such that $n_h(t_0) = m(t_0)$, $n_k(t_0) = q$ and $g_2(m(t_0), q) > m(t_0)$. Indeed, two cases are given when $n(t_0) \notin \mathcal{A}$.

- If $m(t_0) < M(t_0)$, then it suffices to consider an edge (h, k) such that $n_h(t_0) = m(t_0)$ and $n_k(t_0) = q > m(t_0)$, which gives $g_3(m(t_0), p) > m(t_0)$. Note that such an edge exists from the hypothesis of having a connected graph;
- if $m(t_0) = M(t_0)$, necessarily we have that $m(t_0)$ and $M(t_0)$ are odd; then $g_2(m(t_0), m(t_0)) > m(t_0)$.

We define now $\mathcal{I}_a(t) = \{i \in \mathcal{V} : n_i(t) = a\}$. The above discussion implies that $|\mathcal{I}_{m(t_0)}(t_0+1)| < |\mathcal{I}_{m(t_0)}(t_0)|$ with the positive probability of choosing the edge (h, k) and hence that there is also a positive probability that at some finite time $t' > t_0$, $|\mathcal{I}_{m(t_0)}(t')| = 0$, that is $m(t') > m(t_0)$. Iterating this argument and recalling that (see (3.17)) $M(t) \leq M(t_0) + r_{M(t_0)}$ for all $t \geq t_0$, it follows that there exists $\tau \in \mathbb{N}$ such that $\mathbb{P}\left[n(t_0 + \tau) \in \mathcal{A}\right] > 0$.

This proves the thesis. $\square$

We can now go back to the original system. The following corollary follows immediately from the definition of $n(t)$.

Corollary 3.4 *Let $x(t)$ evolve according to (3.14). Then almost surely there exists $T_{con} \in \mathbb{N}$ and $\alpha \in \mathbb{Z}$ such that $x_i(t) = \alpha$ for all $i \in \mathcal{V}$ and for all $t \geq T_{con}$.*

We have already underlined the fact that this strategy does not preserve the initial average, in general. Providing some probabilistic estimation of the distance of the consensus point from the initial average is a challenging problem: we limits our analysis to the following simulation. In Figure 3.2 we plot the variable z that is defined as follows. In the globally quantized strategy we have that, almost surely $\lim_{t\to\infty} = \alpha\mathbf{1}$ for some random integer α. Let $z = |\alpha - 1/N\mathbf{1}^* x(0)|$. In words, z represents the distance from the consensus point to which the globally quantized strategy leads the systems and the average of the initial condition. We have depicted the value of z for a family of random geometric graphs [76] of increasing size from $N = 10$ up to $N = 80$. The initial condition $x_i(0)$ is chosen randomly inside the interval $[-100, 100]$ for all $1 \leq i \leq N$. Moreover for each N, z is

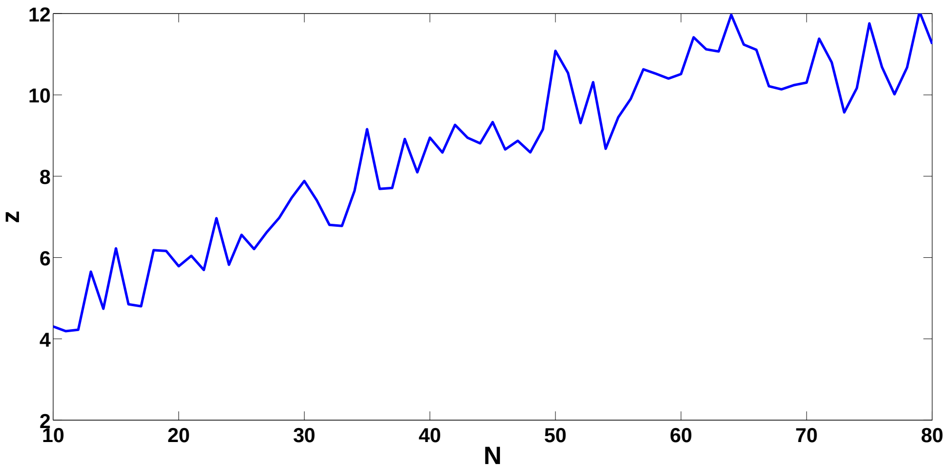

Figure 3.2: Behavior of z for a family of random geometric graphs in case of deterministic quantizers and of globally quantized strategy.

computed as the mean of 100 trials. We can see that the value of z is increasing in N and assumes values that are not negligible with respect to the quantization step size.

Remark 3.3 Using the quantizer q_T instead of q_d would induce a dynamics

$$g_7(h,k) = \left\lfloor \frac{h}{2} \right\rfloor + \left\lceil \frac{k}{2} \right\rceil,$$

which can be studied in an analogous way.

3.2.3 Speed of convergence

Providing some insights on the speed of convergence of (3.5) and of (3.14) is quite hard in general. In Figure 3.3 and Figure 3.4 we report, respectively, a comparison between the partially quantized strategy (3.5) and the gossip algorithm with exchange of perfect information (1.7) and between the globally quantized strategy (3.14) and again the gossip algorithm with exchange of perfect information (1.7). The simulations are made on the same random geometric graphs considered in Figure 3.1, and the initial conditions are randomly chosen inside the interval $[-100, 100]$.

For both strategies we plotted the behavior of the variable $d(t)$ defined in (3.7).

From the Figure 3.3 and Figure 3.4 we can infer that the speed of convergence toward the steady state of the quantized strategies (3.14) and (3.5) is similar to the one of the gossip algorithm with perfect exchange of information. This numerical evidence is not completely understood yet, but some interesting preliminary results appear in [40].

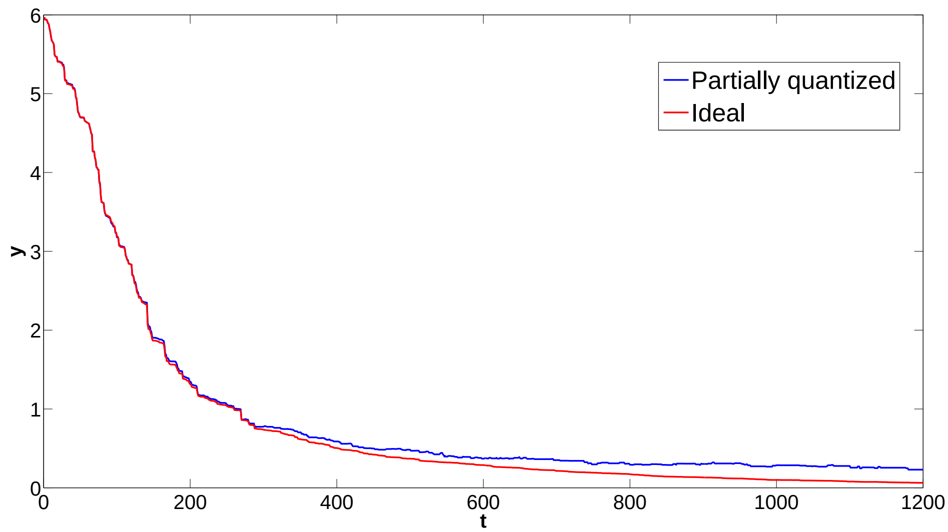

Figure 3.3: Behavior of d, when using the *partially quantized* strategy, for a connected random geometric graph with $N = 50$. Note that since the *partially quantized* strategy does not converge to a consensus, $d(t)$ does not go to 0.

Remark 3.4 If, depending on the application, one can not relax the convergence requirement, we could suggest the following heuristic solution to the consensus problem, which combines the positive features of both strategies,

$$x(t+1) = Pq_d(x(t)) + \varepsilon(t)(x(t) - q_d(x(t))),$$

where $\varepsilon(t),\ t \geq 0$, is a nonnegative sequence such that $\varepsilon(t) \leq 1,\ \forall\, t \geq 0$ and $\lim_{t\to\infty} \varepsilon(t) = 0$.

3.3 Gossip algorithms with probabilistic quantizers

In this section we assume that the information exchanged between the systems is quantized by means of the probabilistic quantizer q_p described in (1.13), namely $\hat{x}_i(t) = q_p(x_i(t))$. We recall the statistics of q_p, as illustrated in Lemma 1.7. Moreover, we make the following natural assumption

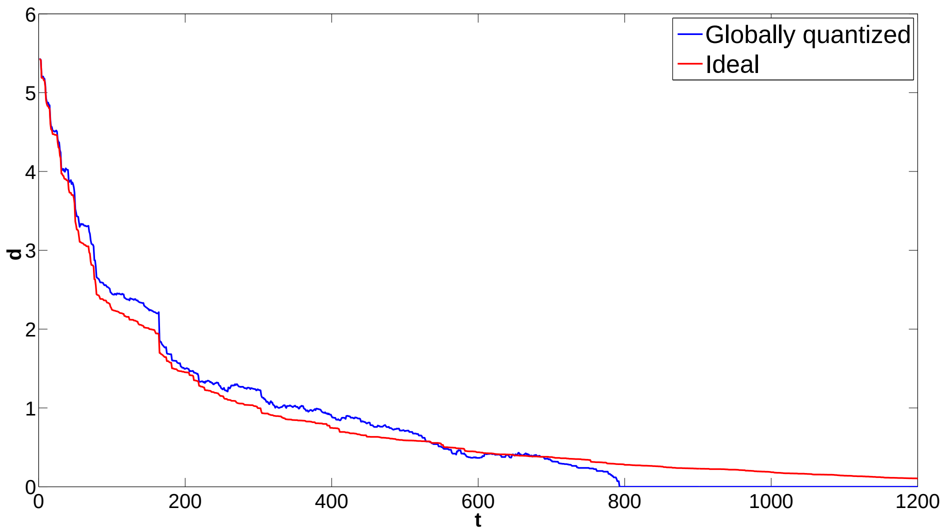

Figure 3.4: Behavior of d, when using the *globally quantized* strategy, for a connected random geometric graph with $N = 50$. In this case, accordingly to the theoretical result stated in Corollary 3.4, $d(t)$ tends to 0.

Assumption 3.2 *Given the values $x_i(t)$ for all $i \in V$, the random variables $q_p(x_i(t))$, as i varies, form an independent set. Moreover, for every $i \neq j$, given $x_i(t)$, $q_p(x_i(t))$ is independent from $x_j(t)$.*

As before, we will now separately analyze the partially and globally quantized strategies.

3.3.1 Partially quantized strategy

The algorithm for partially quantized strategy, when the edge (i, j) is chosen, can be written as

$$\begin{aligned} x_i(t+1) &= x_i(t) - \frac{1}{2}q_p(x_i(t)) + \frac{1}{2}q_p(x_j(t)) \\ x_j(t+1) &= x_j(t) - \frac{1}{2}q_p(x_j(t)) + \frac{1}{2}q_p(x_i(t)). \end{aligned} \tag{3.18}$$

Similarly to the partially quantized strategy via deterministic quantizers (3.5), also (3.18) does not reach the consensus in general. Again we report a simulation showing this fact. In Figure 3.5 the behavior of the quantity $d(t)$, defined in (3.7), is depicted for the same connected random geometric graph considered in Figure 3.1. Note that the quantity $d(t)$ stays visibly away from 0, meaning that the average consensus is not reached.

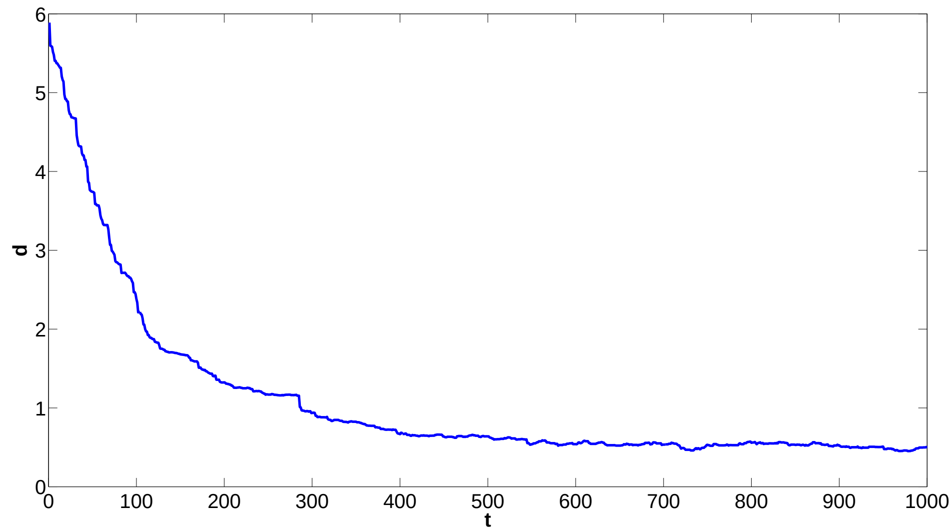

Figure 3.5: Behavior of d for a connected random geometric graph with $N = 50$.

The analysis of (3.18) is more complicate than for the corresponding law (3.5). This is mainly due to the lack of convexity properties which were used in the analysis of (3.5). The following example shows this type of difficulty.

Example 3.1 Consider (3.5) and assume that the edge (i, j) has been selected at time t. Without loss of generality assume that $x_i(t) \leq x_j(t)$. Then, by convexity arguments, we have that $\lfloor x_i(t) \rfloor \leq x_i(t+1), x_j(t+1) \leq \lceil x_j(t) \rceil$. This is no longer true for (3.18). As a numerical example assume that $x_i(t) = 3.4$ and $x_j(t) = 3.6$. Then with probability $1/4$ we will have that $q_p(x_i(t)) = 4$ and $q_p(x_j(t)) = 3$. In this case, by (3.18), we have that $x_i(t+1) = 2.9$ and that $x_j(t+1) = 4.1$. Hence, $x_i(t+1), x_j(t+1)$ do not belong to the interval $[\lfloor x_i(t) \rfloor, \lceil x_j(t) \rceil]$.

For this reason, we do not develop a symbolic analysis for this algorithm, and we do not prove convergence in finite time. By simulations we can see that (3.18) does not drive the states of the systems inside the same bin of quantization, as the corresponding strategy (3.5) using deterministic quantizers. In Figure 3.6, we depict the behavior of the quantity

$$s(t) = \max_{1 \leq i,j \leq N} |x_i(t) - x_j(t)|.$$

for the same random geometric graph considered in Figure 3.5. In this simulation we assume that the initial condition $x_i(0)$ is randomly chosen inside the interval $[-10, 10]$. Note that s asymptotically oscillates around 2. Interesting results on (3.18), in terms of both the asymptotic distance from the initial average and the speed of convergence, can be provided by a mean-square analysis. In the sequel of this subsection, we assume that the initial condition $x(0)$ satisfies the following condition.

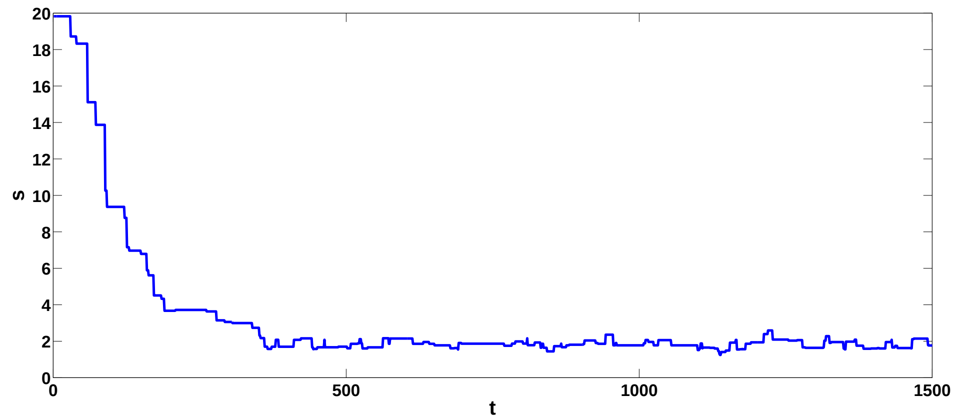

Figure 3.6: Behavior of s for a connected random geometric graph with $N = 50$.

Assumption 3.3 *The initial condition $x(0)$ is a random variable such that $\mathbb{E}[x(0)] = 0$ and $\mathbb{E}[x(0)x^*(0)] = \sigma_0^2 I$ for some $\sigma_0^2 > 0$.*

We start by observing that (3.18) can be rewritten as

$$x(t+1) = P(t)x(t) + (P(t) - I)\,(q_p(x(t)) - x(t)) \tag{3.19}$$

Define

$$e(t) = q_p(x(t)) - x(t).$$

the quantization error and recall the definition of $y(t)$ given in (3.6). From (3.19), using the fact that $P(t)$ is symmetric and stochastic, we easily obtain the following recursive relation in terms of the variables $e(t)$ and $y(t)$:

$$y(t+1) = P(t)y(t) + (P(t) - I)e(t). \tag{3.20}$$

In order to perform an asymptotic analysis of (3.20) it is convenient to introduce the following matrices. Let

$$\Sigma_{yy}(t) = \mathbb{E}\left[y(t)y^*(t)\right], \quad \Sigma_{ee}(t) = \mathbb{E}\left[e(t)e(t)^*\right], \quad \Sigma_{ye}(t) = \mathbb{E}\left[y(t)e(t)^*\right].$$

Equation (3.20) leads to the following recursive equation in terms of the above matrices

$$\begin{aligned}\Sigma_{yy}(t+1) = \mathbb{E}\left[P(t)\Sigma_{yy}(t)P(t)\right] + \mathbb{E}\left[P(t)\Sigma_{ye}(t)(P(t)-I)\right] + \\ + \mathbb{E}\left[(P(t)-I)\Sigma_{ye}^* P(t)\right] + (P(t)-I)\,\Sigma_{ee}(t)\,(P(t)-I)\,.\end{aligned} \tag{3.21}$$

From the fact that $x(0)$ is a random variable satisfying Assumption 3.3, it immediately follows that

$$\Sigma_{yy}(0) = \sigma_0^2\left(I - N^{-1}\,\mathbf{1}\mathbf{1}^*\right). \tag{3.22}$$

The following proposition states some correlation properties of the variables y and e.

Proposition 3.5 *Consider the variables $y(t)$ and $e(t)$ above defined. Then*

$$\mathbb{E}\left[e(t)\right] = 0 \qquad \textit{and} \qquad \Sigma_{ee}(t) = \operatorname{diag}\left\{\sigma_1^2(t), \dots, \sigma_N^2(t)\right\} \tag{3.23}$$

where $\sigma_i^2(t) = \mathbb{E}\left[e_i^2(t)\right]$ is such that $\sigma_i^2(t) \leq 1/4$ for all $1 \leq i \leq N$ and for all $t \geq 0$. Moreover

$$\Sigma_{ye}(t) = 0, \tag{3.24}$$

for all $t \geq 0$.

Proof: Using Lemma 1.7, we have that

$$\begin{aligned}
\mathbb{E}\left[e_i(t)\right] &= \mathbb{E}\left[\,\mathbb{E}\left[q_p(x_i(t)) - x_i(t)|x_i(t)\right]\right]\\
&= \mathbb{E}\left[\,\mathbb{E}\left[q_p(x_i(t))|x_i(t)\right] - x_i(t)\right]\\
&= \mathbb{E}\left[x_i(t) - x_i(t)\right]\\
&= 0.
\end{aligned} \tag{3.25}$$

Moreover, for $i \neq j$, using Assumption 3.2,

$$\begin{aligned}
\mathbb{E}\left[e_i(t)e_j(t)\right] &= \mathbb{E}\left[\mathbb{E}\left[e_i(t)e_j(t)\,|\,x_i(t),\, x_j(t)\right]\right]\\
&= \mathbb{E}\left[\mathbb{E}\left[e_i(t)\,|\,x_i(t),\, x_j(t)\right]\mathbb{E}\left[e_j(t)\,|\,x_i(t),\, x_j(t)\right]\right]\\
&= \mathbb{E}\left[\mathbb{E}\left[e_i(t)\,|\,x_i(t)\right]\mathbb{E}\left[e_j(t)\,|\,x_j(t)\right]\right]\\
&= 0
\end{aligned} \tag{3.26}$$

If $i = j$, using again Lemma 1.7, we have that

$$\begin{aligned}
\mathbb{E}\left[e_i^2(t)\right] &= \mathbb{E}\left[\left(q_p(x_i(t)) - x_i(t)\right)^2\right]\\
&= \mathbb{E}\left[\,\mathbb{E}\left[\left(q_p(x_i(t)) - x_i(t)\right)^2|x_i(t)\right]\right]\\
&\leq \mathbb{E}\left[\frac{1}{4}\right]\\
&= \frac{1}{4}
\end{aligned} \tag{3.27}$$

An argument similar than the one above used to prove that $\mathbb{E}\left[e_i(t)e_j(t)\right] = 0$ allows to prove that $\mathbb{E}\left[x_i(t)e_j(t)\right] = 0$ for any $i \neq j$. This easily yields (3.24). □

From the above properties we have that (3.21) can be rewritten as

$$\Sigma_{yy}(t+1) = \mathbb{E}\left[P(t)\Sigma_{yy}(t)P(t)\right] + \mathbb{E}\left[\left(P(t) - I\right)\Sigma_{ee}(t)\left(P(t) - I\right)\right]. \tag{3.28}$$

To estimate the asymptotic distance from the initial average, we introduce the cost function

$$J(W) = \limsup_{t\to\infty} \sqrt{\frac{1}{N}\mathbb{E}[\|y(t)\|^2]}. \tag{3.29}$$

The cost depends on the selection probabilities W, and, thanks to the above definitions, can be computed as

$$J(W) = \limsup_{t\to\infty} \sqrt{\frac{1}{N}\operatorname{tr}\{\Sigma_{yy}(t)\}}. \tag{3.30}$$

We can rewrite the evolution law (3.28) as

$$\Sigma_{yy}(t+1) = \mathcal{A}(\Sigma_{yy}(t)) + \mathcal{B}(\Sigma_{ee}(t)),$$

where $\mathcal{A}$ and $\mathcal{B}$ are linear operators from $\mathbb{R}^{N\times N}$ to itself. Namely, given a matrix M, $\mathcal{A}(M) = \mathbb{E}\left[P(t)MP(t)\right]$ and $\mathcal{B}(M) = \mathbb{E}\left[(P(t)-I)\, M\, (P(t)-I)\right]$.

It is useful to remark that $\mathcal{A}$ is actually the evolution on Σ_{yy} for the gossip algorithm [9], in the absence of quantization error, while $\mathcal{B}$ can be regarded as a disturbance due to the quantization error. From [37], we know that in the case of no quantization the system converges almost surely to consensus. This implies that $\mathcal{A}$ is an asymptotically stable operator when restricted to the subspace $\mathcal{S} = \{M \in \mathbb{R}^{N\times N} : \mathbf{1}^*M\mathbf{1} = 0\}$. Since $\mathbf{1}^*\mathcal{B}(M)\mathbf{1} = 0$ for any matrix M and $\Sigma_{yy}(0) \in \mathcal{S}$, we have that $\Sigma_{yy}(t) \in \mathcal{S}$ for all $t \geq 0$. As a consequence $\Sigma_{yy}(t)$ converges for $t \to +\infty$ and in the definition of $J(W)$ in (3.30) $\limsup$ can be replaced by $\lim$.

It is actually a general fact that, in most cases, systems with quantization can be regarded as disturbed versions of non-quantized systems. This approach has been taken in [40], to show that the speed of convergence of gossip consensus algorithms is essentially the same with or without quantization, as long as the states are far from consensus.

Providing an expression for $J(W)$ is quite hard in general. We then try to simplify the problem by introducing the following auxiliary system

$$\bar{\Sigma}(t+1) = \mathbb{E}\left[P(t)\bar{\Sigma}(t)P(t)\right] + \frac{1}{4}\mathbb{E}\left[(P(t)-I)^2\right], \tag{3.31}$$

where $\bar{\Sigma}(0) = \Sigma_{yy}(0)$, and the following cost function

$$\bar{J} = \limsup_{t\to\infty} \sqrt{\frac{1}{N}\operatorname{tr}\{\bar{\Sigma}(t)\}}.$$

In principle, $\bar{J}$ should depend on W, too. However, we are going to prove that this is not the case. We have the following comparison result.

Proposition 3.6 *Consider the cost functions $J(W)$ and $\bar{J}$. We have that*

$$J(W) \leq \bar{J}.$$

Proof: To prove the statement we show, by induction on t, that $\bar{\Sigma}(t) \geq \Sigma_{yy}(t)$ for all $t \geq 0$, where the inequality is meant in matricial sense, that is, $\bar{\Sigma}(t) - \Sigma_{yy}(t)$ is a semidefinite positive matrix.

Since $\bar{\Sigma}(0) = \Sigma_{yy}(0)$ the assertion is true for $t = 0$. Assume now that $\bar{\Sigma}(t) \geq \Sigma_{yy}(t)$ is true for a generic t. We have that

$$\begin{aligned}
&\bar{\Sigma}(t+1) - \Sigma_{yy}(t+1) \\
&\qquad = \mathbb{E}\left[P(t)\bar{\Sigma}(t)P(t)\right] + \frac{1}{4}\mathbb{E}\left[(P(t)-I)^2\right] \\
&\qquad\quad - \left(\mathbb{E}\left[P(t)\Sigma_{yy}(t)P(t)\right] + \mathbb{E}\left[(P(t)-I)\,\Sigma_{ee}(t)\,(P(t)-I)\right]\right) \\
&\qquad = \mathbb{E}\left[P(t)(\bar{\Sigma}(t)-\Sigma_{yy}(t))P(t)\right] + \mathbb{E}\left[(P(t)-I)\left(\frac{1}{4}I - \Sigma_{ee}(t)\right)(P(t)-I)\right].
\end{aligned}$$

Since by inductive hypothesis $\bar{\Sigma}(t) \geq \Sigma_{yy}(t)$ and since by Proposition 3.5 we know that $\Sigma_{ee}(t) \leq \frac{1}{4}I$ for all $t \geq 0$, we have that $\bar{\Sigma}(t+1) - \Sigma_{yy}(t+1) \geq 0$. □

Observe now that, since $P(t)^2 = P(t)$ we obtain that $\mathbb{E}[(I - P(t))^2] = I - \mathbb{E}[P(t)]$. From this fact we obtain the following result.

Proposition 3.7 *Given the above definitions and (3.28),*

$$\lim_{t\to\infty} \bar{\Sigma}(t) = \frac{1}{4}\left(I - \frac{1}{N}\mathbf{1}\mathbf{1}^*\right).$$

Proof: Define the matrix $\bar{\mathcal{B}} = \mathbb{E}[(I - P(t))^2]$ Since $\bar{\Sigma}_{yy}(0) \in S$, and $\mathcal{A}$ is asymptotically stable if restricted to the subspace $\mathcal{S}$, then

$$\lim_{t\to\infty} \bar{\Sigma}(t) = \sum_{t=0}^{+\infty} \mathcal{A}^{(t)}(\bar{\mathcal{B}}).$$

This is the only fixed point of the iteration law (3.31). Thus we are left to prove that $\Sigma^* = \frac{1}{4}\left(I - \frac{1}{N}\mathbf{1}\mathbf{1}^*\right)$ is a fixed point, that is $\Sigma^* = \mathcal{A}(\Sigma^*) + \bar{\mathcal{B}}$. This is true, because

$$\begin{aligned}
\mathcal{A}(\Sigma^*) + \bar{\mathcal{B}} &= \frac{1}{4}\mathbb{E}\left[P(t)\left(I - \frac{1}{N}\mathbf{1}\mathbf{1}^*\right)P(t)\right] + \frac{1}{4}(I - \mathbb{E}[P(t)]) \\
&= \frac{1}{4}\left\{\mathbb{E}\left[P(t)^2\right] - \frac{1}{N}\mathbf{1}\mathbf{1}^* + I - \mathbb{E}[P(t)]\right\} = \frac{1}{4}\left\{I - \frac{1}{N}\mathbf{1}\mathbf{1}^*\right\}.
\end{aligned}$$

□

Corollary 3.8 *For every probability matrix W we have that $J(W) \leq \frac{1}{2}$.*

Proof: From the above proposition we can argue that $\bar{J} = \frac{1}{2}\sqrt{\frac{N-1}{N}}$, and since $J(W) \leq \bar{J}$, we can conclude. □

From these theorems we draw a strong conclusion about the convergence of the algorithm. In spite of missing consensus in the strict sense, the asymptotical mean squared error of the algorithm is smaller than the size of the quantization bin, and has a bound which does not depend on the number of the agents, nor on the topology of the graph, nor on the probability of the edges selection.

3.3.2 Globally quantized strategy

The algorithm for the globally quantized strategy, when the edge (i,j) is chosen, can be written as

$$\begin{aligned} x_i(t+1) &= \frac{1}{2}q_p(x_i(t)) + \frac{1}{2}q_p(x_j(t)) \\ x_j(t+1) &= \frac{1}{2}q_p(x_j(t)) + \frac{1}{2}q_p(x_i(t)). \end{aligned} \tag{3.32}$$

Below we prove that the law (3.32), as the law (3.14), drives almost surely the systems to exact consensus at an integer value. Moreover, we show by simulations, that the consensus point, even if (3.32) does not preserve the average of the state, is rather close to the average of the initial condition. This represents a significant improvement with respect to the strategy (3.14), that, as seen in Figure 3.2, leads to a consensus point whose distance from the average of the initial condition, is not negligible in general.

With the globally quantized strategy (3.32), like with (3.18), we have to deal with two sorts of randomness, since the interacting pair is randomly selected, and the quantization map is itself random. This makes the analysis of (3.32) more complicate than the analysis of (3.14). However, again, we are able to prove the convergence by a *symbolic dynamics* approach.

Let again $n_i(t) = \lfloor 2x_i(t) \rfloor$ for all $i \in V$ and let $n(t) = [n_1(t), \dots, n_N(t)]^*$. Before finding a recursive equation for $n(t)$, we need to introduce the following random variable. Let

$$T_{all} = \inf\{t : \text{at time t every node in } V \text{ has been selected at least once}\}$$

T_{all} is an integer random variable which is almost surely finite, because nodes are selected with positive probability. Note that, from (3.32), $x_i(t) \in \{a, a+1/2\}$ for some integer number a, for all $t \geq T_{all}$. This allows us to disregard the evolution before T_{all} and to analyze, for $t > T_{all}$, the symbolic dynamics as follows. For $t \geq T_{all}$, by recalling how the probabilistic quantizer works, we have that

$$q_p(x_i(t)) = \begin{cases} \frac{n_i(t)}{2} & \text{if } n_i(t) \text{ is even} \\ \lceil \frac{n_i(t)}{2} \rceil \quad \text{with probability } 1/2 & \text{if } n_i(t) \text{ is odd} \\ \lfloor \frac{n_i(t)}{2} \rfloor \quad \text{with probability } 1/2 & \end{cases}$$

Let ξ_1 and ξ_2 be two independent Bernoulli random variables with parameter $1/2$ and define $g_4 : \mathbb{Z} \times \mathbb{Z} \to \mathbb{Z}$ by

$$g_4(h,k) = \left\lceil \frac{h}{2} \right\rceil + \left\lceil \frac{k}{2} \right\rceil - \xi_1 r_h - \xi_2 r_k,$$

where r_h denotes the remainder of the division of h by 2. If, at time instant t, the edge (i,j) is selected, then

$$(n_i(t+1), n_j(t+1)) = (g_4(n_i(t), n_j(t)), g_4(n_i(t), n_j(t))). \tag{3.33}$$

The following result characterizes the convergence properties of (3.33). Recall the definition of the set $\mathcal{A}$ in (3.16).

Theorem 3.9 *Let $n(t)$ evolve according to (3.33). For every fixed initial condition $n(0)$, almost surely there exists $T_{con} \in \mathbb{N}$ such that $n(t) \in \mathcal{A}$ for all $t \geq T_{con}$.*

Proof: The proof is quite similar to the proof of Theorem 3.3 and Theorem 3.1, and it is based on proving the following three facts:

(i) each element in the set $\mathcal{A}$ is invariant for the evolution described by (3.33);

(ii) $n(t)$ is a Markov process on a finite number of states;

(iii) there is a positive probability for $n(t)$ to reach a state in $\mathcal{A}$ in a finite number of steps.

Let us now check them in order.

(i) is trivial.

(ii) Markovianity immediately follows from the fact that subsequent random choices of the edges are independent and from (3.33). To prove that the states are finite, define $m(t)$ and $M(t)$ as in (3.10) and (3.11). Let $h \in \mathbb{Z}$. Then, from the structure of g_4 we have that

- $g_4(h, h) = h$ if h is even;
- $h - 1 \leq g_4(h, h) \leq h + 1$ if h is odd.

The above two properties imply that $m(0) - r_{m(0)} \leq n_i(t) \leq M(0) + r_{M(0)}$ for all $i \in V$ and for all $t \geq 0$, This yields (ii).

(iii) Observe that
$$g_4(h, k) = g_3(h, k) - \xi_1 r_h - \xi_2 r_k,$$
where g_3 is the map defining the evolution of (3.15). Hence
$$\mathbb{P}\left[g_4(h, k) = g_3(h, k)\right] \geq \frac{1}{4}.$$
This fact, combined with the fact (iii) proved along the proof of Theorem 3.3, ensures that, also for (3.33), there is a positive probability of reaching a state in $\mathcal{A}$ in a finite time.

□

The above theorem and the previous remarks about T_{all} lead to the following claim about the original system.

Corollary 3.10 *Let $x(t)$ evolve following (3.32). Then almost surely there exists $T_{con} \in \mathbb{N}$ and $\alpha \in \mathbb{Z}$ such that $x_i(t) = \alpha$ for all $i \in V$ and for all $t \geq T_{con}$.*

As for (3.14), it is an open problem to provide a theoretical estimation of the distance between the consensus point to which (3.32) leads the systems, and the average of the initial condition. We limit our analysis to the following simulations. In Figure 3.7 we plot the variable $z = |c - 1/N\mathbf{1}^* x(0)|$, where c is such that $\lim_{t\to\infty} x(t) = c\mathbf{1}$, as a function of N. The variable z represents the distance between the consensus point to which the globally quantized strategy leads the systems and the average of the initial condition. We plot the value of z for a family of random geometric graphs of increasing size from $N = 10$ up to $N = 80$. The initial condition $x_i(0)$ is chosen randomly inside the interval $[-100, 100]$ for all $1 \leq i \leq N$. Moreover for each N, the plotted z is the average of 100

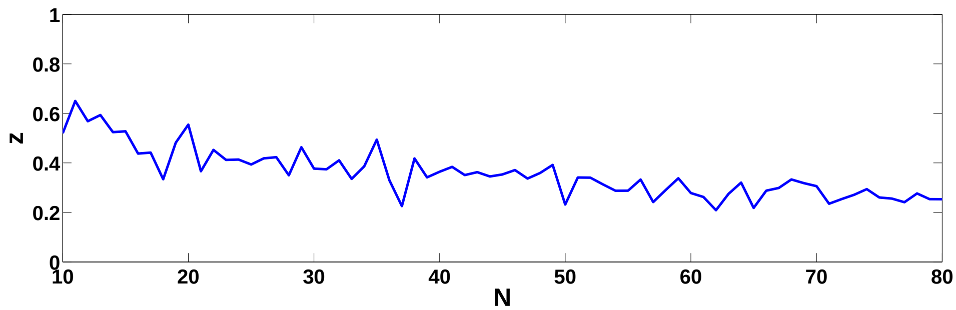

Figure 3.7: Behavior of z for a family of random geometric graphs when considering the globally quantized strategy using probabilistic quantizers.

trials. In Figure 3.8 we provide a comparison between (3.14) and (3.32). The globally quantized strategy using probabilistic quantizers, differently from the globally quantized strategy using deterministic quantizers, seems to reach the consensus very close to the average of the initial condition.

3.4 Conclusion and open questions

In this chapter we studied the gossip algorithm for the consensus problem with quantized communication. In order to face the effects due to the quantization (both deterministic and probabilistic) we proposed here two updating rules: the globally quantized strategy and the partially quantized strategy. In the former the nodes use only quantized information in order to update their state. In the latter they have access also to exact information regarding their own state. We summarize our results in Table 3.1.

We have seen that the partially quantized strategy, with both the deterministic and probabilistic quantizers, does not reach the consensus in general, but maintains the average of the state at each iteration and drives all the states very close to the average of the initial condition. On the other hand, we have shown that the globally strategy leads almost surely to a consensus which, however, does not coincide with the average of the initial condition. We have provided some simulations characterizing the distance between the consensus point and the initial average. While using the deterministic quantizer this

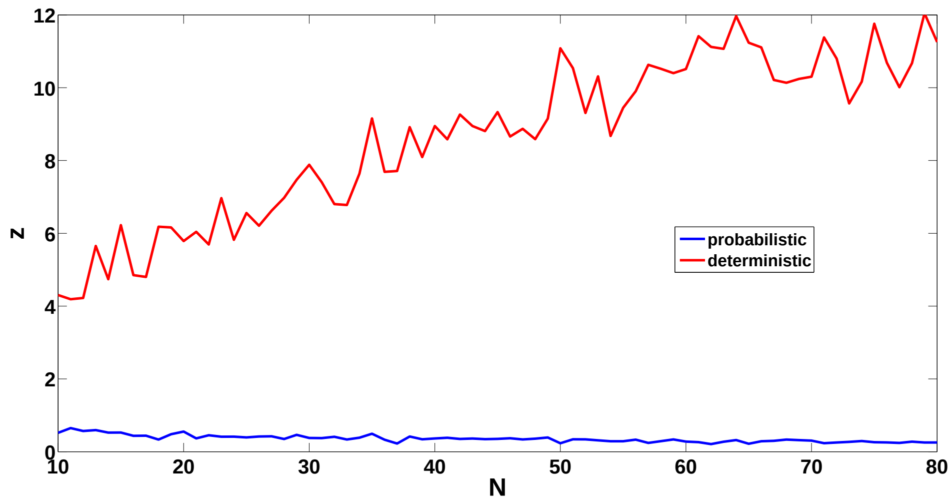

Figure 3.8: Comparison in terms of z between the "deterministic" and the "probabilistic" strategy, for a family of random geometric graphs.

Table 3.1: Summary of results

	Globally Quant.	Partially Quant.
q_d	Finite time conv. to consensus Larger error	Finite time conv. to $N^{-1/2}\|x - x_{\text{ave}}\mathbf{1}\|_2 \leq 1/2$ Average preserved
q_p	Finite time conv. to consensus Smaller error	Asympt. conv. to $N^{-1/2}\sqrt{\mathbb{E}[\|x - x_{\text{ave}}\mathbf{1}\|_2^2]} \leq 1/2$ Average preserved

distance turns out to be not negligible, using the probabilistic quantizer the consensus is reached surprisingly very close to the average of the initial condition. Providing some theoretical insights on this fact will be the object of future research.

Few words are in order to comment on the generality of the results obtained for the probabilistic quantizer q_p. Theorem 3.10 indeed does not use any statistical property of the quantizer, and not even its definition, except the fact that for all $a \in \mathbb{Z}$,

$$q_p(a + 1/2) = \begin{cases} a + 1 & \text{with probability } p \\ a & \text{with probability } 1 - p \end{cases},$$

with $p = 1/2$. Then the result holds, with the same proof, for every quantizer q, which maps $x \in \mathbb{R}$ into $q(x) \in \{\lfloor x \rfloor, \lceil x \rceil\}$, and satisfying the above condition, with $p \in (0, 1)$. On the other hand, the results in Section 3.3.1 are based on the statistics of the quantization error and does not use any other feature of the quantizer: they thus hold for any quantizer having the same statistics. The unbiasedness of the quantizer is essential to such analysis. It is easy to see that the probabilistic quantizer q_p is the *only* unbiased quantizer in the above family, and then it is the only which allows both the analyses. This justifies *a posteriori* our choice of the probabilistic quantizer.

For simplicity reasons we have chosen, defining the update rules (3.2) and (3.4), to consider consensus weights equal to 1/2. Hence, although we believe that the case we studied already shows the significant features of the problem, it is natural to ask whether our results can be extended to general choices of weights. Such extension is likely to be possible, after solving some technical difficulties. Let us consider, for instance, equation (3.4). Such an update rule can be extended considering an update matrix $P(t) = I - \varepsilon E_{ij}$, instead of (1.9). We conjectured in [40] that the so defined algorithm converges for ε a rational number in $(0, 1/2]$. This fact has been proved for $\varepsilon \in (0, 1/2]$ in the pair of forthcoming papers [56, 57].

Finally, another issue which deserves further attention is the speed of convergence of the presented algorithms. Indeed, the non-quantized gossip algorithm [9] is known to asymptotically converge (in a mean squared sense) at exponential speed, with a rate which depends on the matrix W. It is thus natural to conjecture that the convergence of the quantized version will be roughly exponential, as long as the differences states are much larger than the quantization step. Preliminary results in this sense are in [40] and in Section 3.3. However, the granularity effects eventually comes out in the convergence, making the systems converge in finite time to some limit point. Giving a rigorous clarification of this question is an interesting open problem.

Chapter 4

Efficient quantization

This chapter presents a consensus strategy in which the systems can exchange information through a strongly connected digital communication network. We present and discuss an efficient encoding/decoding strategy, based on the so called *zooming-in/zooming-out* quantization method [10, 84, 68], with theoretical and simulation results on its performance. The algorithm and the main convergence result were presented, in slightly different way, in [19]. See also the report [22].

4.1 Encoding-decoding schemes

Let us assume to have a strongly connected graph $\mathcal{G}$ and an average consensus controller K. In the extended communication framework (1.16), the control input $u_i(t)$ of agent i has the following form

$$u_i(t) = \sum_{j=1}^{N} K_{ij} \hat{x}_j^i(t)\,, \tag{4.1}$$

where $\hat{x}_j^i$ is the estimate of the state x_j which has been built by the agent i.

The point now is to describe how the estimates are built. In Chapters 2 and 3, we have assumed that the estimate is just the transmitted message. Now, instead, we look for a less trivial way of constructing estimates.

Suppose that the j-th agent sends to the i-th agent, through a digital channel, at each time instant t, a symbol $s_{ij}(t)$ belonging to a finite or countable alphabet $\mathcal{S}_{ij}$, called the *transmission alphabet*. It is assumed that the channel is reliable, that is each symbol transmitted is received without error. In general, the structure of the coder by which the j-th agent produces the symbol to be sent to the i-th agent can be described by the following equations

$$\begin{cases} \xi_{ij}(t+1) &=& F_{ij}(\xi_{ij}(t), s_{ij}(t)) \\ s_{ij}(t) &=& Q_{ij}(\xi_{ij}(t), x_j(t)) \end{cases} \tag{4.2}$$

where $s_{ij}(t) \in \mathcal{S}_{ij}$, $\xi_{ij}(t) \in \Xi_{ij}$, $Q_{ij} : \Xi_{ij} \times \mathbb{R} \to \mathcal{S}_{ij}$, and $F_{ij} : \Xi \times \mathcal{S}_{ij} \to \Xi_{ij}$. The decoder, run by agent i, is the system

$$\begin{cases} \xi_{ij}(t+1) &= F_{ij}(\xi_{ij}(t), s_{ij}(t)) \\ \hat{x}^i_j(t) &= H_{ij}(\xi_{ij}(t), s_{ij}(t)), \end{cases} \tag{4.3}$$

where $H_{ij} : \Xi_{ij} \times \mathcal{S}_{ij} \to \mathbb{R}$.
The set Ξ_{ij} serves as *state space* for the coder/decoder, whereas the maps F_{ij}, Q_{ij}, H_{ij} represent, respectively, the *coder/decoder dynamics*, the *quantizer function*, and the *decoder function*. Coder and decoder are jointly initialized at $\xi_{ij}(0) = \xi_0$.
In general, one may have different encoders at agent j, according to the various agents the agent j wants to send its data. For the sake of simplicity, we assume in the sequel that each system uses the same encoder for all data transmissions. Thus, each agent j broadcasts the same symbol $s_j(t)$ to all its neighbors. In this case every agent receiving data from j obtains the same estimate of $x_j(t)$, which we shall denote by $\hat{x}_j(t)$. In this way, by letting $F_j = F_{ij}, H_j = H_{ij}, Q_j = Q_{ij}$ and $\Xi_j = \Xi_{ij}$, the previous coder/decoder couple can be represented by the following *state estimator with memory*

$$\begin{cases} \xi_j(t+1) &= F_j(\xi_j(t), s_j(t)) \\ s_j(t) &= Q_j(\xi_j(t), x_j(t)) \\ \hat{x}_j(t) &= H_j(\xi_j(t), s_j(t)) \end{cases} \tag{4.4}$$

Moreover (4.1) assumes the following form

$$u_i(t) = \sum_{j=1}^{N} K_{ij} \hat{x}_j(t), \tag{4.5}$$

and the system evolution is

$$x(t+1) = x(t) + K\hat{x}(t). \tag{4.6}$$

It is worth observing that the controller (4.5) preserves the average of the state $x(t)$ at each time step, since Proposition 1.8 applies.

Remark 4.1 (Need for synchrony) Remark that, from the point of view of the implementation, both agent j and its neighbors have to run identical copies of the system $\xi_j(t)$. This clearly implies an intrinsical weakness of the scheme with respect to failures: if messages are lost, the synchrony of these copies is lost as well, and the algorithm does not work.

In change of this potential drawback, the scheme has very interesting convergence properties.

4.2 Zooming-in zooming-out strategy

Our strategy is inspired by the quantized stabilization technique proposed in [10], which is called zooming-in/zooming-out strategy. In this case the information exchanged between the agents is quantized by scalar uniform quantizers which assume values in a finite set, and can be described as follows. For $m \in \mathbb{N}$ define the set of quantization levels

$$\mathcal{S}_m = \{-1 + \frac{2\ell - 1}{m} \mid \ell \in \{1, \dots, m\}\} \cup \{-1, 1\}.$$

The corresponding *uniform quantizer* $q^{(m)} : \mathbb{R} \to \mathcal{S}_m$ is as follows. Let $x \in \mathbb{R}$ then

$$q^{(m)}(x) = -1 + \frac{2\ell - 1}{m}$$

if $\ell \in \{1, \dots, m\}$ satisfies $-1 + \frac{2(\ell-1)}{m} \leq x \leq -1 + \frac{2\ell}{m}$, otherwise $q^{(m)}(x) = 1$ if $x > 1$ and $q^{(m)}(x) = -1$ if $x < -1$.

It is easy to see that the quantizer $q^{(m)}$ enjoys the following property.

Lemma 4.1 (Quantization error) *Given $z \in \mathbb{R}$ and $l \in \mathbb{R}_{>0}$ such that $|z| \leq l$, it holds*

$$|z - lq^{(m)}\left(\frac{z}{l}\right)| \leq \frac{l}{m}. \tag{4.7}$$

Let $m \in \mathbb{N}$, $k_{\text{in}} \in]0, 1[$, and $k_{\text{out}} \in]1, +\infty[$. The *zooming-in/zooming-out* encoder/decoder, with parameters $m, k_{\text{in}}, k_{\text{out}}$, is defined by the alphabet $\mathcal{S} = \mathcal{S}_m$, the state space $\Xi_j = \mathbb{R} \times \mathbb{R}_{>0}$, such that $\xi_j(t) = (\hat{x}_j(t), l_j(t))$, and the dynamics

$$\begin{aligned}
&\hat{x}_j(0) = 0 \\
&\hat{x}_j(t+1) = \hat{x}_j(t) + l_j(t+1)s_j(t+1) \qquad \forall\, t \geq 0,
\end{aligned}$$

and

$$\begin{aligned}
&l_j(0) = l_0 \in \mathbb{R} \\
&l_j(1) = l_0 \\
&l_j(t+1) = \begin{cases} k_{\text{in}} l_j(t) & \text{if } |s_j(t)| < 1 \\ k_{\text{out}} l_j(t) & \text{if } |s_j(t)| = 1 \end{cases} \qquad \forall\, t > 0.
\end{aligned}$$

The sent symbol is

$$s_j(t) \quad = \quad q^{(m)}\left(\frac{x_j(t) - \hat{x}_j(t-1)}{l_j(t)}\right) \qquad \forall\, t > 0.$$

Let us comment on such definition. Remark that the first component of the coder/decoder state contains $\hat{x}(t)$, the estimate of $x(t)$. The transmitted messages contain a quantized version of the estimation error $x_j(t) - \hat{x}_j(t-1)$ scaled by the factor $l_j(t)$. Accordingly, the second component of the coder/decoder state, l_j, is referred to as the *scaling factor*: it increases when $|x_j(t) - \hat{x}_j(t-1)| > l_j(t)$ ("zooming out step") and decreases when $|x_j(t) - \hat{x}_j(t-1)| \leq l_j(t)$ ("zooming in step").

We can now prove the main theorem of this chapter.

Theorem 4.2 (Convergence) *Consider the system (4.6) with the zooming coding. Let ρ be the essential spectral radius of $I+K$. Suppose that $\rho < k_{\text{in}} < 1$, $m \geq \frac{(4+3k_{\text{in}})\sqrt{N}}{k_{\text{in}}(k_{\text{in}}-\rho)}$ and that $l_0 > \frac{2(\rho+2)\|x(0)\|}{k_{\text{in}}-\frac{3\sqrt{N}}{m}}$. Then, for any initial condition $x(0) \in \mathbb{R}^N$,*

$$\lim_{t\to+\infty} x(t) = \lim_{t\to+\infty} \hat{x}(t) = x_{\text{ave}}\mathbf{1},$$

and the convergence is exponential, with rate not bigger than k_{in}.

Proof: Let us rewrite the overall system as the coupling of the three dynamical systems on $\mathbb{R}^N$,

$$\begin{aligned}
&l_j(0) = l_0 \qquad \forall j \in V\\
&l_j(1) = l_0 \qquad \forall j \in V\\
&l_j(t+1) = \begin{cases} k_{\text{in}} l_j(t) & \text{if } |x_j(t+1)-\hat{x}_j(t)| < 1\\ k_{\text{out}} l_j(t) & \text{if } |x_j(t+1)-\hat{x}_j(t)| \geq 1 \end{cases} \qquad \forall j \in V, t \in \mathbb{N}
\end{aligned}$$

$$\begin{aligned}
&x(0) = x_0\\
&x(t+1) = x(t) + K\hat{x}(t) \qquad \forall t \in \mathbb{Z}_{\geq 0}
\end{aligned}$$

$$\begin{aligned}
&\hat{x}_j(0) = 0 \qquad \forall j \in V\\
&\hat{x}_j(t+1) = \hat{x}_j(t) + l_j(t+1)q^{(m)}\left(\frac{x_j(t+1)-\hat{x}_j(t)}{l_j(t+1)}\right) \qquad \forall j \in V, t \in \mathbb{Z}_{\geq 0}.
\end{aligned}$$

Their coupling implies that the following facts are equivalent:

(i) $\lim_{t\to+\infty} l(t) = 0$;

(ii) $\lim_{t\to+\infty} x(t) = x_{\text{ave}}\mathbf{1}$;

(iii) $\lim_{t\to+\infty} (\hat{x}(t) - x(t)) = 0$.

The idea of the proof is the following: we show that under the assumptions there always happen zooming in steps. This implies $l(t) = k_{\text{in}}^{t-1} l_0 \mathbf{1}$ and implies in turn convergence to average consensus.

By definition, there are only zooming in steps if

$$|x_j(t+1) - \hat{x}_j(t)| < l_j(t+1) \quad \forall i \in V, t \in \mathbb{Z}_{\geq 0} \tag{4.8}$$

or equivalently if

$$|x_j(t+1) - \hat{x}_j(t)| < k_{\text{in}}^t l_0 \quad \forall i \in V, t \in \mathbb{Z}_{\geq 0}. \tag{4.9}$$

Let us define $\tilde{y}(t) = Kx(t)$ and $e(t) = x(t) - \hat{x}(t)$, and remark that (4.9) is equivalent to

$$|\tilde{y}(t) - (K - I)e(t)| < k_{\text{in}}^t l_0, \tag{4.10}$$

where the inequality is meant componentwise.

To prove this fact, we proceed by strong induction over t, proving that for all $t \in \mathbb{Z}_{\geq 0}$

$$\|\tilde{y}(t)\| \leq k_{\text{in}}^t l_0 \tag{4.11}$$

$$\|\tilde{y}(t) - (K - I)e(t)\| \leq k_{\text{in}}^t l_0. \tag{4.12}$$

Since $\|\tilde{y}(t) - (K - I)e(t)\|_\infty \leq \|\tilde{y}(t) - (K - I)e(t)\|$, the inequality (4.12) implies (4.9) and then convergence.

Let us thus check the validity of Equations (4.11) and (4.12) for $t = 0$, $t = 1$, and for a general $t + 1$, with $t \in \mathbb{N}$, given it true for all $t' < t$. We shall use the recursion

$$\tilde{y}(t + 1) = (I + K)\tilde{y}(t) - K^2 e(t),$$

as well as Lemma 4.1, which implies $\|e(t)\| \leq \frac{\sqrt{N}}{m} l_0$.

Using triangle and submultiplicative inequality of norms, and $\|I + K\| = \rho$, $\|K\| \leq 2$, and $\|K - I\| \leq 3$, we get that

$t = 0$. (4.11) holds if $\|Kx_0\| < l_0$, and (4.12) holds if $\|x_0\|_\infty < l_0$.

$t = 1$. (4.12) holds if

$$2(2 + \rho)\|x_0\| + 3\frac{\sqrt{N}}{m} l_0 < k_{\text{in}} l_0, \tag{4.13}$$

and (4.11) holds if

$$\|Kx_0\| < k_{\text{in}} l_0. \tag{4.14}$$

$t > 1$. Assuming (4.11) and (4.12) true for time steps t and $t - 1$, we obtain that (4.12) holds for the time step $t + 1$ if

$$\rho k_{\text{in}}^t + 4\frac{\sqrt{N}}{m} k_{\text{in}}^{t-1} l_0 + 3\frac{\sqrt{N}}{m} k_{\text{in}}^t l_0 < k_{\text{in}}^{t+1} l_0,$$

that is if

$$\rho k_{\text{in}} + 4\frac{\sqrt{N}}{m} + 3\frac{\sqrt{N}}{m} k_{\text{in}} < k_{\text{in}}^2. \tag{4.15}$$

Instead, (4.11) holds if

$$\rho k_{\text{in}} + 4\frac{\sqrt{N}}{m} < k_{\text{in}}^2. \tag{4.16}$$

The system of conditions (4.13)-(4.14)-(4.15)-(4.16) is satisfied when the assumptions of the theorem hold. □

This result is quite conservative in proving convergence, since it guarantees that since the beginning only zooming in steps happen. This is clearly restrictive, since there is no need for $l(t)$ to be monotonic to converge to zero.

4.3 Simulation results

In this section we show by simulations the properties of the algorithm in terms of (speed of) convergence, as it depends on the parameters of the method, $k_{\text{in}}, k_{\text{out}}, m$, and on the topology of the graph. Indeed, the convergence result we obtained, Theorem 4.2, gives sufficient conditions on the parameters, depending on ρ, and on N, and thus on the graph. Convergence is possible, provided the conditions, at exponential speed, with rate $k_{\text{in}} > \rho$. Simulations are worth of interest, since they demonstrate that convergence is possible also *outside the scope of Theorem 4.2*, and with a *speed which can be faster than the linear algorithm with ideal communication* (1.5).

Roughly speaking, we can identify two different regimes for the evolution of the algorithm, depending on the zooming-in rate k_{in} which is enforced. If $k_{\text{in}} > \rho$, the zooming-in follows the natural contraction rate of the system, and according to the spirit of the proof of Theorem 4.2, eventually (almost) only zooming-in steps happen, and the rate of convergence is (no better than) k_{in}. The choice of k_{out} plays essentially no role. If instead $k_{\text{in}} < \rho$, the state does not contract fast enough for a stationary sequence of zooming-ins to establish. Then, zooming-in and zooming-out steps alternate in a complicated way, leading anyhow to consensus, with an evolution which depends on both k_{in} and k_{out}: the convergence can be faster than in the ideal case. Such phenomenon is rather surprising at a first glance, but we can explain it intuitively if we consider that the zooming-in/zooming-out algorithm takes advantage of the use of memory, and thus the sent messages are actually more informative.

Most of the simulations regard ring graphs, because concentrating on a simple example of a poorly connected graph helps to highlight the features of the algorithm. Weights of matrices P have been chosen according to a maximum degree rule, as explained at page 17. Hence a ring of $N = 20$ nodes induces a matrix whose essential spectral radius is $\rho = 0.9673$. The initial conditions have been generated according to a gaussian distribution, and all simulations with 20 nodes start from the same initial condition.

The transition between the two regimes is explored in Figure 4.1, confirming the role of ρ. Instead, Figures 4.2 and 4.2 consider the role of the number of quantization levels m. In the $k_{\text{in}} > \rho$ regime, at least $m = 3$ is required for convergence, and increasing m leads to improve the rate, approaching k_{in}. In the $k_{\text{in}} < \rho$ regime, $m = 1$ is enough for convergence, and increasing m leads to slightly improve the rate. Since $|\mathcal{S}_m| = m+2$, then 5 and 3 symbols are sufficient for convergence, respectively. If, taking advantage of the synchrony, one encodes the middle level, which is zero whenever m is odd, as a no-signal, the required number of level can be obtained with only two bits (one bit, respectively).

The scaling properties of the algorithm are a main point of interest. Figure 4.4 considers how the performance of the algorithm (with fixed parameters) degrades as N increases, in a sequence of ring graphs. The plot shows that m need not to depend on N as required Theorem 4.2. Instead, Figure 4.5 considers two cases in which the zooming factors k_{in} and k_{out} are depends on the graph, in such a way to keep the evolution in the same regime, in spite of the growth of N.

Not to restrict ourselves to ring graphs, we show some simulations regarding random geometric graphs in Figure 4.6. The plots show how the performance depends on the

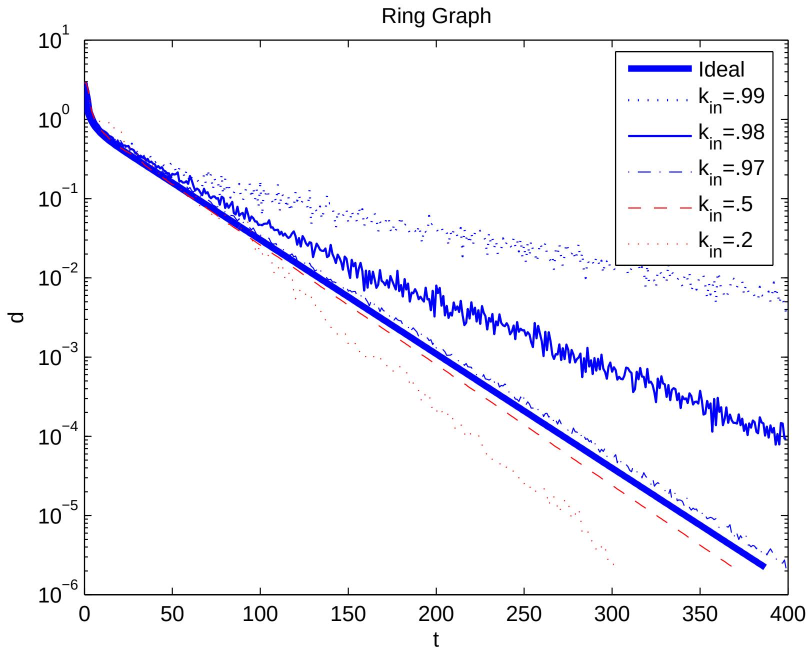

Figure 4.1: Ring graph. Dependence on k_{in}: remark the threshold $\rho = 0.9673$. $N = 20$, $m = 6$, $k_{\text{out}} = 2$.

topology of the geometric graph: different samples have different spectral radiuses, and then different convergence rates in the ideal case. Instead, in most cases the zooming algorithm is in the $k_{\text{in}} < \rho$ regime, and then the performance is roughly independent of the sample. This suggests that in this regime the algorithm is little sensitive to the spectral radius, and rises the question of which is a significant graph theoretic index.

4.4 Conclusions

With theoretical and experimental results we showed that using the zooming scheme the average consensus problem can be efficiently solved although the agents can send only quantized information. Indeed the systems converge to average consensus at exponential speed, with a convergence rate which can be chosen as small as the convergence rate of consensus with perfect exchange of information, provided the number of quantization levels is large enough. Though the theoretical results are quite conservative the efficiency of these methods is apparent from simulations, which show the method to converge outside the scope of current results, and converge at a rate faster than the ideal algorithm. Improving the theoretical analysis is thus a major concern, especially about the issue of scalability in the number of agents. We shall note that recently [6] a modification of the presented zooming algorithm has been proved to converge under a different condition.

We have already underlined in Remark 4.1 that the algorithm requires the states

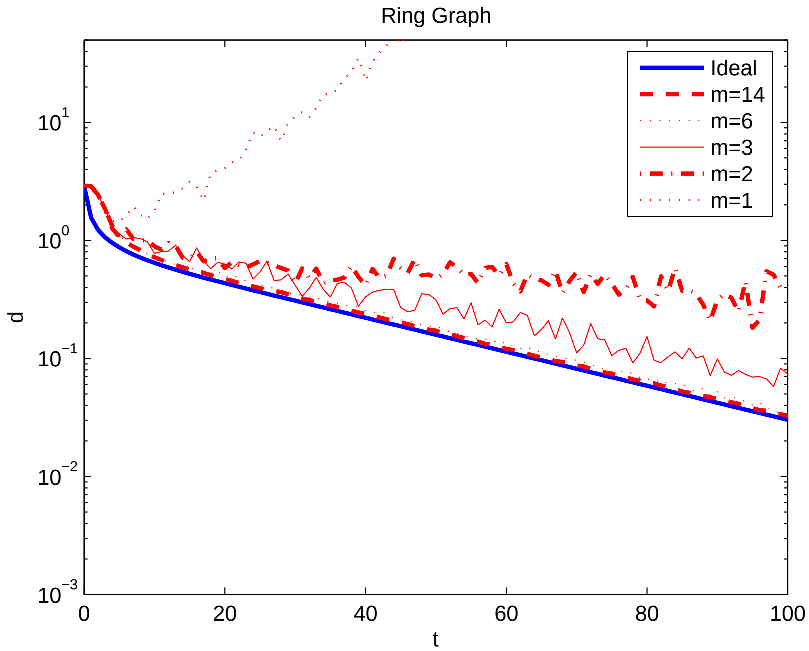

Figure 4.2: Ring graph, in the $k_{\text{in}} > \rho$ regime. $N = 20$, $k_{\text{in}} = 0.97$, $k_{\text{out}} = 2$.

estimates to be shared among neighbors. This implies that the algorithm is not suitable, in the present form, for implementation on a digital noisy channel, unless we allow a feedback on the channel state. A natural development of this research is thus designing algorithms to solve the average consensus problem at exponential speed over a network of digital noisy channels.

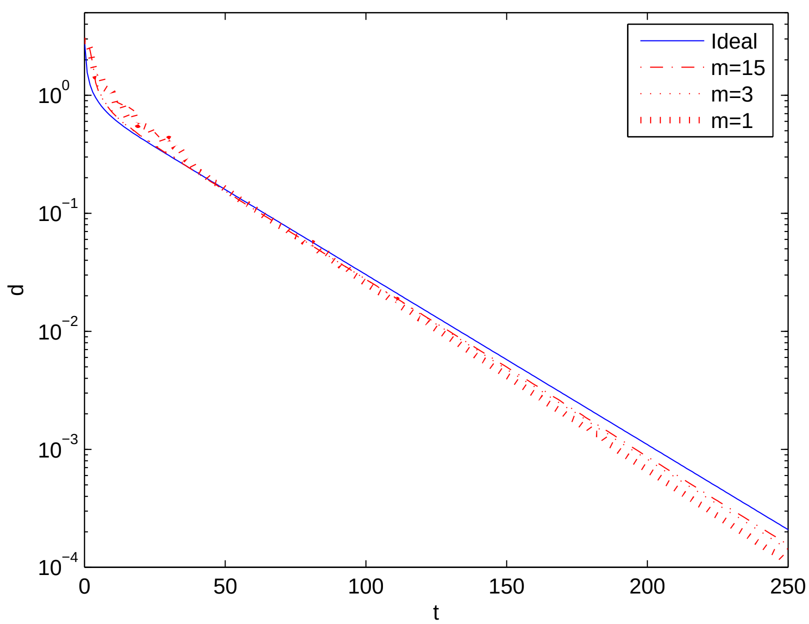

Figure 4.3: Ring graph, in the $k_{\rm in} < \rho$ regime. $N = 20$, $k_{\rm in} = 0.9$, $k_{\rm out} = 2$.

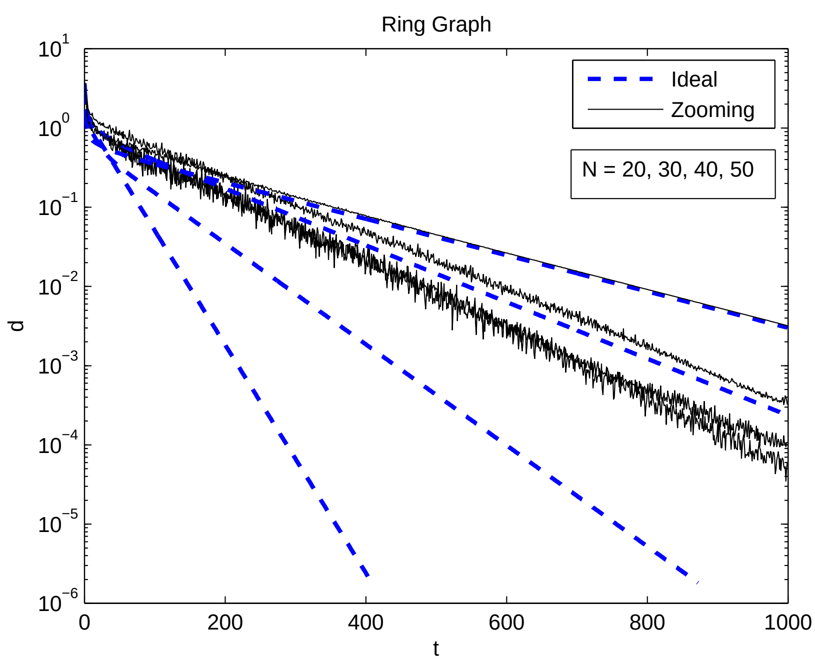

Figure 4.4: Ring graph, with $m = 3$, $k_{\rm in} = 0.99$, $k_{\rm out} = 2$. For $N > 20$ we are out of the scope of Theorem 4.2, and for $N > 30$, we have $k_{\rm in} < \rho$.

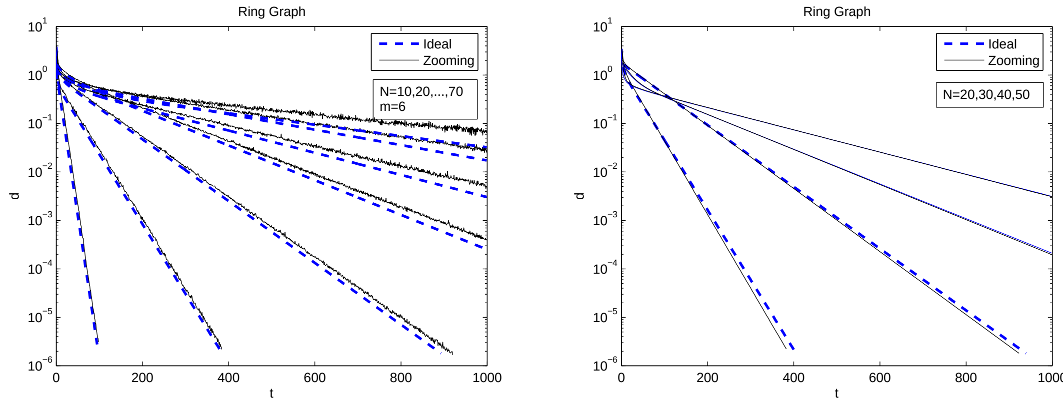

Figure 4.5: Ring graph, $m = 6$, for increasing N. Left plot assumes $k_{\text{in}} = 1.01\rho$ and $k_{\text{out}} = 2$, yielding the $k_{\text{in}} > \rho$ regime. Right plot assumes $k_{\text{in}} = 0.5\rho$, $k_{\text{out}} = k_{\text{in}}^{-1}$, yielding the $k_{\text{in}} < \rho$ regime.

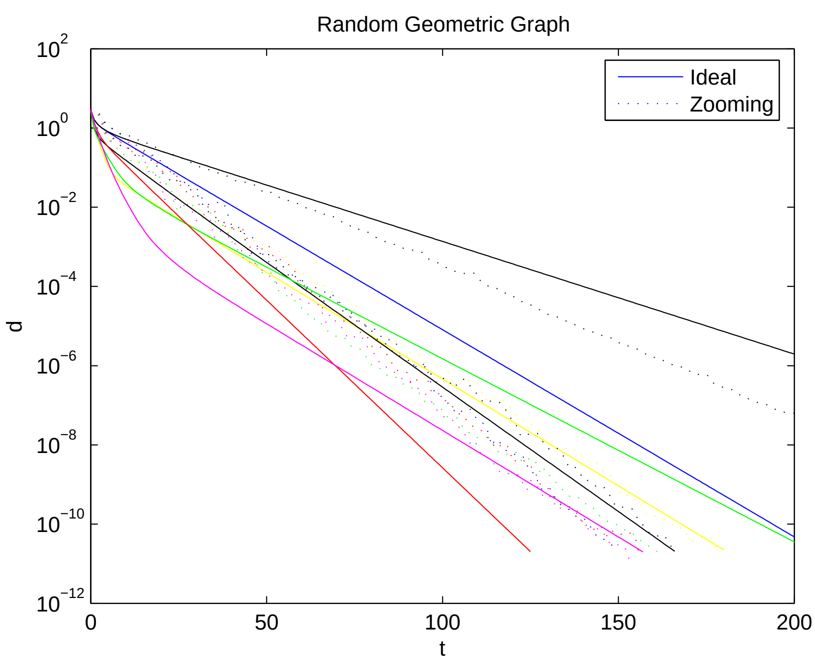

Figure 4.6: Different performance on different samples of random geometric graph with $R = 0.5$. $m = 3$, $k_{\text{in}} = 0.3$, $k_{\text{out}} = 2$.

Chapter 5

Perspectives on the consensus problem

During this first part we have studied the average consensus problems over networks of digital links. We have shown that quantized information can be effectively used in consensus algorithms. If we use uniform static quantizers, convergence to average consensus can not be achieved with arbitrary precision. This can be due either to a corruption of the average, or to a lack of convergence. These phenomena have been studied, in Chapters 2 and 6, for algorithms involving synchronous and gossip communication. Abstracting from the many results obtained, a general observation is that the use of a randomized quantizer can improve the performance. If instead we design efficient quantization techniques, as in Chapter 4, we can obtain similar performance as in the ideal case, at the price of using additional memory and computational resources.

This research is open to many potential developments: we list here some of the most interesting and promising.

(i) *Functionals of P.* It is well known that the speed of convergence of the linear ideal consensus algorithm depends on the spectral radius of P. It comes out from the results of this thesis (2.19), that the asymptotic performance of the algorithms with quantized communication depends in a more elaborated way on the eigenvalues of P. Such functionals of P are of great interest both theoretically and in practice, to compare the relative advantages of different topologies. A recent work considering various functionals of P is [46]. Up to now they have been studied on Cayley topologies: extending this study to more general graphs, e.g., random geometric graphs, would be very interesting. Moreover, it would be interesting to study the optimization of $\Phi(P)$, for a given topology, as a function of the edge weights.

(ii) *Speed of convergence.* Most of our attention has been given to the degradation of the asymptotical performance. However, since in practice the algorithms are supposed to run for a finite time, it is important to have results about the speed of convergence of the algorithms with quantization, or about their performance after a finite horizon of time. A general intuition, expressed in the literature and confirmed from the results presented here, is that the speed of convergence of the quantized algorithms is roughly the same as of the ideal case, in the sense that the overall error decreases exponentially in time, until the achievable precision is reached: later, there is no further improvement. However, giving clear analytical results on this point is an open question, in spite of various attempts in the literature [55, 3, 40, 23, 93].

(iii) *High performance quantized algorithms.* The zooming in-zooming out scheme of Chapter 4 is able to solve the average consensus problem at exponential speed, with a convergence rate which can be chosen as small as the convergence rate of consensus with ideal exchange of information, or even smaller (in simulations). We would like to extend these results to avoid, if possible, the technical limitation of the dependence of the required number of levels on the number of agents, and to understand the potential improvements, with respect to the ideal case, which can come from encoding/decoding schemes with memory.

(iv) *Digital noisy channels.* All the digital channels we have considered are lossless, reliable and without delay. Relaxing these assumptions, and providing algorithms to solve the average consensus problem at exponential speed over a network of digital noisy channels [17], is probably the greatest open issue at the moment.

(v) *Applications.* Consensus algorithms are not only interesting in themselves, but also as a step towards more ambitious goals of coordinated control and optimization: estimation [87, 90, 13, 42], Kalman filtering [15], coverage (Chapters 6 and 7), optimization [71].

Part II

Partitions optimization

Chapter 6

Gossip coverage optimization: dynamical systems on the space of partitions

Introduction

Motivation and scope

This chapter deals with distributed coverage control and partitioning policies for networks of robots. Coverage control has the goal of optimizing the configuration of a group of robots, or agents, in a given environment: the optimization problem is expressed in terms of a cost function, depending on the agents' positions.

Statement of contributions

First, we describe coverage control algorithms in a novel way. Classically, the state space for the coverage algorithms are the agent positions: based on their positions, the agents apportion the environment into regions, which are assigned to each agent. In our approach, the agents positions are no longer the main concern: the state space is a space of partitions of the given environment. We discuss important properties of such a space, namely its compactness with respect to a suitable metric, and the continuity of several functions defined on it.

Second, as key motivating application, we devise a novel algorithm for coverage optimization, a "gossip" algorithm, in which only one pair of agents communicates per time step. We do this, because we know that reducing the communication burden is a critical issue for coverage control: indeed pairwise communication can be more effective in practical situations if connections between agents are not guaranteed to be fully reliable.

Third, we provide convergence theorems which extend the LaSalle invariance principle to a special class of set-valued maps on metric spaces. Applying these extensions of the LaSalle invariance principle and the properties of the state of partitions, we are able to give conditions for the proposed algorithm to converge to the critical points of a natural cost functional.

Related works

Coverage control problems have been recently solved in a distributed way in [29], which set up the classical agent-based approach. A broad discussion about coverage control is presented in [11, 63]; other related works include [1] on sensor-based algorithms, [52] on dynamic coverage, [47] on estimation of stochastic spatial fields, [74] on equitable partitioning policies, and [92, 12, 78] on nonconvex environments. The pairwise "gossip" approach to agents communication has been already considered for the consensus problem in recent works [9, 65].

Chapter structure

The chapter is structured as follows. In Section 6.1 we formally describe the coverage control problem, we introduce the space of partitions and we review the classic results, in our perspective. In Section 6.2 we present our algorithm. The convergence of the algorithm is stated in Section 6.2.3, and proved in Section 6.5.1. The proof is based on a LaSalle invariance principle for switching maps, proved in Section 6.3. To verify that the invariance principle can be applied to the proposed algorithm, it is necessary to study the space of partitions. A suitable metric is introduced on it, so that the space is proved to be compact, in Section 6.4, and several functions and maps defined on it are proved to be continuous, in Section 6.5.

Notation

Given a subset Q of the Euclidean space $\mathbb{R}^d$, we let interior(A) denote its interior, $\overline{A}$ denote its closure, and ∂A denote its boundary. Given two sets X and Y, a *set-valued map*, denoted by $f : X \rightrightarrows Y$, associates to an element of X a subset of Y.

6.1 Coverage optimization via multicenter functions

We are given a group of robots (also called agents) with limited communication and sensing capabilities, and an environment, and we want the agents to deploy in the area in an optimal way. The optimality is described by a suitable cost function to be minimized. The environment is apportioned into smaller regions, each assigned to an agent. The partition, and the agents configuration, are iteratively updated in a way to minimize a cost functional, which depends on the current partition and the agents' positions.

6.1.1 Partitions and multicenter optimization

In this and following sections partitions are defined as follows.

Definition 6.1 (Partition) *Let Q be a compact convex subset of $\mathbb{R}^d$ with non-empty interior. An N*-partition *of Q, denoted by $v = \{v_i\}_{i=1}^N$, is a collection of N subsets of Q with the following properties:*

(i) each set v_i, $i \in \{1, \dots, N\}$, is closed, has non-empty interior, and its boundary has measure[1] *zero;*

(ii) $\text{interior}(v_i) \cap \text{interior}(v_j)$ *is empty whenever $i \neq j$; and*

(iii) $\cup_{i \in \{1,\dots,N\}} v_i = Q$.

We let $\mathcal{V}_N$ denote the set of N-partitions of Q.

Let $p = (p_1, \dots, p_N) \in Q^N$ denote the position of N agents in the environment Q. Given any $v \in \mathcal{V}_N$ and almost any[2] $p \in Q^N$, each agent is naturally in one-to-one correspondence with an element of v; specifically we sometime refer to v_i as the *dominance region* of agent $i \in \{1, \dots, N\}$.

Given a bounded measurable positive function $\phi : Q \to \mathbb{R}_{>0}$, called a *density function*, define the *multicenter function* $\mathcal{H}_{\text{multicenter}} : \mathcal{V}_N \times Q^N \to \mathbb{R}_{>0}$ by

$$\mathcal{H}_{\text{multicenter}}(v, p) = \sum_{i=1}^{N} \int_{v_i} \|p_i - q\|^2 \phi(q) dq. \tag{6.1}$$

We aim to minimize this function with respect to both the partition v and the locations p.

Remark 6.1 (Locational optimization) The function $\mathcal{H}_{\text{multicenter}}$ has the following interpretation. Given an agent at location p_i, assume that $\|p_i - q\|^2$ is the cost incurred by agent i to "service" an event taking place at point q. Events take place inside the environment Q with likelihood ϕ. Accordingly, the multicenter function $\mathcal{H}_{\text{multicenter}}$ quantifies how well the environment Q is partitioned and how well the agents are placed inside Q. This and related optimal sensor placement problems are studied in locational and geometric optimization, spatial resource allocation, quantization theory, clustering analysis, and statistical pattern recognition; see [11, Chapter 2] and references therein. □

Among all the possible ways of partitioning a subset of $\mathbb{R}^d$, there is one which is worth of special attention. Define the *set of partly coincident locations* $S_N = \{p \in Q^N \mid p_i = p_j \text{ for some } i, j \in \{1, \dots, N\}, i \neq j\}$. Given $p \in Q^N \setminus S_N$, the *Voronoi partition of Q generated by p*, denoted by $V(p)$, is the collection of the *Voronoi regions* $\{V_i(p)\}_{i=1}^N$, defined by

$$V_i(p) = \{q \in Q \mid \|q - p_i\| \leq \|q - p_j\| \text{ for all } j \neq i\}. \tag{6.2}$$

In other words, the Voronoi partition is a map $V : Q^N \setminus S_N \to \mathcal{V}_N$. The regions $V_i(p)$, $i \in \{1, \dots, N\}$, are convex and, if Q is a polytope, they are polytopes. Now, given two distinct points q_1 and q_2 in $\mathbb{R}^d$, define the $(q_1; q_2)$*-bisector half-space* by

$$H_{\text{bisector}}(q_1; q_2) = \{q \in \mathbb{R}^d \mid \|q - q_1\| \leq \|q - q_2\|\}. \tag{6.3}$$

[1]As measure we mean the Lebesgue measure in $\mathbb{R}^d$.
[2]That is, for every $p \in Q^N$ but a set of measure zero.

In other words, $H_{\mathrm{bisector}}(q_1; q_2)$ is the closed half-space containing q_1 whose boundary is the hyperplane bisecting[3] the segment from q_1 to q_2. Note that $H_{\mathrm{bisector}}(q_1; q_2) \neq H_{\mathrm{bisector}}(q_2; q_1)$ and that Voronoi partition of Q satisfies

$$V_i(p_1, \dots, p_n) = Q \cap \big(\cap_{j \neq i} H_{\mathrm{bisector}}(p_i; p_j) \big).$$

Each region equipped with a density function possesses a point of special interest, called the centroid. The *centroid* of a measurable set $A \subset Q$ with non-empty interior is the point in A defined by

$$\mathrm{Cd}(A) = \Big(\int_A \phi(q) dq \Big)^{-1} \int_A q \phi(q) dq. \tag{6.4}$$

With a slight notational abuse, given $v \in \mathcal{V}_N$, we let

$$\mathrm{Cd}(v) = (\mathrm{Cd}(v_1), \dots, \mathrm{Cd}(v_N)) \in Q^N$$

be the vector of regions centroids. Voronoi partitions and centroids have the following optimality properties.

Proposition 6.1 (Properties of $\mathcal{H}_{\mathbf{multicenter}}$) *For any partition $v \in \mathcal{V}_N$ and any point set $p \in Q^N \setminus S_N$,*

$$\mathcal{H}_{\mathrm{multicenter}}(V(p), p) \leq \mathcal{H}_{\mathrm{multicenter}}(v, p), \tag{6.5}$$

$$\mathcal{H}_{\mathrm{multicenter}}(v, \mathrm{Cd}(v)) \leq \mathcal{H}_{\mathrm{multicenter}}(v, p). \tag{6.6}$$

Furthermore, inequality (6.5) *is strict if any entry of $V(p)$ differs from the corresponding entry of v by a set with non-empty interior, and inequality* (6.6) *is strict if* $\mathrm{Cd}(v)$ *differs from p.*

These statements are proved in [11, Propositions 2.14 and 2.15]. These results motivate the following definition: A partition $v^* \in \mathcal{V}_N$ is a *centroidal Voronoi partition* if $v^* = V(\mathrm{Cd}(v^*))$.

Additionally, it is of interest to consider two variations of the multicenter function. We define $\mathcal{H}_{\mathrm{Voronoi}} : Q^N \setminus S_N \to \mathbb{R}_{\geq 0}$ and $\mathcal{H}_{\mathrm{centroid}} : \mathcal{V}_N \to \mathbb{R}_{\geq 0}$ by

$$\mathcal{H}_{\mathrm{Voronoi}}(p) = \mathcal{H}_{\mathrm{multicenter}}(V(p), p) = \sum_{i=1}^{N} \int_{V_i(p)} \|q - p_i\|^2 \phi(q) dq, \tag{6.7}$$

$$\mathcal{H}_{\mathrm{centroid}}(v) = \mathcal{H}_{\mathrm{multicenter}}(v, \mathrm{Cd}(v)) = \sum_{i=1}^{N} \int_{v_i} \|q - \mathrm{Cd}(v_i)\|^2 \phi(q) dq. \tag{6.8}$$

The function $\mathcal{H}_{\mathrm{centroid}}$ is novel here, plays a key role in later developments and has the following property that is an immediate consequence of Proposition 6.1. Given a partition v with $\mathrm{Cd}(v) \notin S_N$,

$$\mathcal{H}_{\mathrm{centroid}}(V(\mathrm{Cd}(v))) \leq \mathcal{H}_{\mathrm{centroid}}(v), \tag{6.9}$$

and this inequality is strict if any entry of $V(\mathrm{Cd}(v))$ differs from the corresponding entry of v by a set with non-empty interior.

[3] A hyperplane bisects a segment if it is perpendicular to and passes through the midpoint of the segment.

6.1.2 Distributed coverage control and its limitations

Given a group of robots with limited inter-robot communication, the objective of coverage control algorithms is to move the robots in order to minimize $\mathcal{H}_{\mathrm{multicenter}}$. To formalize the notion of limited communication, we introduce a useful graph. The *Delaunay graph* associated to the distinct positions $p \in Q^N \setminus S_N$ is the undirected graph with node set $\{p_i\}_{i=1}^N$ and with the following edges: (p_i, p_j) is an edge if and only if $V_i(p) \cap V_j(p)$ is non-empty. In other words, two agents are neighbors if and only if their corresponding Voronoi regions intersect, see Figure 6.1.

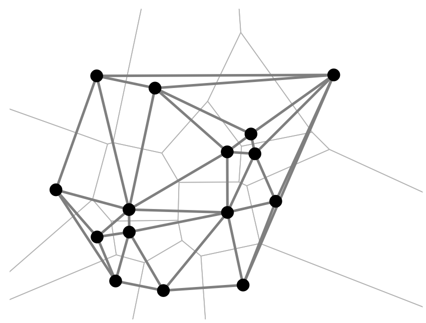

Figure 6.1: The Voronoi partition and the corresponding Delaunay graph

The *distributed coverage algorithm* studied in [29], and originally introduced by Lloyd [59], is described as follows. At each discrete time instant $t \in \mathbb{Z}_{\geq 0}$, each agent i performs the following tasks: (i) it transmits its position and receives the positions of its neighbors in the Delaunay graph; (ii) it computes its Voronoi region with the information received; (iii) it computes the centroid of its Voronoi region; and (iv) it moves to the centroid. In mathematical terms, we have

$$p(t+1) = \mathrm{Cd}(V(p(t))). \tag{6.10}$$

Because of the smoothness of the various maps, compactness of Q, and monotonicity properties in Proposition 6.1, one can show that the solutions of (6.10) converge asymptotically to the set of the centroidal Voronoi partitions. Note that this distributed coverage algorithm requires the robotic network to have the following three properties:

(i) each agent can talk with all its neighbors in the Delaunay graph;

(ii) the agents are synchronized and they all talk at the same time; and

(iii) each communication link is fully reliable and features no transmission delay.

Delaunay graphs have the following undesirable properties: for worst-case robot positions, a node might have $N-1$ neighbors and the distance between two neighbors may be arbitrarily large. In short summary, synchronized reliable communication along all edges of a Delaunay graph is a strong requirement.

In this chapter we are interested in reducing the synchronization and communication requirements of distributed coverage algorithms. In particular, is it possible to optimize agents positions and environment partition with asynchronous, unreliable, delayed communication? What if the communication model is that of *gossiping agents*, that is, a model in which only a pair of robots can communicate at any time? How do we overcome the limitation that Voronoi partitions generated by moving agents can not be computed with only asynchronous pairwise communication. Giving such a question an affirmative answer is the object of this chapter.

6.2 Partitions-based gossip coverage algorithm

We give here a general description of our approach. In the partitions-based approach, the position of the robot essentially plays no role anymore and we instead describe how to update the dominance regions. Designing coverage algorithms as dynamical systems on partitions has an important advantage: we do not restrict our attention to Voronoi partitions. However, algorithms based on Voronoi partitions can also be treated in our approach.

Example 6.1 (Lloyd algorithm in the partitions-based approach) The distributed coverage algorithm (6.10) changes the agents locations so as to incrementally minimize the function $\mathcal{H}_{\text{Voronoi}}$, while the partition of the environment is a function of the agent positions. We change this perspective and take a dual approach: we consider as a coverage algorithm an algorithm that evolves partitions. From this partitions-based viewpoint, the distributed coverage algorithm is rather an iterated map on $\mathcal{V}_N$ and the cost function to be incrementally minimized is $\mathcal{H}_{\text{centroid}}$ rather than $\mathcal{H}_{\text{Voronoi}}$. Equation (6.10) is rewritten as

$$v(t+1) = V(\mathrm{Cd}(v(t))). \tag{6.11}$$

The solutions of (6.11) converge asymptotically to the set of the centroidal Voronoi partitions. This result, which comes obviously from the analogous fact about the dynamical system (6.10), can be re-obtained using the techniques developed in the present chapter. In this sense our approach extends the one in [29]. □

6.2.1 The partitions-based gossip coverage algorithm

In this section we present the *partitions-based gossip coverage algorithm* that, at each iteration, requires only a pair of adjacent regions to communicate and change. Before presenting the algorithm, we adopt the following convention: we allow communication between adjacent regions. The following definition generalizes the notion of Delaunay graph to graphs between partitions.

Definition 6.2 (Adjacency graph generated by a partition) *The adjacency graph*[4]

[4] This graph is called the dual graph in the theory of planar graphs.

associated to the partition $v \in \mathcal{V}_N$ is the undirected graph with node set v and with the edge set $\mathcal{E}(v)$ defined as follows: (v_i, v_j) is an edge if and only if the two regions v_i and v_j are adjacent in the sense that $\overline{\mathrm{interior}(v_i)} \cap \overline{\mathrm{interior}(v_j)}$ is non-empty.

Recalling the notion of bisector half-space from equation (6.3), the *partitions-based gossip coverage algorithm* is stated as follows.

> At each time $t \in \mathbb{Z}_{\geq 0}$, each agent i maintains the dominance region $v_i(t)$; the collection $v(t) = \{v_1(t), \dots, v_N(t)\}$ is an N-partition of Q and it is initialized arbitrarily. At each $t \in \mathbb{Z}_{\geq 0}$ a communicating pair is selected, say $(i, j) \in \mathcal{E}(v(t))$. Every agent $k \notin \{i, j\}$ sets $v_k(t+1) := v_k(t)$, whereas agents i and j perform the following tasks:
>
> agent i transmits to agent j its dominance region $v_i(t)$ and vice-versa
> both agents compute the centroids $\mathrm{Cd}(v_i(t))$ and $\mathrm{Cd}(v_j(t))$
> **if** $\mathrm{Cd}(v_i(t)) = \mathrm{Cd}(v_j(t))$ **then**
> $\quad v_i(t+1) := v_i(t)$ and $v_j(t+1) := v_j(t)$
> **else**
> $\quad v_i(t+1) := \big(v_i(t) \cup v_j(t)\big) \cap H_{\mathrm{bisector}}\big(\mathrm{Cd}(v_i(t)); \mathrm{Cd}(v_j(t))\big)$
> $\quad v_j(t+1) := \big(v_i(t) \cup v_j(t)\big) \cap H_{\mathrm{bisector}}\big(\mathrm{Cd}(v_j(t)); \mathrm{Cd}(v_i(t))\big)$
> **end if**

In other words, when two agents with distinct centroids communicate, they update their regions as follows: the union of the two regions is divided by the hyperplane bisecting the segment between the two centroids; see Figure 6.2. As a consequence, $\{v_i(t+1), v_j(t+1)\}$ is a partition of the set $v_i(t) \cup v_j(t)$ and, if the centroids $\mathrm{Cd}(v_i(t))$, $\mathrm{Cd}(v_j(t))$ are distinct, it is the Voronoi partition of that set generated by the centroids.

To establish that the algorithm is well-posed, it is essential to guarantee that the sequence of collections $v(t)$ generated by the algorithm is an N-partition at all times t, that is, satisfies the three properties in Definition 6.1. The proof of this fact is postponed to Lemma 6.15.

For any pair $(i, j) \in \{1, \dots, N\}^2$, we define the map $T_{ij} : \mathcal{V}_N \to \mathcal{V}_N$ by $\bar{v} = T_{ij}(v)$, where

$$\begin{aligned} \bar{v}_i &= \big(v_i \cup v_j\big) \cap H_{\mathrm{bisector}}\big(\mathrm{Cd}(v_i); \mathrm{Cd}(v_j)\big), \\ \bar{v}_j &= \big(v_i \cup v_j\big) \cap H_{\mathrm{bisector}}\big(\mathrm{Cd}(v_j); \mathrm{Cd}(v_i)\big), \\ \bar{v}_k &= v_k, \qquad \text{for all } k \notin \{i, j\}, \end{aligned}$$

if $\mathrm{Cd}(v_i(t)) \neq \mathrm{Cd}(v_j(t))$, and $\bar{v} = v$ otherwise. Note that T_{ij} is defined for any two regions i, j not necessarily adjacent. The dynamical system on the space of partitions is therefore described by

$$v(t+1) = T_{ij}(v(t)), \qquad \text{for some } (i, j) \in \mathcal{E}(v(t)), \tag{6.12}$$

together with a rule describing what edge is selected at each time. It is also useful to consider the family of maps defined by all edges $(i, j) \in \mathcal{E}(v)$ and introduce the set-valued map $T : \mathcal{V}_N \rightrightarrows \mathcal{V}_N$ defined by $T(v) = \{T_{ij}(v) \mid (i, j) \in \mathcal{E}(v)\}$. Then the evolution under

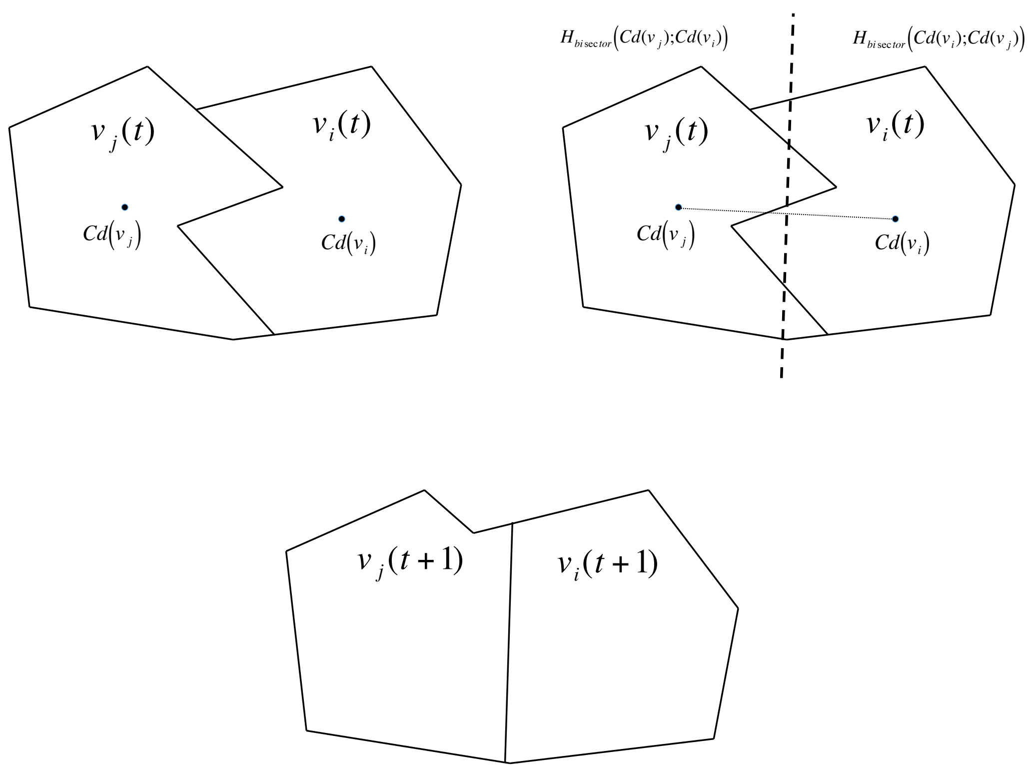

Figure 6.2: One iteration of the partitions-based gossip algorithm. Step 1: a pair of adjacent region is selected and their centroids are computed. Step 2: the bisector half-space is computed. Step 3: the resulting regions are computed.

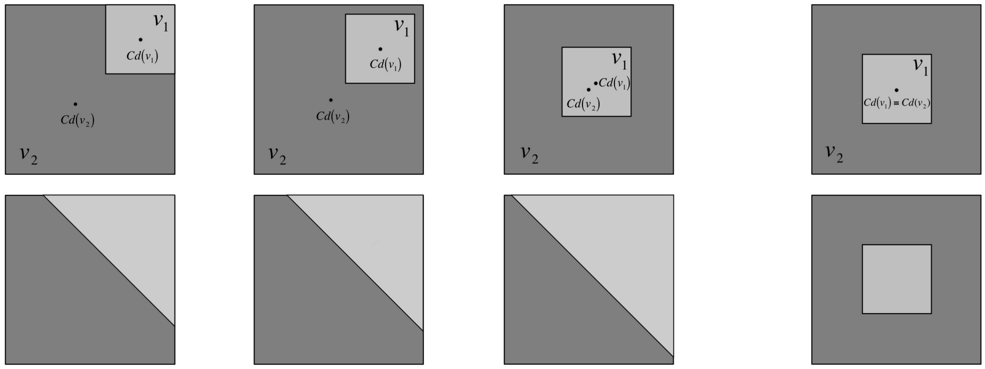

Figure 6.3: Illustration of Example 6.2 about the lack of continuity of the map T_{12}. In the first row, the first three figures show a sequence of partitions $\{v_1(t), v_2(t)\}_{t\in\mathbb{N}}$ whose limiting partition $(v_1(\infty), v_2(\infty))$, in the fourth figure, is given by two concentric sets. In the second row, the first three figures show the sequence $\{T(v_1(t), v_2(t))\}_{t\in\mathbb{N}}$; the fourth figure shows $T_{12}(v_1(\infty), v_2(\infty))$ that is equal to $(v_1(\infty), v_2(\infty))$ assuming that the environment has uniform density. Equation (6.14) is now easily verified.

the gossip coverage control algorithm is one of the solutions to the non-deterministic set-valued dynamical system

$$v(t+1) \in T(v(t)). \tag{6.13}$$

In the following sections we analyze the dynamical system described by equation (6.12) or (6.13).

6.2.2 Designing a continuous algorithm

The maps T_{ij} and T do not satisfy certain continuity properties, as illustrated by the following example.

Example 6.2 For $N = 2$, consider the map T_{12} and a sequence of partitions $\{v_1(t), v_2(t)\}_{t\in\mathbb{N}}$, with the properties that $\lim_{t\to\infty} v_1(t) = v_1(\infty)$, $\lim_{t\to\infty} v_2(t) = v_2(\infty)$, and $\mathrm{Cd}(v_1(\infty)) = \mathrm{Cd}(v_2(\infty))$. These properties imply that $T_{12}(v_1(\infty), v_2(\infty)) = \{v_1(\infty), v_2(\infty)\}$. Now, if at each time t the partition $T_{12}(v_1(t), v_2(t))$ differs from the partition $\{v_1(t), v_2(t)\}$ by regions of finite area (as for example depicted in Figure 6.3), then

$$\lim_{t\to\infty} T_{12}(v_1(t), v_2(t)) \neq T_{12}\big(\lim_{t\to\infty} v_1(t), \lim_{t\to\infty} v_2(t)\big) = \{v_1(\infty), v_2(\infty)\}. \tag{6.14}$$

□

Similar discontinuities arise also for sequences of partitions in which two regions become adjacent, or loose adjacency, in the limit, as illustrated in Figure 6.4.

We need continuity properties in our convergence analysis and, therefore, our analysis begins by introducing a minor modification of T that instead does have these properties. For any pair i, j, $i \neq j$, we introduce a new map T_{ij}^{δ}, parameterized by the positive

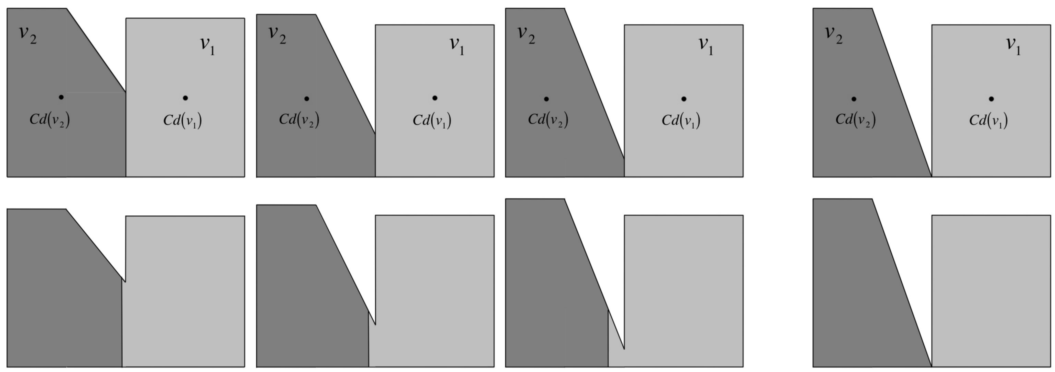

Figure 6.4: Illustration of the lack of continuity of the map T_{12}. In the first row, the first three figures show a sequence of partitions $\{v_1(t), v_2(t)\}_{t\in\mathbb{N}}$ whose limiting partition $(v_1(\infty), v_2(\infty))$, in the fourth figure, is given by two disconnected sets. In the second row, the first three figures show the sequence $\{T(v_1(t), v_2(t))\}_{t\in\mathbb{N}}$; the fourth figure shows $T_{12}(v_1(\infty), v_2(\infty))$, which is equal to $(v_1(\infty), v_2(\infty))$ assuming that the environment has uniform density. Equation (6.14) is now easily verified.

real number δ, that is a smoothed version of T_{ij}. To design the map T_{ij}^{δ} we need some geometric notions. Given $v = (v_1, \dots, v_N) \in \mathcal{V}_N$, consider two regions v_i and v_j such that $\mathrm{Cd}(v_i) \neq \mathrm{Cd}(v_j)$. Define the set R_i by

$$R_i = v_i \cap H_{\text{bisector}}(\mathrm{Cd}(v_j); \mathrm{Cd}(v_i)) = \{q \in v_i \mid \|q - \mathrm{Cd}(v_i)\| \geq \|q - \mathrm{Cd}(v_j)\|\},$$

and similarly $R_j = v_j \cap H_{\text{bisector}}(\mathrm{Cd}(v_i); \mathrm{Cd}(v_j))$. Let $\gamma_\perp$ be the hyperplane bisecting the segment from $\mathrm{Cd}(v_i)$ to $\mathrm{Cd}(v_j)$, that is,

$$\gamma_\perp = \partial H_{\text{bisector}}(\mathrm{Cd}(v_i); \mathrm{Cd}(v_j)) = \{q \in Q \mid \|q - \mathrm{Cd}(v_j)\| = \|q - \mathrm{Cd}(v_i)\|\}.$$

For each $q \in R_i \cup R_j$ there exists only one hyperplane γ parallel to $\gamma_\perp$ such that $q \in \gamma$; we denote this hyperplane as γ_q. Consider now two points, $\bar{p}_i \in R_i$ and $\bar{p}_j \in R_j$, such that

$$\bar{p}_i \in \underset{q\in \text{interior}\,(R_i)}{\operatorname{argmax}} \min_{y\in\gamma_\perp} \|q - y\|, \quad \text{and} \quad \bar{p}_j \in \underset{q\in \text{interior}\,(R_j)}{\operatorname{argmax}} \min_{y\in\gamma_\perp} \|q - y\|.$$

Given two subsets A and B of Q, we define

$$\mathrm{dist}(A, B) = \inf_{(a,b)\in A\times B} \|a - b\|.$$

Pick $\delta > 0$ and define

$$\beta(v_i, v_j) = \mathrm{sat}_\delta\left(\|\,\mathrm{Cd}(v_i) - \mathrm{Cd}(v_j)\|\right) \cdot \left(1 - \mathrm{sat}_\delta(\mathrm{dist}(\mathrm{interior}(v_i), \mathrm{interior}(v_j)))\right),$$

where $\mathrm{sat}_\delta : \mathbb{R}_{\geq 0} \to [0, 1]$ is defined by $\mathrm{sat}_\delta(x) = x/\delta$ if $x \in [0, \delta]$ and $\mathrm{sat}_\delta(x) = 1$ if $x > \delta$. Remark that $\beta(v_i, v_j) \in [0, 1]$Let

$$\tilde{R}_i = \{q \in R_i \mid \mathrm{dist}(\bar{p}_i, \gamma_q) \leq \beta(v_i, v_j)\,\mathrm{dist}(\bar{p}_i, \gamma_\perp)\},$$

and

$$\tilde{R}_j = \{q \in R_j \mid \operatorname{dist}(\bar{p}_j, \gamma_q) \leq \beta(v_i, v_j) \operatorname{dist}(\bar{p}_j, \gamma_\perp)\}.$$

Finally, we define the map $T^\delta_{ij} : \mathcal{V}_N \to \mathcal{V}_N$ by $\bar{v} = T^\delta_{ij}(v)$, where

$$\begin{aligned}
\bar{v}_i &= \Big(v_i \setminus \operatorname{interior}(\tilde{R}_i)\Big) \cup \tilde{R}_j, \\
\bar{v}_j &= \Big(v_j \setminus \operatorname{interior}(\tilde{R}_j)\Big) \cup \tilde{R}_i, \\
\bar{v}_k &= v_k, \qquad \text{for all } k \notin \{i, j\},
\end{aligned}$$

if $\operatorname{Cd}(v_i(t)) \neq \operatorname{Cd}(v_j(t))$, and $\bar{v} = v$ otherwise. Define the set-valued map $T^\delta : \mathcal{V}_N \rightrightarrows \mathcal{V}_N$ by

$$T^\delta = \{T^\delta_{ij}(v) \mid (i, j) \in \{1, \ldots, N\}^2,\ i \neq j\}. \tag{6.15}$$

As we discussed for the map T, the map T^δ is well-posed; the proof of this fact is postponed to Lemma 6.16. The map T^δ is illustrated in Figure 6.5 and in the following remark.

Remark 6.2 (Properties of the continuous algorithm) (i) Given $v \in \mathcal{V}_N$ consider the two regions v_i and v_j. If $\beta(v_i, v_j) = 1$, namely, if the distance between the regions v_i and v_j is zero (v_i and v_j are adjacent) and the distance between $\operatorname{Cd}(v_i)$ and $\operatorname{Cd}(v_j)$ is larger than δ, then $T^\delta_{ij}(v) = T_{ij}(v)$, that is, the map T^δ_{ij} reduces to the map T_{ij} introduced in the previous section.

Additionally, if $\beta(v_i, v_j) = 0$, namely, either $\operatorname{Cd}(v_i)$ and $\operatorname{Cd}(v_j)$ coincide or the distance between v_i and v_j is larger than δ, then $T^\delta_{ij}(v) = v$, that is, the map T^δ_{ij} leaves the regions unchanged.

(ii) When δ is very small, the map T^δ is very similar to the map T presented in the previous section.

(iii) The map T^δ has some continuity properties; we present this formal statement in Lemma 6.16.

(iv) It is worth noting that the "smoothed map" T^δ can be seen not only as a mathematical stratagem, but also as an implementation of the coverage algorithm for agents having arbitrarily small, but finite, communication ranges and moving inside their regions □

6.2.3 Analysis and convergence results

Before stating the main convergence results of the present chapter we introduce the following definition.

Definition 6.3 *A partition $v \in \mathcal{V}_N$ is said to be a* mixed equal-centroidal and centroidal-Voronoi partition *if, for all $(i, j) \in \mathcal{G}(v)$, either $\operatorname{Cd}(v_i) = \operatorname{Cd}(v_j)$ or (v_i, v_j) is a centroidal Voronoi partition of $v_i \cup v_j$.*

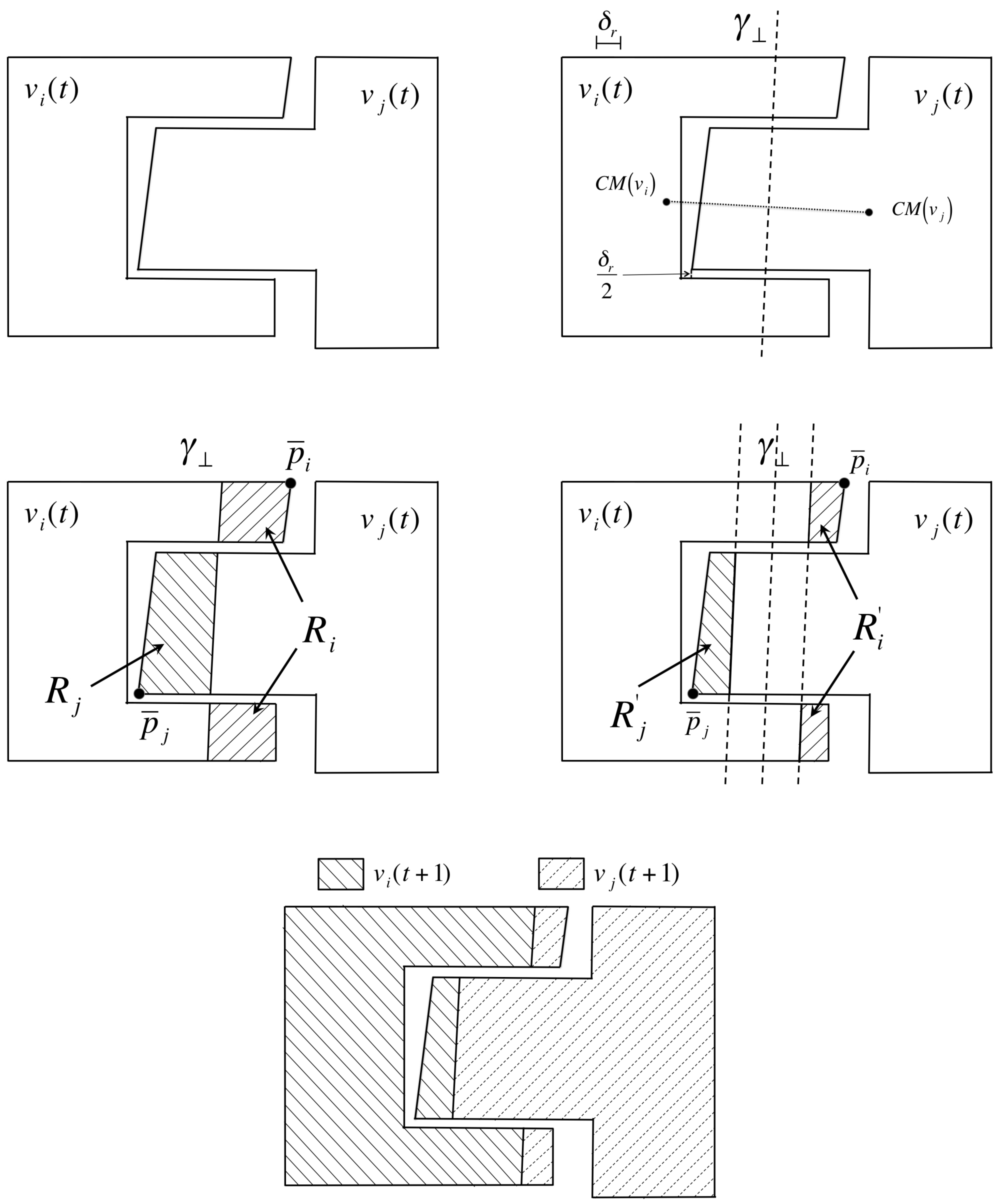

Figure 6.5: Demonstration of the update rule T^{δ} for a pair of close but not adjacent regions.

Denoting by $\mathcal{M}$ the set of mixed equal-centroidal and centroidal Voronoi partitions, we can equivalently write

$$\mathcal{M} = \{v \in \mathcal{V}_N \mid v = T_{ij} \, \forall i, j \in \{1, \dots, N\}, \, j \neq i\}. \tag{6.16}$$

Lemma 6.2 *Given a partition $v \in \mathcal{V}_N$, if (v_i, v_j) is not a centroidal Voronoi partition of $v_i \cup v_j$, then it exists $k \in \{1, \dots, N\}$ such that $(k, j) \in \mathcal{E}(v)$ and (v_k, v_j) is not a centroidal Voronoi partition of $v_k \cup v_j$.*

Proof: If $(i, j) \in \mathcal{E}(v)$ there is nothing to prove. Let us then assume $(i, j) \notin \mathcal{E}(v)$. Let $q \in v_i$ be closer to $\mathrm{Cd}(v_j)$ than to $\mathrm{Cd}(v_i)$, and let $j' \in \{1, \dots, N\}$ be such that $\mathrm{Cd}(v_{j'})$ is the closest centroid to q. If $(i, j') \in \mathcal{E}(v)$, we are done. Assume then that $(i, j') \notin \mathcal{E}(v)$. Then, there is an interposing region, say v_k, adjacent to v_j and different from v_i. By convexity of Q, we can find $q' \in v_k$ on the segment joining q and $\mathrm{Cd}(v_{j'})$. Then, by construction, q' is closer to $\mathrm{Cd}(v_{j'})$ than to any other centroid of regions. This concludes the proof. □

Remark 6.3 The above lemma means that forcing a partition, pairwise along the edges its adjacency graph, to be centroidal Voronoi, guarantees the same between all pairs of regions. We conclude that the set of the centroidal Voronoi partitions in $\mathcal{V}_N$ is a subset of $\mathcal{M}$. □

Now, we are ready to provide sufficient conditions for convergence of the algorithm defined in (6.15).

Theorem 6.3 (Worst case convergence) *Let $\delta > 0$, $v(0) \in \mathcal{V}_N$, and $v(t+1) \in T^{\delta}(v(t))$ for $t \in \mathbb{N}$. Assume that it exists $D \in \mathbb{N}$ such that for each pair $(i, j) \in \{1, \dots, N\}^2$, $i \neq j$, and for each time $t \in \mathbb{N}$, there exists $\bar{t} \leq t + D$ such that*

$$v(\bar{t} + 1) = T^{\delta}_{ij}(v(\bar{t})).$$

Then, all trajectories $v(t)$ converge to the set of the mixed equal-centroidal and centroidal-Voronoi partitions.

6.2.4 Simulation results and implementation remarks

We have extensively simulated the partitions-based gossip coverage algorithm described by (6.12), assuming that at each iteration, an edge is randomly chosen among all the edges belonging to the current set of edges of the adjacency graph. From these simulations, the effectiveness of the algorithm above introduced appears evident; all solutions converge to a partition belonging to the set of the *centroidal Voronoi partitions* (Figure 6.6). Indeed, modulo the event of regions having the very same centroid, Theorem 6.3 proves convergence to centroidal Voronoi partitions, as for the synchronous algorithm (6.10). In simulations such negative event does not display in generic conditions.

It is worth to remark that, although the algorithm allows the regions to be non convex, and even not connected, we have convergence to Voronoi regions, which are convex, and connected. Moreover, from the point of view of the applications, a connected region can be covered by an agent in a more natural way. These reasons suggest keeping the regions connected when applying the algorithm. Although we do not have theoretical results in this sense, we simulated a modification of the algorithm T which keeps the dominance regions connected: during the update step, every connected component is traded between the interacting regions only if this can be done without loosing connectivity. Simulations, in Figure 6.7, show that such an algorithm, if applied to partition a convex polygonal 2-dimensional area, leads to centroidal Voronoi partitions as well.

For the purpose of this article, simulations have been implemented as a *Matlab* program, using the *General Polygon Clipping C* language library to perform operations on polygons. Relevant issues for implementation, in simulations, as well as in actual robots, include the need for dealing with numerical rounding errors and the possibility for the polygons to have, in principle, a large number of vertices. Indeed, if we assume that every robot knows Q and ϕ, the information to be exchanged in pairs is the regions geometry, which can be summarized as a list of ordered lists of vertices. Thus an increase in their number affects both the computation and the communication burden on the agents. However, in practice the regions rarely get complicated shapes.

6.3 LaSalle invariance principle for set-valued maps

In this section we consider discrete-time continuous-space set-valued dynamical system defined on metric spaces. Our goal is to provide some extensions of the classical LaSalle Invariance Principle [60]; we refer the interested reader to [5, 51, 82] for recent Lasalle invariance principles for switched continuous-time and hybrid systems, respectively. We start by reviewing some notions including set-valued dynamical systems, closedness properties, Lyapunov function, and weak positive invariance [60].

Consider a metric space (X, d), where X is a set and d a metric on X. A set-valued map defined on (X, d) is a map $T : X \rightrightarrows X$ which associates to each point in X a subset in X, with the property that $T(x) \neq \emptyset$ for all $x \in X$. Note that a map from X to X can be interpreted as a singleton-valued map. An evolution of the dynamical system determined by T over X is any sequence $\{x_m\}_{m\in\mathbb{N}_0} \subset X$ with the property that

$$x_{m+1} \in T(x_m),\ m \in \mathbb{N}_0.$$

In other words, given any initial $x_0 \in X$, a trajectory of T is computed by recursively setting x_{m+1} equal to an arbitrary element in $T(x_m)$. An algorithm is therefore a nondeterministic discrete-time dynamical system.

Next, we introduce a notion of continuity for set-valued maps. A set valued map T is *closed* at $x \in X$ if for all pairs of convergent sequences $x_k \to x$ and $x'_k \to x'$ such that $x'_k \in T(x_k)$, one has that $x' \in T(x)$. An algorithm is closed on $W \subset X$ if it is closed at w, for all $w \in W$. In particular, every continuous map $T : X \to X$ is closed on X. A set C is weakly positively invariant with respect to T if, for any $x_0 \in C$, there exists $x \in T(x_0)$ such that $x \in C$. Let $U : X \to R$. We say that U is a Lyapunov function for T on W if (i) U is continuous on W and (ii) $U(w') \leq U(w)$ for all $w' \in T(w)$ and all $w \in W$. We are

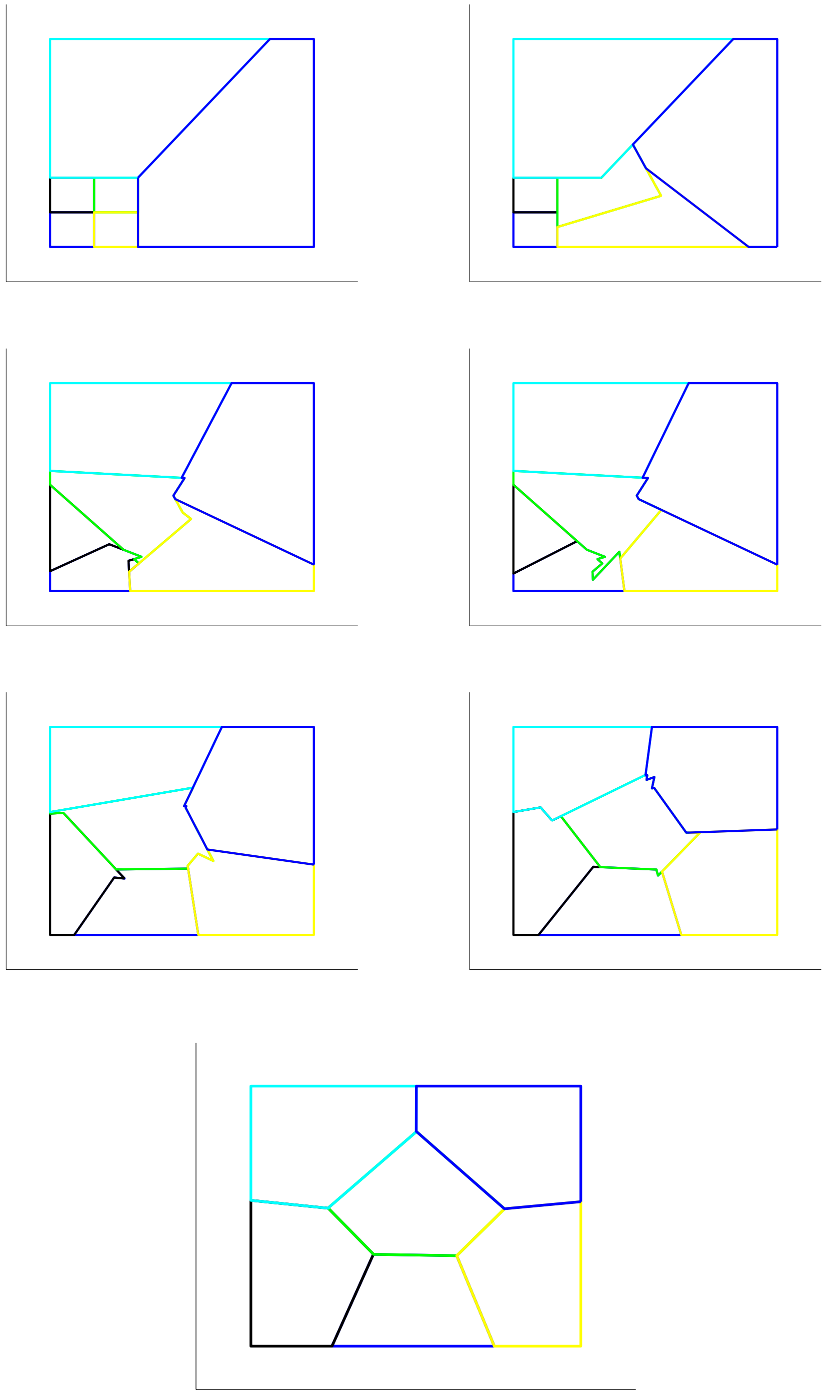

Figure 6.6: Simulations for algorithm T. The figure shows snapshots of a time evolution of the partitions, for $t = 0, 5, 20, 25, 50, 80, 150$. The density function ϕ is uniform. In consequence of this, the algorithm makes the regions converge towards an even partition of the given environment. Remark that the dominance regions can loose connectivity during the evolution.

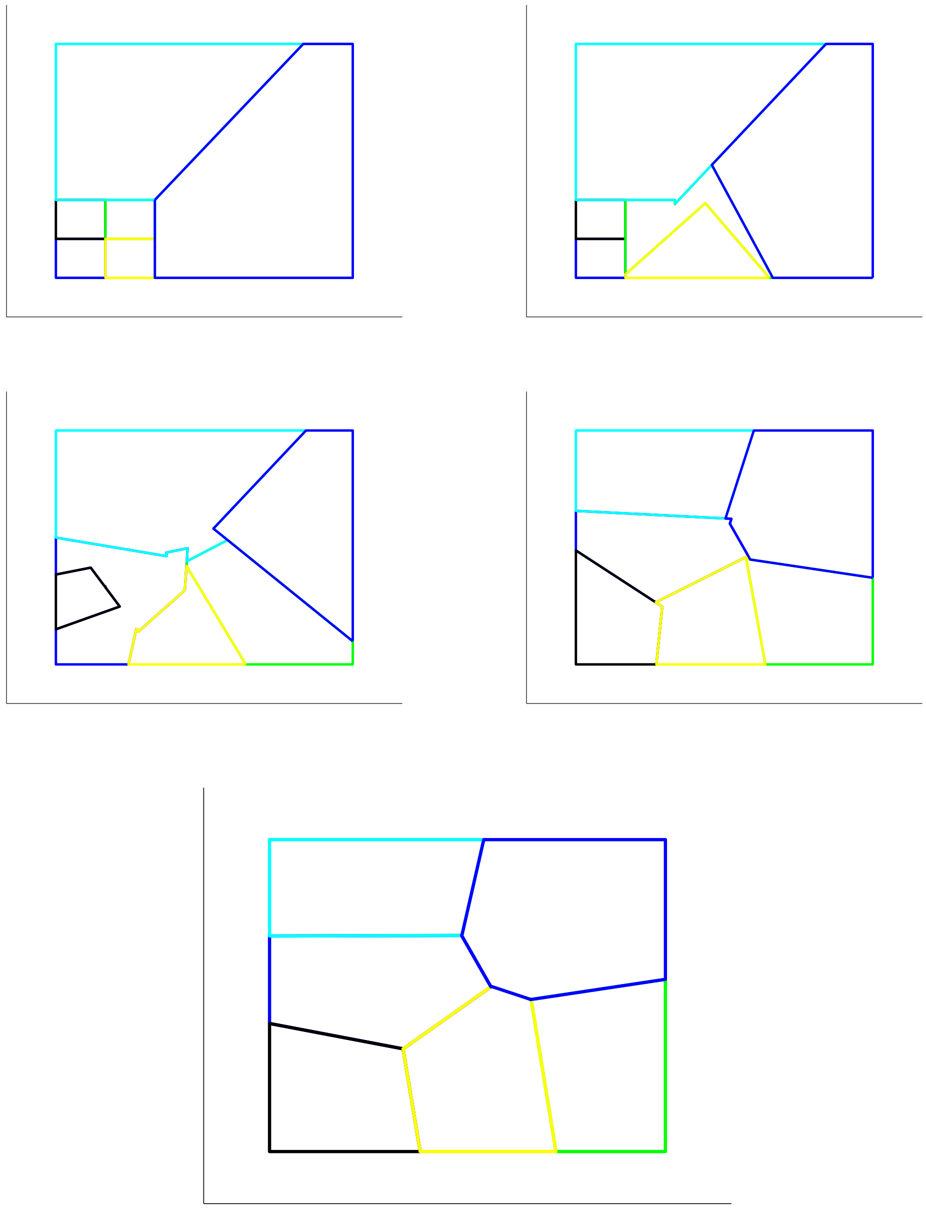

Figure 6.7: Simulations for algorithm T modified to keep regions connected. The figure shows snapshots of a time evolution of the partitions, for $t = 0, 5, 25, 50, 150$. The density function ϕ is uniform. In consequence of this, the algorithm makes the regions converge towards an even partition of the given environment. Remark that the dominance regions are not convex during the evolution, but are constrained to be connected.

ready now to state the following result which extends on one hand the LaSalle invariance principle to set-valued maps defined on metric spaces, and, on the other hand, the Global Convergence Theorem in [60] to more general Lyapunov functions. Its proof follows the lines of the proof of Theorem C.1 in [27], and is thus omitted.

Lemma 6.4 *Let (X, d) be a given metric space. Let $W \subseteq X$ be such that its closure $\overline{W}$ is compact. Let T be a* closed *set-valued map on W and let U be a Liapunov function for T on W. Let $w_0 \in W$, and assume the sequence $\{w_n \,|\, n \in \mathbb{N}_0\}$ defined via $w_{n+1} \in T(w_n)$ is in W. Then there exists $c \in \mathbb{R}$ such that*

$$w_n \to M \cap U^{-1}(c),$$

where M is the largest weakly positively invariant set contained in

$$\{x \in X \mid \exists\, y \in T(x) \quad \textit{such that} \quad U(y) = U(x)\} \cap \overline{W}\}.$$

In our chapter, given the metric space (X, d) and $W \subseteq X$, we deal with a particular set valued map $T : W \to 2^W$, which is defined by a collection of singleton-valued maps $T_1, \ldots, T_m : W \to W$. Specifically, given $w \in W$,

$$T(w) = \{T_1(w), \ldots, T_m(w)\}.$$

The following result provides an useful characterization of the closedness of this kind of set-valued maps, reducing it to the continuity of ordinary maps.

Lemma 6.5 *Let (X, d) be a given metric space. Let $W \subseteq X$ be such that its closure $\overline{W}$ is compact. Let $T_1, \ldots, T_m : W \to W$ be a finite collection of continuous functions. Define the set-valued map $T : W \to 2^W$ by*

$$T(w) = \{T_1(w), \ldots, T_m(w)\}.$$

Then the map T is closed on W.

Proof: Let $w_n \to w$ and $w'_n \to w'$ be a pair of convergent sequences in W, such that $w'_n \in T(w_n)$. We want to prove that the continuity of the maps $T_1, \ldots, T_m$ implies also $w' \in T(w)$.

We start by observing, that, by hypothesis, for all $n \in \mathbb{N}_0$ there exists $i_n \in \{1, \ldots, m\}$ such that $w'_n = T_{i_n}(w_n)$. Then, since the set $\{1, \ldots, m\}$ is finite, there exists $\bar{i} \in \{1, \ldots, m\}$ such that $\bar{i}$ appears infinitely many times in the sequence $\{i_n\}_{n\in\mathbb{N}}$. Consider the subsequences $\{w_{n_l}\} \subseteq \{w_n\}$ and $\{w'_{n_l}\} \subseteq \{w'_n\}$, such that $w'_{n_l} = T_{\bar{i}}(w_{n_l})$. Clearly, we have that $w_{n_l} \to w$ and $w'_{n_l} \to w'$, where from the continuity of $T_{\bar{i}}$ it follows that $w' = T_{\bar{i}}(w)$. Thus $w' \in T(w)$ and the claim is proved. □

Theorem 6.6 is a stronger version of Lemma 6.4, for a particular class of ”switching” maps.

Theorem 6.6 *Let (X, d) be a given metric space. Let $W \subseteq X$ be such that its closure $\overline{W}$ is compact. Let $T_1, \ldots, T_m : W \to W$ be a finite collection of continuous functions. Define the set-valued map $T : W \rightrightarrows W$ by*

$$T(w) = \{T_1(w), \ldots, T_m(w)\},$$

and let $U : X \to \mathbb{R}$ be a Liapunov function for T on W. Assume that, if $w' \in T(w)$ and $U(w') = U(w)$ then $w' = w$. Let $w_0 \in W$, and consider the sequence $\{w_n | n \in \mathbb{N}_0\}$ defined via $w_{n+1} \in T(w_n)$. Assume that there exists $D \in \mathbb{N}$ such that, for any $n \in \mathbb{N}_0$ and $j \in \{1, \ldots, m\}$, there exists $\bar{n} \in [n, n+D]$ such that $x_{\bar{n}+1} = T_j(x_{\bar{n}})$.

Then there exists $c \in \mathbb{R}$ such that

$$x_n \to K \cap U^{-1}(c),$$

where K is the largest weakly positively invariant set contained in

$$\{x \in X \mid U(y) = U(x) \quad \forall\, y \in T(x)\} \cap \overline{W}.$$

Proof: Let M be the largest weakly positively invariant set contained in

$$\{x \in X \mid \exists\, y \in T(x) \quad \text{such that} \quad U(y) = U(x)\} \cap \overline{W}.$$

Since T is closed by Lemma 6.5, the assumptions of Lemma 6.4 are met; hence there exists $c \in \mathbb{R}$ such that $U(w_n) \to c$ and $w_n \to M \cap U^{-1}(c)$.

Let $\Omega(w_n) \subseteq \overline{W}$ denote the ω-limit set of the sequence $\{w_n | n \in \mathbb{N}_0\}$. To prove the statement of the theorem we need to show that $\Omega(w_n) \subseteq K \cap U^{-1}(c)$. We proceed by contradiction. To this aim let $\bar{w} \in M \cap U^{-1}(c) \setminus K \cap U^{-1}(c)$ and let $\{w_{n_h} | h \in \mathbb{N}_0\}$ be a subsequence such that $w_{n_h} \to \bar{w}$. Without loss of generality, assume that there exists j', $1 \le j' < m$, such that $U(T_i(\bar{w})) = U(\bar{w})$ if $1 \le i \le j'$, and $U(T_i(\bar{w})) \neq U(\bar{w})$, if $j'+1 \le i \le m$. By hypothesis this implies also that $T_i(\bar{w}) = \bar{w}$, if $1 \le i \le j'$, and $T_i(\bar{w}) \neq \bar{w}$, if $j'+1 \le i \le m$. By the continuity of the maps T_i we can argue that there exists $\delta \in \mathbb{R}_{>0}$ such that, if $j'+1 \le i \le m$ then $T_i(w) \neq w$ for all $w \in B_\delta(\bar{w}) = \{w \in W : d(w, \bar{w}) \le \delta\}$.

Let now

$$\gamma = \min_{j'+1 \le i \le m} \left\{ \min_{w \in B_\delta(\bar{w})} \{U(w) - U(T_i(w))\} \right\}.$$

By hypothesis we have that, if $j'+1 \le i \le m$ then $U(T_i(w)) < U(w)$ for all $w \in B_\delta$. Hence, since $B_\delta(\bar{w})$ is closed, and U and the maps T_i are continuous, we can deduce that $\gamma > 0$. Consider now the set $\{T_i\}_{i=1,\ldots,j'}$; this is a collection of continuous maps having $\bar{w}$ as fixed point. Then, there exists a suitable $\varepsilon \in \mathbb{R}_{>0}$ such that, given any D-upla $(j_1, \ldots, j_D) \in \{1, \ldots j'\}^D$, we have that $T_{j_1} \circ T_{j_2} \circ T_{j_3} \circ \ldots \circ T_{j_D}(w) \in B_\delta(\bar{w})$ for all $w \in B_\varepsilon(\bar{w}) = \{w \in W \mid d(w, \bar{w}) \le \varepsilon\}$.

Let now $\bar{h}$ such that the element $w_{n_{\bar{h}}}$ in the subsequence $\{w_{n_h} | h \in \mathbb{N}_0\}$ satisfies $d\left(w_{n_{\bar{h}}}, \bar{w}\right) < \varepsilon$ and $U(w_{n_{\bar{h}}}) - c < \gamma$. Let

$$s = \min\{t \in [1, D] \mid \exists j \in \left\{j'+1, \ldots, m\right\} \text{ such that } w_{n_{\bar{h}}+t+1} = T_j\left(w_{n_{\bar{h}}+t}\right)\}$$

Observe that $U\left(w_{n_{\bar{h}}+s}\right) - c < \delta$ and $U\left(w_{n_{\bar{h}}+s}\right) - U\left(T_j\left(w_{n_{\bar{h}}+s}\right)\right) \ge \delta$ implying that $U\left(T_j\left(w_{n_{\bar{h}}+s}\right)\right) < c$. This is a contradiction. □

6.4 The set of partitions

Motivated by the results in Section 6.3, we introduce a metric structure on the set of partitions; specifically, we show how the set of partitions, endowed with a suitable metric, is a compact metric space. Additionally, we show the continuity properties of certain maps. In this section, and only in this section, the assumptions on Q are relaxed to give more general results: we assume that $Q \subset \mathbb{R}^d$ is *compact* and *connected* and has non-empty interior.

Recall that a metric space is a set S endowed with a scalar function of $S \times S$, called a metric, that satisfies the following properties: non-negativity, identity of indiscernibles, symmetry, and triangle inequality.

Metric structure of the set of closed subsets

Let $\mathcal{C}$ denote the *set of the closed subsets* of Q. We want to introduce a suitable metric and topology on $\mathcal{C}$; since the cost functions defined in Section 6.1 are insensitive to sets of zero measure, we look for a metric with the same property.

Let λ be the Lebesgue measure of a subset of $\mathbb{R}^d$. Given two subsets $A, B \in \mathcal{C}$, define their *symmetric difference* by $A\Delta B = (A \cup B) \setminus (A \cap B)$ and their *symmetric distance* $d_\Delta : \mathcal{C} \times \mathcal{C} \to \mathbb{R}_{\geq 0}$ by

$$d_\Delta(A, B) = \lambda(A\Delta B).$$

In other words, the symmetric distance is the measure of the symmetric difference of the two sets. Given these definitions, it is useful to identify sets that differ by a set of measure zero. More formally, let us write $A \sim B$ whenever $\lambda(A\Delta B) = 0$, and remark that $\sim$ is an equivalence relationship. In what follows we will study the *quotient set of closed subsets* $\mathcal{C}^* = \mathcal{C}/\sim$. The next result is the main result of this section.

Theorem 6.7 (Compactness of $\mathcal{C}^*$) *The pair $(\mathcal{C}^*, d_\Delta)$ is a metric space. Moreover, with the topology induced by the metric d_Δ, the set $\mathcal{C}^*$ is compact.*

Proof: It is easy to verify that d_Δ is a metric on $\mathcal{C}^*$. Instead, proving the compactness of the space is not trivial. First, we introduce a definition of convergence in $\mathcal{C}$, according to [67], which is not based on the metric d_Δ. Given a sequence of sets $\{E_n\}_{n\in\mathbb{N}}$, define

$$\begin{aligned}
\liminf_{n\to\infty} E_n \\
&= \{x \in Q \text{ such that if } U \text{ is an open subset of } Q \text{ and } x \in U,\ \text{then } U \cap E_n \neq \emptyset \\
&\qquad \text{for all but finitely many } n\},\\
\limsup_{n\to\infty} E_n \\
&= \{x \in Q \text{ such that if } U \text{ is an open subset of } Q \text{ and } x \in U,\ \text{then } U \cap E_n \neq \emptyset \\
&\qquad \text{for infinitely many } n\}.
\end{aligned}$$

We say that $\lim_{n\to\infty} E_n = E$ whenever $\liminf_{n\to\infty} E_n = \limsup_{n\to\infty} E_n = E$. It is known (e.g., see [67, Theorem 0.8]) that $\mathcal{C}$, with this definition of convergence, is a compact metric space. Thus, also $\mathcal{C}^*$, equipped with the quotient topology, is compact. It is almost immediate to check that, if $\lim_{n\to\infty} E_n = E$, then $\lim_{n\to\infty} \lambda(E_n \Delta E) = 0$. In other words, convergence as previously defined implies convergence in the topology induced by the d_Δ metric. From this we argue that $(\mathcal{C}^*, d_\Delta)$ is compact as well. $\square$

We now present some other properties of certain closed subsets. A set $C \in \mathcal{C}$ is *regularly closed* if $\overline{(\text{interior}(C))} = C$. Given a closed set $C \in \mathcal{C}$, let $\overline{\text{interior}(C)}$ be its *regularization.*

Lemma 6.8 *Let A and B be closed subsets of Q, that is, $A, B \in \mathcal{C}$.*

(i) If $d_\Delta(A, B) = 0$, then $\text{interior}(A) = \text{interior}(B)$.

(ii) If $d_\Delta(A, B) = 0$ and A, B are regularly closed, then $A = B$.

Proof: First, we prove (i). From the assumption, there exists a set M, having $\lambda(M) = 0$, such that $(B \setminus A) \subset M$. Then $B \subset (A \cup M)$. This implies $\text{interior}(B) \subset \text{interior}(A \cup M) = \text{interior}\,(A)$, because if there be a point belonging to $\text{interior}(A \cup M)$ and not to interior (A), then there be a whole open ball contained in the former, and then their measure should differ. The same argument applied to $A \setminus B$ completes the proof of (i). Statement (ii) is an immediate consequence. $\square$

Remark 6.4 (The set of regularly closed subsets) Since using a quotient space can be little intuitive, one would like to identify one representative in each class, and thus find a subspace of $\mathcal{C}$, which can be identified with $\mathcal{C}^*$. In view of Lemma 6.8(ii), a natural candidate is the set of regularly closed subsets. To show the identification is correct, one can prove that, if $C \in \mathcal{C}$, then $\lambda(C \Delta \overline{\text{interior}\,(C)}) = 0$. This would imply that the regularization of any closed set belongs to its class. Since $(C \Delta \overline{\text{interior}\,(C)}) \subset \partial C$, the claim is true if the set C had zero-measure boundary. However, this is not always the case for closed sets. Actually, the regularization of a closed set can differ from it by a set of finite measure. One can find [7] an open set A such that $\text{interior}\,(\overline{A}) = A$ and $\lambda(\partial A) > 0$. Thus, if we let $C = \overline{A}$, then $C \Delta \overline{\text{interior}\,(C)} = C \setminus \overline{\text{interior}\,(C)} = \partial A$, which has finite measure. Then the desired identification fails for general closed sets. $\square$

Metric structure of the set of partitions

Let us now introduce a metric structure on the space of partitions. The space of partitions $\mathcal{V}_N$, introduced in Section 6.1, is mapped by the canonical projection into a $\mathcal{V}_N^*$, whose components belong to $\mathcal{C}^*$. The metric d_Δ naturally extends to a metric on the product space $(\mathcal{C}^*)^N$ and on $\mathcal{V}_N^*$ as follows. The symmetric distance on partitions $d_\Delta : \mathcal{V}_N \times \mathcal{V}_N \to \mathbb{R}_{\geq 0}$ is defined by

$$d_\Delta(u, v) = \sum_{i=1}^{N} d_\Delta(u_i, v_i), \tag{6.17}$$

where $u = \{u_i\}_{i=1}^N$ and $v = \{v_i\}_{i=1}^N$ are any two N-partitions. The compactness of the space of partitions is then a simple consequence of Theorem 6.7.

Corollary 6.9 (Metric structure and compactness of $\mathcal{V}_N^*$) *The pair $(\mathcal{V}_N^*, d_\Delta)$ is a metric space. Moreover, with the topology induced by the metric d_Δ, the closure of $\mathcal{V}_N^*$ is a compact set.*

In the rest of the chapter, $\mathcal{V}_N^*$ and $\mathcal{V}_N$ are treated as one and the same: one may think to $\mathcal{V}_N$ as the space of the actual dynamics for the robots, and $\mathcal{V}_N^*$ as a space which is introduced for analysis purposes. Note that, thanks to the definition of $\mathcal{V}_N$, $\mathcal{V}_N^*$ can as well be depicted as a space of "partitions" made of regularly closed sets, representing the actual regions. Indeed, the identification between equivalence classes and regularly closed representatives, which can not be done (Remark 6.4) for general closed sets, can be done for sets satisfying the assumptions in Definition 6.1.

It can be checked[5] that all scalar functions of $\mathcal{C}$ or $\mathcal{V}_N$, that we use in this chapter, are independent of the representative which is chosen, but only depends on the equivalence class, that is, the functions are defined up to sets of measure zero. Thus, not only a sequence in $\mathcal{V}_N$ is mapped into a sequence in $\mathcal{V}_N^*$, but the dynamics in $\mathcal{V}_N$ induces a dynamics in $\mathcal{V}_N^*$; it is the latter dynamics that we are able to study. Some additional useful equivalence properties are stated as follows.

Corollary 6.10 *If u and v are two N-partitions with $d_\Delta(u, v) = 0$, then*

(i) the adjacency graph of u is equal to the adjacency graph of v;

(ii) u and v have the same regularization; and

(iii) if each set in u and v is regularly closed, then $u = v$.

6.5 Continuity of relevant functions and maps

Recall that the compact connected set Q is equipped with a bounded measurable positive function $\phi : Q \to \mathbb{R}_{>0}$, and that the centroid of any subset of Q is computed with respect to ϕ, see equation (6.4). Define the diameter of Q and the infinity norm of ϕ by $\operatorname{diam}(Q) = \{\|x - y\| \mid x, y \in Q\}$ and $\|\phi\|_\infty = \max_{x \in Q} \phi(x)$, respectively.

Lemma 6.11 *The centroid map $\operatorname{Cd} : \mathcal{C}^* \setminus \{\emptyset\} \to Q$, defined in equation (6.4), is continuous. Namely, given $A, A' \in \mathcal{C}^*$ such that $\lambda(A \cap A') > 0$,*

$$\|\operatorname{Cd}(A) - \operatorname{Cd}(A')\| \leq \frac{3\sqrt{d}\,\operatorname{diam}(Q)\,\|\phi\|_\infty}{\int_{A \cap A'} \phi(x)dx} d_\Delta(A, A'). \tag{6.18}$$

Proof: Let $p_i = (\operatorname{Cd}(A))_i$ and $p'_i = (\operatorname{Cd}(A'))_i$ be the i-th component of the two centroids. Let x_i be the i-th component of a generic point $x \in Q$. To keep the notation

[5] In some cases, this is obvious. We work out explicitly the non-obvious examples later in this section.

short, let $d\phi$ denote $\phi(x)dx$. Then

$$\begin{aligned}
|p_i - p'_i| &= \left| \frac{\int_A x_i d\phi}{\int_A d\phi} - \frac{\int_{A'} x_i d\phi}{\int_{A'} d\phi} \right| = \left| \frac{\int_A x_i d\phi \int_{A'} d\phi - \int_{A'} x_i d\phi \int_A d\phi}{\int_A d\phi \int_{A'} d\phi} \right| \\
&= \left| \frac{\left(\int_{A\setminus A'} x_i d\phi - \int_{A'\setminus A} x_i d\phi\right) \int_{A'\cap A} d\phi}{\int_A d\phi \int_{A'} d\phi} - \frac{\left(\int_{A\setminus A'} d\phi - \int_{A'\setminus A} d\phi\right) \int_{A'\cap A} x_i d\phi}{\int_A d\phi \int_{A'} d\phi} \right. \\
&\quad \left. + \frac{\int_{A\setminus A'} x_i d\phi \int_{A'\setminus A} d\phi - \int_{A'\setminus A} x_i d\phi \int_{A\setminus A'} d\phi}{\int_A d\phi \int_{A'} d\phi} \right| \\
&\leq \frac{\left|\left(\int_{A\setminus A'} x_i d\phi - \int_{A'\setminus A} x_i d\phi\right)\right| \int_{A'\cap A} d\phi}{\int_A d\phi \int_{A'} d\phi} + \frac{\left|\left(\int_{A\setminus A'} d\phi - \int_{A'\setminus A} d\phi\right)\right| \int_{A'\cap A} x_i d\phi}{\int_A d\phi \int_{A'} d\phi} \\
&\quad + \left| \frac{\int_{A\setminus A'} x_i d\phi \int_{A'\setminus A} d\phi - \int_{A'\setminus A} x_i d\phi \int_{A\setminus A'} d\phi}{\int_A d\phi \int_{A'} d\phi} \right|. \qquad (6.19)
\end{aligned}$$

We consider the denominator and write

$$\begin{aligned}
\int_A d\phi \int_{A'} d\phi &= \left(\int_{A\cap A'} d\phi + \int_{A\setminus A'} d\phi \right) \left(\int_{A'\cap A} d\phi + \int_{A'\setminus A} d\phi \right) \\
&= \left(\int_{A\cap A'} d\phi \right)^2 + \int_{A\cap A'} d\phi \int_{A\Delta A'} d\phi + \int_{A\setminus A'} d\phi \int_{A'\setminus A} d\phi.
\end{aligned}$$

Since the summands are all positive, each of them is smaller than their sum. Using this expansion, we upper bound each of the three summands in (6.19). The first term in (6.19) is upper bounded as follows:

$$\frac{\left|\int_{A'\setminus A} x_i d\phi - \int_{A\setminus A'} x_i d\phi\right| \int_{A'\cap A} d\phi}{\left(\int_{A\cap A'} d\phi\right)^2} \leq \frac{\int_{A\Delta A'} |x_i| d\phi}{\int_{A\cap A'} d\phi} \leq \frac{\operatorname{diam}(Q)\, \|\phi\|_\infty}{\int_{A\cap A'} d\phi} \lambda(A\Delta A'),$$

where we have assumed, without loss of generality, that Q contains the origin. The second term can be bounded similarly, and we upper bound the third term as follows:

$$\begin{aligned}
\frac{\left|\int_{A\setminus A'} x_i d\phi \int_{A'\setminus A} \phi dx - \int_{A'\setminus A} x_i d\phi \int_{A\setminus A'} \phi dx\right|}{\int_{A\cap A'} d\phi \int_{A\Delta A'} d\phi} &\leq \frac{\int_{A\Delta A'} d\phi \int_{A\Delta A'} |x_i| d\phi}{\int_{A\cap A'} d\phi \int_{A\Delta A'} d\phi} \\
&\leq \frac{\int_{A\Delta A'} |x_i| d\phi}{\int_{A\cap A'} d\phi} \\
&\leq \frac{\operatorname{diam}(Q)\, \|\phi\|_\infty}{\int_{A\cap A'} d\phi} \lambda(A\Delta A').
\end{aligned}$$

Finally, we obtain $|p_i - p'_i| \leq \dfrac{3\operatorname{diam}(Q)\|\phi\|_\infty}{\int_{A\cap A'} d\phi} \lambda(A\Delta A')$, from which the claim follows, since $Q \subset \mathbb{R}^d$. $\square$

Lemma 6.12 *The multicenter function $\mathcal{H}_{\text{centroid}} : \mathcal{V}_N \to \mathbb{R}_{\geq 0}$, defined in equation* (6.8)*, is Lipschitz continuous.*

Proof: Let u, v be two N-partitions. Then

$$
\begin{aligned}
|\mathcal{H}_{\text{centroid}}(u) - &\mathcal{H}_{\text{centroid}}(v)| \\
&= \left| \sum_{i=1}^{N} \int_{u_i} \|x - \operatorname{Cd}(u_i)\|^2 \phi(x) dx - \sum_{i=1}^{N} \int_{v_i} \|y - \operatorname{Cd}(v_i)\|^2 \phi(y) dy \right| \\
&\leq \sum_{i=1}^{N} \left| \int_{u_i} \|x - \operatorname{Cd}(u_i)\|^2 \phi(x) dx - \int_{v_i} \|y - \operatorname{Cd}(v_i)\|^2 \phi(y) dy \right| \\
&\leq \sum_{i=1}^{N} \Bigg(\left| \int_{u_i \setminus v_i} \|x - \operatorname{Cd}(u_i)\|^2 \phi(x) dx - \int_{v_i \setminus u_i} \|x - \operatorname{Cd}(v_i)\|^2 \phi(x) dx \right| \\
&\qquad + \left| \int_{u_i \cap v_i} (\|x - \operatorname{Cd}(u_i)\|^2 - \|x - \operatorname{Cd}(v_i)\|^2) \phi(x) dx \right| \Bigg) \\
&\leq \sum_{i=1}^{N} \Bigg(\left| \int_{u_i \setminus v_i} \|x - \operatorname{Cd}(u_i)\|^2 \phi(x) dx \right| + \left| \int_{v_i \setminus u_i} \|x - \operatorname{Cd}(v_i)\| \phi(x) dx \right| \\
&\qquad + \int_{u_i \cap v_i} \big| \|x - \operatorname{Cd}(u_i)\|^2 - \|x - \operatorname{Cd}(v_i)\|^2 \big| \phi(x) dx \Bigg) \\
&\leq \sum_{i=1}^{N} \left(\operatorname{diam}(Q)^2 \int_{u_i \Delta v_i} \phi(x) dx + \int_{u_i \cap v_i} 2 \operatorname{diam}(Q) \| \operatorname{Cd}(u_i) - \operatorname{Cd}(v_i) \| d\phi \right).
\end{aligned}
$$

In the last inequality we used the following fact:

$$
\begin{aligned}
\big| \|x - \operatorname{Cd}(v_i)\|^2 - &\|x - \operatorname{Cd}(v_j)\|^2 \big| \\
&\leq (\|x - \operatorname{Cd}(v_i)\| + \|x - \operatorname{Cd}(v_j)\|) \, \big| \|x - \operatorname{Cd}(v_i)\| - \|x - \operatorname{Cd}(v_j)\| \big| \\
&\leq 2 \operatorname{diam}(Q) \| \operatorname{Cd}(v_i) - \operatorname{Cd}(v_j) \|.
\end{aligned}
$$

Then using (6.18), we obtain

$$
\begin{aligned}
\sum_{i=1}^{N} &\left[\operatorname{diam}(Q)^2 \|\phi\|_\infty \, d_\Delta(u_i, v_i) + 2 \operatorname{diam}(Q) \| \operatorname{Cd}(v_i) - \operatorname{Cd}(u_i) \| \int_{u_i \cap v_i} d\phi \right] \\
&\leq \sum_{i=1}^{N} \left[\operatorname{diam}(Q)^2 \|\phi\|_\infty d_\Delta(u_i, v_i) + 6\sqrt{d} \, \operatorname{diam}(Q)^2 \|\phi\|_\infty d_\Delta(u_i, v_i) \right],
\end{aligned}
$$

and, finally,

$$
|\mathcal{H}_{\text{centroid}}(u) - \mathcal{H}_{\text{centroid}}(v)| \leq \operatorname{diam}(Q)^2 (6\sqrt{d} + 1) \, \|\phi\|_\infty \, d_\Delta(u, v).
$$

This bound implies Lipschitz continuity. □

Before we have studied the continuity of cost functions. From now on, we focus on the evolution maps. The following result is straightforward.

Lemma 6.13 (Properties of boundary) *Let A, B subsets of a metric space. The following three properties hold:*

(i) $\partial(A \cup B) \subset (\partial A \cup \partial B)$*;*

(ii) $\partial(A \cap B) \subset (\partial A \cup \partial B)$*;*

(iii) $\partial(A \setminus B) \subset (\partial A \cup \partial B)$*.*

Lemma 6.14 *The Voronoi map $V : Q^N \setminus S_N \to \mathcal{V}_N$, defined in equation* (6.2)*, is continuous.*

Proof: For simplicity, we discuss the two dimensional case. We first prove the continuity when $N = 2$: let p_1 and p_2 be points in Q. Let $2l = \|p_1 - p_2\|$. Since $p_1 \neq p_2$, $l > 0$. Up to isometries, we can assume that, in the Euclidean plane (x, y), $p_1 = (-l, 0)$ and $p_2 = (l, 0)$. Let d_1 and d_2 be the distances from the origin of points p_1 and p_2, respectively. It is clear that the two Voronoi regions of p_1 and p_2 are separated by the locus of points $\{x \in Q \mid \|x - p_1\| = \|x - p_2\|\}$, that is the vertical axis. Now, we assume that the positions of p_1 and p_2 are perturbed by a quantity less than or equal to δ, with $0 < \delta < l$. By effect of the perturbation, the axis separating the two Voronoi regions is perturbed, but it is contained in the locus of points $Y_{1\,2}(\delta) = \{x \in Q \mid |\|x - p_1\| - \|x - p_2\|| \leq 2\delta\}$. By definition, this is the set comprised between the two branches of the hyperbola whose equation is $\frac{x^2}{\delta^2} - \frac{y^2}{l^2 - \delta^2} = 1$. By elementary geometric considerations, the area of this region can be upper bounded by

$$\begin{aligned}
\lambda(Y_{12}(\delta)) &\leq 2\,\delta\,2\ \mathrm{diam}(Q) + 4\,\mathrm{diam}(Q)^2 \frac{\delta/l}{\sqrt{1 - \frac{\delta^2}{l^2}}} \\
&\leq 4\ \mathrm{diam}(Q)\Big(1 + \frac{\mathrm{diam}(Q)}{l}\Big)\delta.
\end{aligned}$$

This bound implies the continuity. The case in which $N > 2$ follows because, moving all points by at most δ, the change in all the regions is upper bounded by $\bigcup_{1 \leq i,j \leq N} Y_{ij}(\delta)$, which vanishes zero as $\delta \to 0^+$. □

Lemma 6.15 *Let A, B be closed subsets of Q, having non empty and disjoint interiors, and* $\mathrm{Cd}(A) \neq \mathrm{Cd}(B)$*. Then the map*

$$(A, B) \mapsto (A \cup B \cap H_{\mathrm{bisector}}(\mathrm{Cd}\,(A), \mathrm{Cd}\,(B)), A \cup B \cap H_{\mathrm{bisector}}(\mathrm{Cd}\,(B), \mathrm{Cd}\,(A)))$$

is continuous in $\mathcal{C}^*$, and the images have non-empty and disjoint interiors, and zero-measure boundaries.

Proof: The result contains several claims which deserve to be checked. Let A' and B' in $\mathcal{C}^*$ be the images of A and B. First, we need to prove that $A' \neq \emptyset$ and $B' \neq \emptyset$. This is a corollary of the following claim. Let $\mathrm{CC}(A)$ be the convex hull of the points of A, and let $x \in \mathrm{CC}(A)$ and $y \in A \cup B$. Then it exists v, a vertex of $\mathrm{CC}(A)$, such that $\|v - x\| < \|v - y\|$.

The fact that the images have non-empty interior can be proved as follows. Let x be an internal point of A. Then, either it is an interior point of A' or B', or x lies on the splitting axis. In the latter case, a point can be found in a neighborhood of x which is internal to A'.

The interiors are disjoint because $x \in \mathrm{interior}(A')$ if and only if $x \in \mathrm{interior}(A \cup B)$ and $\|x - \mathrm{Cd}(A)\| < \|x - \mathrm{Cd}(B)\|$. Thus clearly $x \notin \mathrm{interior}(B')$.

Since the images are obtained by intersecting the given sets with an half-space, which has zero-measure boundary, Lemma 6.13 implies that the boundaries of the images have zero measure, as well.

Since continuity is claimed in the quotient set, we need to check that the map does not depend on the representatives. Let us then take two other representatives, $C \sim A$ and $D \sim B$. Then there exist two zero-measure sets M and L such that $C \subset (A \cup M)$ and $D \subset (B \cup L')$. Since the centroids do not depend on zero-measure sets, the images of C and D are $C' \subset (A' \cup M \cup L)$ and $D' \subset (B' \cup M \cup L)$. The reverse argument, as in the proof of Lemma6.8, yields the claim.

Let us now discuss the continuity of the map. Let $\tilde{A}$ and $\tilde{B}$ be two sets in $\mathcal{C}$, and $\tilde{A}'$, $\tilde{B}'$ their images. We aim to upper bound the distance $d_\Delta(\tilde{A}', A') + d_\Delta(\tilde{B}', B')$. To do this, we combine the bounds which give continuity in the proofs of Lemmas 6.11 6.12 6.14. The difference of the images depends on the difference of the two argument regions both directly and via the induced difference between the centroids. Recalling the proof of Lemma 6.14, let δ be a bound on how much the two centroids are displaced. Then the region $Y(\delta) = \{x \in Q \mid |\|x - \mathrm{Cd}(A)\| - \|x - \mathrm{Cd}(B)\|| \leq 2\delta\}$ has to be included in the upper bound on the difference in the images. Combining these considerations we get

$$
\begin{aligned}
&d_\Delta(\tilde{A}', A') + d_\Delta(\tilde{B}', B') \\
&\quad \leq \Big(d_\Delta(\tilde{A}, A) + d_\Delta(\tilde{B}, B)\Big) + \lambda\Big(Y(\max\{\|\mathrm{Cd}(A) - \mathrm{Cd}(\tilde{A})\|, \|\mathrm{Cd}(B) - \mathrm{Cd}(\tilde{B})\|\})\Big) \\
&\quad \leq \left(1 + 12\sqrt{d}\,\mathrm{diam}(Q)^2\|\phi\|_\infty\left(1 + \frac{\mathrm{diam}(Q)}{\|\mathrm{Cd}(A) - \mathrm{Cd}(B)\|}\right)\frac{\int_{A\cap\tilde{A}}\phi dx + \int_{B\cap\tilde{B}}\phi dx}{\int_{A\cap\tilde{A}}\phi dx \int_{B\cap\tilde{B}}\phi dx}\right) \\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad \times \Big(d_\Delta(\tilde{A}, A) + d_\Delta(\tilde{B}, B)\Big).
\end{aligned}
$$

The bound implies continuity. □

Lemma 6.16 *For all $\delta > 0$, $(i,j) \in \{1,\ldots,N\}^2$, the map $T_{ij}^\delta : \mathcal{V}_N \to \mathcal{V}_N$, defined in Section 6.2.2, is well-posed and continuous.*

Proof: We have to establish three facts: (i) the image of the map is in $\mathcal{V}_N$; (ii) the map is defined up to sets of measure zero; and (iii) the map is continuous. Adopting the notations of Section 6.2.2, for example, we let $v \in \mathcal{V}_N$ and $\bar{v} = T^{\delta}_{ij}(v)$, we check these claims as follows.

(i) For all i, the set $\bar{v}_i$ is obviously closed, and by Lemma 6.13 $\lambda(\partial v_i) = 0$. For any pair $i, j \neq i$, similarly to Lemma 6.15, $\text{interior}(\bar{v}_i) \cap \text{interior}(\bar{v}_j) = \emptyset$, and since $\bar{v}_i \cup \bar{v}_j = v_i \cup v_j$, the collection of sets $\bar{v}$ covers Q.

(ii) The map is defined up to sets of measure zero. Indeed, the bisector half-space is by definition insensitive to sets of zero measure, and so are $\bar{p}_i$ and β, thanks to Lemma 6.8.

(iii) We use the following fact: let A, B, C and A', B', C' sets with $B \subset A$, $C \cap A = \emptyset$ and $B' \subset A'$, $C' \cap A' = \emptyset$, then

$$\left[((A \setminus B) \cup C) \Delta ((A' \setminus B') \cup C')\right] \subset \left[(A \Delta A') \cup (B \Delta B') \cup (C \Delta C')\right].$$

This implies that

$$\left[\bar{v}_i \Delta \bar{v}'_i\right] \subset [(v_i \Delta v'_i) \cup (\tilde{R}_i \Delta \tilde{R}'_i) \cup (\tilde{R}_j \Delta \tilde{R}'_j)].$$

Thus, we have to check that $\tilde{R}_i$ and $\tilde{R}_j$ depend continuously on v_i and v_j. But this fact is a simple consequence of $\tilde{R}_j$ being a subset of R_j, which depends continuously on v, by Lemma 6.15.

□

6.5.1 Convergence proofs

We are now almost ready to complete the proof of the convergence result presented in Section 6.2.3. The following result clarifies the relationship between $\mathcal{H}_{\text{centroid}}$ and T^{δ}.

Lemma 6.17 *Let $v \in \mathcal{V}_N$. Then $\mathcal{H}_{\text{centroid}}(T_{ij}(v)) \leq \mathcal{H}_{\text{centroid}}(v)$, and moreover $\mathcal{H}_{\text{centroid}}(T_{ij}(v)) = \mathcal{H}_{\text{centroid}}(v)$ if and only if $T_{ij}(v) = v$. The same hold if T_{ij} is replaced by T^{δ}_{ij}.*

We do not give the details of the proof: the result is a simple consequence of how the maps T_{ij} and T^{δ}_{ij} have been defined, using (6.9) and Proposition 6.1.

Theorem 6.3 then follows putting pieces together.

Proof: [Proof of Theorem 6.3] The goal is to apply Theorem 6.6 to the sequence $v(t)$ with $v(0) \in \mathcal{V}_N$ and $v(t+1) \in T^{\delta}(v(t))$, for all $t \in \mathbb{N}_0$. Note that from Lemma 6.16, the system evolves, through continuous mappings, in $\mathcal{V}_N$, which by Corollary 6.9 is subset of a compact metric space. The result follows because $\mathcal{H}_{\text{centroid}}$ is a Liapunov function for the system, meeting the assumptions of Theorem 6.6, thanks to Lemmas 6.12 and 6.17.

□

Chapter 7

Gossip in one dimensional partitions

In this chapter we study gossip partition optimization problems in one dimensional domains. With respect to the general case, the one dimensional case enjoys many properties which allow to make an easier and more complete analysis, which includes encouraging results on the speed of convergence of the algorithm. We are not surprised to obtain such results in the one dimensional case only, since also the evaluation of the speed of convergence of the Lloyd deployment (6.11) algorithm is an open question in higher dimensions [62].

7.1 Special properties

The first simplification is that, up to isometries and rescaling, we can assume that the domain to apportion be the unit interval $[0, 1]$. Moreover, we are going to show that we do not need to consider the most general space of partitions, but only partitions of $[0, 1]$ in N intervals. Let us consider the set of such partitions,

$$\mathcal{U}_N = \{\{[u_i, u_{i+1}]\} \mid 0 = u_0 \leq u_1, \dots, \leq u_i \leq u_{i+1} \leq \dots, \leq u_N = 1\}, \qquad (7.1)$$

with the understanding that the region $[u_{i-1}, u_i]$ is assigned to agent $i \in \{1, \dots, N\}$. It is clear that $\mathcal{U}_N$ is a subset of $\mathcal{V}_N$, and that the Voronoi partitions belong to $\mathcal{U}_N$. Indeed, the map $V : [0, 1]^N \to \mathcal{U}_N$ maps $p \in [0, 1]^N$ into $u = V(p)$ such that

$$u_i = \frac{p_i + p_{i+1}}{2} \quad \forall i \in \{1, \dots, N - 1\}.$$

Then $\mathcal{U}_N$ is a natural candidate as a state space for a partition optimization algorithm.

The metric structure of $\mathcal{U}_N$ is described by the following result.

Proposition 7.1 *Let $d_1 : \mathcal{U}_N \times \mathcal{U}_N \to \mathbb{R}$ be*

$$d_1(u, u') = \sum_{i=1}^{N-1} |u_i - u'_i|.$$

Then d_1 is a metric on $\mathcal{U}_N$, equivalent to d_Δ. Moreover $\mathcal{U}_N$ is compact.

Proof: Given two partitions $v, v' \in \mathcal{U}_N$,

$$d_H(v,v') = \sum_{i=1}^N \max\{|v_{i-1} - v'_{i-1}|, |v_i - v'_i|\},$$
$$d_1(v,v') = \sum_{i=1}^N |v_i - v'_i|,$$
$$d_\Delta(v,v') = \sum_{i=1}^N \min\{|v_{i-1} - v'_{i-1}| + |v_i - v'_i|, v_i - v_{i-1} + v'_i - v'_{i-1}\}.$$

From these formulas it is clear that

$$d_1(v,v') \le d_H(v,v') \le 2d_1(v,v'),$$

and

$$d_\Delta(v,v') \le 2d_1(v,v').$$

Then the Hausdorff metric d_H and the metric d_1 are equivalent, while the above bounds leave the possibility that the topology induced by d_Δ be coarser. However, the intervals are convex sets, and we know from [48] that in such a case the symmetric difference and Hausdorff metrics induce the same topology. We can also see this if we remark that

$$\min\{|v_{i-1} - v'_{i-1}| + |v_i - v'_i|, v_i - v_{i-1} + v'_i - v'_{i-1}\} = |v_{i-1} - v'_{i-1}| + |v_i - v'_i|$$

whenever $[v_{i-1}, v_i] \cap [v'_{i-1}, v'_i] \neq \emptyset$, and this condition is eventually true for subsequent element of a converging sequence of partitions.

Since we have shown that $\mathcal{U}_N$ is isomorphic to a subset of $\mathcal{V}_N$, the compactness of $\mathcal{U}_N$ follows from $\mathcal{U}_N$ being closed. $\square$

Remark 7.1 Remark that $u \in \mathcal{U}_N$ can also be seen as a vector $(u_i)_{i=1}^{N-1}$. This natural mapping into a subset of $[0,1]^{N-1}$ is an isomorphism, and this identification will be used in the sequel.

Remark 7.2 (Centroid map) The centroid map Cd, restricted to the intervals, has the simple form

$$\mathrm{Cd}([a,b]) = \frac{\int_a^b x\phi(x)dx}{\int_a^b \phi(x)dx}. \tag{7.2}$$

The map is well defined and continuous if $a < b$. The ill-definition of the map Cd for sets of zero measure con be resolved defining, for $a \in [0,1]$, $\mathrm{Cd}([a,a]) = a$. Such extension is clearly continuous.

Remark 7.3 (Cost function) The cost function is

$$\mathcal{H}_{\text{centroid}}(v) = \sum_{i=1}^{N} \int_{v_{i-1}}^{v_i} |x - \mathrm{Cd}([v_{i-1}, v_i])|^2 \phi(x) dx. \tag{7.3}$$

We proved in Lemma 6.12 that $\mathcal{H}_{\text{centroid}}$ is Lipschitz continuous. A stronger result can be proved in one dimension.

Proposition 7.2 *$\mathcal{H}_{\text{centroid}}$ is differentiable in $\mathcal{U}_N$, and its derivatives are, for $j = 1, \dots, N$,*

$$\frac{\partial \mathcal{H}_{\text{centroid}}}{\partial v_j} = 2\phi(v_j) \left(\mathrm{Cd}([v_j, v_{j+1}]) - \mathrm{Cd}([v_{j-1}, v_j])\right) \left(v_j - \frac{\mathrm{Cd}([v_j, v_{j+1}]) + \mathrm{Cd}([v_{j-1}, v_j])}{2}\right).$$

Proof: Recall the standard formula

$$\frac{d}{dx} \int_{\alpha(x)}^{\beta(x)} f(x,y) dy = \int_{\alpha(x)}^{\beta(x)} \frac{\partial f(x,y)}{\partial x} dy + \frac{d\beta(x)}{dx} f(x, \beta(x)) - \frac{d\alpha(x)}{dx} f(x, \alpha(x)).$$

Using this formula, from (7.2) we compute the derivatives

$$\frac{\partial\, \mathrm{Cd}([v_{i-1}, v_i])}{\partial v_i} = \frac{\phi(v_i)}{\int_{v_{i-1}}^{v_i} \phi(x) dx} \left(v_i - \mathrm{Cd}([v_{i-1}, v_i])\right)$$

and

$$\begin{aligned}
&\frac{\partial}{\partial v_i} \int_{v_{i-1}}^{v_i} \left(x - \mathrm{Cd}([v_{i-1}, v_i])\right)^2 \phi(x) dx \\
&\quad = \int_{v_{i-1}}^{v_i} \frac{\partial}{\partial v_i} \left(x - \mathrm{Cd}([v_{i-1}, v_i])\right)^2 \phi(x) dx + \phi(v_i) \left(v_i - \mathrm{Cd}([v_{i-1}, v_i])\right)^2 \\
&\quad = \phi(v_i) \left(v_i - \mathrm{Cd}([v_{i-1}, v_i])\right)^2 .
\end{aligned}$$

Since v_j is the extremum of two consecutive intervals,

$$\begin{aligned}
\frac{\partial \mathcal{H}_{\text{centroid}}}{\partial v_j} &= \phi(v_j) \left[(v_j - \mathrm{Cd}(v_{j-1}, v_j))^2 - (v_j - \mathrm{Cd}(v_j, v_{j+1}))^2\right] \\
&= \phi(v_j) \left(\mathrm{Cd}([v_j, v_{j+1}]) - \mathrm{Cd}([v_{j-1}, v_j])\right) \left(v_j - \frac{\mathrm{Cd}([v_j, v_{j+1}]) + \mathrm{Cd}([v_{j-1}, v_j])}{2}\right).
\end{aligned}$$

□

Remark 7.4 Since $\mathcal{H}_{\text{centroid}}$ is continuous and positive on the compact $\mathcal{U}_N$, it attains a positive global minimum. Moreover, the critical points of $\mathcal{H}_{\text{centroid}}$ are exactly the centroidal Voronoi partitions.

Gossip on a line graph

An important reason for the simplicity of the one dimensional case lies in the topology of the adjacency graphs.

Remark 7.5 (Adjacency graph) If $u \in \mathcal{U}_N$, then $\mathcal{G}(u)$ is the line graph of N nodes.

As a consequence of this, the pairwise algorithm (6.12), applied to partitions in $\mathcal{U}_N$, always features a map $T_{i\,i+1}$. For this reason, we simplify the notation, defining the map $T_i : \mathcal{U}_N \to \mathcal{U}_N$ such that $v' = T_i(v)$ is given by

$$v'_i = \frac{\mathrm{Cd}([v_{i-1}, v_i]) + \mathrm{Cd}([v_i, v_{i+1}])}{2},$$
$$v'_k = v_k \text{ if } k \neq i.$$

Correspondingly, the set-valued map $T : \mathcal{U}_N \rightrightarrows \mathcal{U}_N$ is such that v is mapped into $\{T_i(v)\}_{i=1}^N$, and we shall consider coverage algorithm generating a sequence of partitions $\{v(t)\}_{t\in\mathbb{N}_0}$ belonging $\mathcal{U}_N$: if agent i and $i+1$ interact at time t,

$$v(t+1) = T_i(v(t)).$$

Clearly, the maps T_i are continuous on $\mathcal{U}_N$, and then, by Lemma 6.5, T is a closed map. Moreover, the application of T preserves the communication graph. In force of the above remarks, Theorem 6.3 applies to sequences generated by T from $v(0) \in \mathcal{U}_N$, and then T drives the system to centroidal Voronoi partitions, that are the critical points of $\mathcal{H}_{\text{centroid}}$.

Remark 7.6 Whereas in the higher dimensional case the evolution of the partitions under a gossip algorithm can not be restricted to convex sets, thanks to the special properties of the one dimensional case we have been able to define a dynamics involving intervals, which enjoy the property of being convex.

7.2 Uniform density: equitable partitioning and average consensus

Formulas in the above sections are much simpler when the density is uniform, and stronger results can be obtained. Let us then assume in the sequel that ϕ is constant, and namely, with no further loss of generality, $\phi(x) = 1$. The map T is such that $v' = T_i(v)$ is given by

$$v'_i = \frac{v_{i-i} + 2v_i + v_{i+1}}{4}.$$

Proposition 7.3 *If $\phi(x) = 1$, then*

$$\begin{aligned}\mathcal{H}_{\text{centroid}}(v) &= \sum_{i=1}^{N} \int_{v_{i-1}}^{v_i} \left(\frac{v_{i-1}+v_i}{2} - x\right)^2 dx \\ &= \frac{1}{12}(v_i - v_{i-1})^3, \end{aligned} \tag{7.4}$$

and $\mathcal{H}_{\text{centroid}}$ *is a convex function.*

Proof: Specializing Proposition 7.2 or differentiating (7.4) we easily obtain

$$\begin{aligned}
\frac{\partial \mathcal{H}_{\text{centroid}}}{\partial v_j} &= \frac{v_{j+1} - v_{j-1}}{2}\left(v_j - \frac{v_{j-1} + v_{j+1}}{2}\right) \quad \text{for } j = 1, \dots, N-1, \\
\frac{\partial^2 \mathcal{H}_{\text{centroid}}}{\partial v_j^2} &= \frac{v_{j+1} - v_{j-1}}{2} \quad \text{for } j = 1, \dots, N-1, \\
\frac{\partial^2 \mathcal{H}_{\text{centroid}}}{\partial v_j \partial v_{j+1}} &= -\frac{v_{j+1} - v_j}{2} \quad \text{for } j = 1, \dots, N-2, \\
\frac{\partial^2 \mathcal{H}_{\text{centroid}}}{\partial v_j \partial v_{j-1}} &= -\frac{v_j - v_{j-1}}{2} \quad \text{for } j = 2, \dots, N-1.
\end{aligned}$$

The second derivatives can be arranged in an Hessian matrix $H \in \mathbb{R}^{N-1 \times N-1}$, which turns out to be symmetric and tridiagonal, with diagonal entries $d_i = \frac{1}{2}(v_{j+1} - v_{j-1})$, for $i = 1, \dots, N-1$, and super-diagonal entries $s_i = -\frac{1}{2}(v_{j+1} - v_j)$, for $i = 1, \dots, N-2$. Since $|s_i| + |s_{i-1}| \leq d_i$, from Gershgorin Theorem we deduce that its eigenvalues are all positive. Then $\mathcal{H}_{\text{centroid}}$ is a convex function. □

Some more remarks are in order.

Remark 7.7 (i) The critical points of $\mathcal{H}_{\text{centroid}}$ are such that $v_j = \frac{v_{j+1}+v_{j-1}}{2}$ $\forall j \in \{1, \dots, N-1\}$. Remark that this corresponds to the configuration with equispaced points $v_i^* = \frac{i-1}{N}$, and $\mathcal{H}_{\text{centroid}}(v^*) = \frac{1}{12}\frac{1^2}{N^2}$. Such a partition is called *equitable* [74].

(ii) $\mathcal{U}_N$ is the intersection of an $(N-1)$-hypercube and of $N-1$ half-spaces, and then a convex set. Then, since $\mathcal{H}_{\text{centroid}}$ is a convex function, the critical point is the (unique) global minimum.

(iii) Since $\phi(x) = 1$, the configurations with equispaced points are the centroidal Voronoi partitions of $[0, 1]$. Indeed, $v^* = T_i(v^*)$ for all $i \in \{1, \dots, N\}$.

Motivated by the form of the critical points of $\mathcal{H}_{\text{centroid}}$, we can also suggest a different algorithm, defined by the map $T_i^\infty : \mathcal{U}_N \to \mathcal{U}_N$, such that $v' = T_i^\infty(v)$ is

$$\begin{aligned}
v'_i &= \frac{v_{i-1} + v_{i+1}}{2} \\
v'_k &= v_k \quad \text{for } k \neq i\ .
\end{aligned}$$

It is clear that also T_i^∞ is continuous for all i, and then the set-valued map $T^\infty : \mathcal{U}_N \rightrightarrows \mathcal{U}_N$, such that $T^\infty(v) = \{T_i^\infty(v)\}_{i=1,\dots,N}$, is closed. Moreover, $\mathcal{H}_{\text{centroid}}$ is a Liapunov function for T^∞. Then we can conclude that Theorem 6.3 applies also to sequences generated by T^∞.

Coverage and consensus

Further insights can be obtained if we remark that the one dimensional deployment problem with uniform density can be cast into a consensus problem. This allows to find stronger convergence results, and some extra information about the speed of convergence of the algorithm.

Remark 7.8 (Coverage as optimization on the simplex) Let Σ be the $(N-1)$-dimensional standard simplex. Let $f_1 : \mathcal{U}_N \to \Sigma$ be such that $x = f_1(u)$ is the vector of components $x_i = u_i - u_{n-1}$, for $i = \{1, \dots, N\}$. Then f_1 is one-to-one and onto. As a consequence, for any ϕ, every one-dimensional coverage algorithm is equivalent to an optimization algorithm evolving on the simplex.

In the case $\phi(x) \equiv 1$, the optimization goal is to reach an equispaced distribution on $[0, 1]$, that is, on the simplex, to make every x_i reach the value $1/N$. This is exactly an *average consensus problem.* Average consensus problems, which have been studied in the Part I of the thesis, can be solved by pairwise communication, with either a deterministic [65] or a randomized [9] communication schedule.

Now we explicitly show that the proposed deployment algorithms correspond on the simplex to average consensus algorithms, which can be studied thanks to the tools recently developed in [24, 36].

Synchronous Lloyd algorithm

We start with the synchronous algorithm (6.11). In the simplex, it induces the map $E : \Sigma_{N-1} \to \Sigma_{N-1}$, with $x' = E(x)$ given by

$$
\begin{aligned}
x'_1 &= \frac{3x_1 + x_2}{4} \\
x'_i &= \frac{x_{i-1} + 2x_i + x_{i+1}}{4} \qquad \text{for all } i = 2, \dots, N-1 \\
x'_N &= \frac{x_{N-1} + 3x_N}{4}
\end{aligned}
$$

We can write

$$x(t+1) = E\, x(t), \tag{7.5}$$

where E is a suitable tridiagonal symmetric matrix.

Proposition 7.4 *Let $x(t)$ evolve following* (7.5). *Then $\lim_{t\to\infty} x(t) = x^*$, with $x^*_i = 1/N$ for all $i = 1, \dots, N$.*

Proof: Let an *augmented tridiagonal Toepliz matrix* be a matrix

$$
\mathrm{ATrid}^+(a, b) = \begin{pmatrix}
b+a & a & 0 & \dots & 0 \\
a & b & a & \dots & \vdots \\
\vdots & \dots & \dots & \dots & \vdots \\
0 & \dots & a & b & a \\
0 & \dots & 0 & a & b+a
\end{pmatrix}.
$$

According to [14], the matrix $\mathrm{ATrid}^+(a,b)$ has eigenvalues $b+2a\cos\frac{l\pi}{N}$, for $l=0,\dots,N-1$. Then, $E=\mathrm{ATrid}^+(1/4,1/2)$ has eigenvalues

$$\frac{1}{2}+\frac{1}{2}\cos\frac{l\pi}{N} \quad l=0,\dots,N-1.$$

This implies [24] the claim. □

We can define[1] the rate of convergence of such a consensus algorithm as

$$R_{synch}=\frac{1}{N}\sup_{x(0)}\limsup_{t\to\infty}\left(\|x(t)-x^*\|^2\right)^{1/t}.$$

The rate, from Proposition 1.5, is equal to the square of the second largest (in modulus) eigenvalue of E, that is

$$R_{synch}=\frac{1}{4}\left(1+\cos\frac{\pi}{N}\right)^2 \asymp 1-\frac{\pi^2}{2N^2}, \text{ for } N\to\infty.$$

Pairwise gossip algorithm

The same rewriting can be done for the two pairwise algorithms proposed above, T and T^∞. Indeed, the application of T_i and T_i^∞ can be read on the simplex as follows. Let $v'=T_i(v)$. Correspondingly, on the simplex we have $x'=S_i(x)$ with

$$\begin{aligned} x'_i &= \frac{3x_i+x_{i+1}}{4} \\ x'_{i+1} &= \frac{x_i+3x_{i+1}}{4} \\ x'_k &= x_k \text{ otherwise .} \end{aligned}$$

Similarly, denoting $v''=T_i^\infty(v)$, we have on the simplex $x''=S_i^\infty(x)$ with

$$x''_i=x''_{i+1}=\frac{x_i+x_{i+1}}{2},$$

while all other components are unchanged.

From now on, we concentrate on the algorithm S^∞: the algorithm S can be analyzed in a similar fashion. The evolution law of S^∞ can be rewritten more compactly as $x(t+1)=A(t)x(t)$. At each time step $A(t)\in\{A_i\}_{i=1}^{N-1}$, where A_i is the N-dimensional matrix

$$A_i=\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1/2 & 1/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1/2 & 1/2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix},$$

[1] This definition differs by a square from the one in Chapter 1.

in which the block $\begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}$ appears on the diagonal in position $i, i+1$.

This allows to prove strong convergence results on the gossip coverage algorithms in dimension one.

Theorem 7.5 *Let $u(0) \in \mathcal{U}_N$ and $u(t+1) \in T(u(t))$ for all $t \in \mathbb{N}_0$. Assume $\phi \equiv 1$. Then $u(t)$ converges exponentially to the only centroidal Voronoi partition, if and only if, for all $i \in \{1, \dots, N\}$, $u(t+1) = T_i(u(t))$ for infinitely many t.*

Proof: Since the underlying topology is the line graph, Proposition 2 of [65] implies that the system S makes $x(t)$ converge to x^* if and only if every edge is activated an infinite number of times. $\square$

Moreover, we consider a randomized schedule, in which each pair of adjacent regions is randomly selected with probability p_i at every time step $t \geq 0$, that is

$$\mathbb{P}\left[A(t) = A_i\right] = p_i. \tag{7.6}$$

Theorem 7.6 *Let $u(0) \in \mathcal{U}_N$ and $u(t+1) \in T(u(t))$ for all $t \in \mathbb{N}_0$. Assume $\phi \equiv 1$. Assume there exists $\bar{p}$, $0 < \bar{p} < 1$, such that, for all $i \in \{1, \dots, N\}$,*

$$\mathbb{P}\left[u(t+1) = T_i(u(t))\right] \geq \bar{p}.$$

Then $u(t)$ converges almost surely to the only centroidal Voronoi partition.

Proof: It is easy to check that $\mathbb{E}[A(t)]$ is a doubly stochastic and irreducible matrix, and every A_i has positive diagonal. Hence, we can argue the result applying Corollary 3.2 in [36]. $\square$

For simplicity, let us assume in the sequel that $p_i = \frac{1}{N-1}$ for all $i \in \{1, \dots, N\}$. Then

$$\mathbb{E}[A(t)] = I + \frac{1}{2}\frac{1}{N-1}\,\mathrm{ATrid}^+(1, -2).$$

Following [36], we develop a mean squared error analysis of the probabilistic algorithm to study its speed of convergence. Let us define the speed of convergence of the randomized algorithm as

$$R = \sup_{x(0)} \limsup_{t\to\infty} \mathbb{E}\left[\frac{1}{N}\|x(t) - x^*\|^2\right]^{1/t}.$$

Then we can prove the following result.

Theorem 7.7 *For the algorithm T^∞ with a randomized schedule* (7.6) *and $p_i = \frac{1}{N-1}$,*

$$1 - \frac{2}{N-1}\left(1 - \cos\frac{\pi}{N}\right) \leq R_{T^\infty} \leq 1 - \frac{1}{N-1}\left(1 - \cos\frac{\pi}{N}\right).$$

Then $R_{T^\infty} = 1 - \Theta(N^{-3})$ for $N \to \infty$.

Proof: The proof is based on Proposition 4.4 in [36]. If we denote by $\rho(M)$ the second largest in modulus eigenvalue of the matrix M, then a randomized algorithm has a rate R estimated by

$$\rho(\mathbb{E}[A(t)])^2 \leq R \leq \rho(\mathbb{E}[A(t)^2]).$$

Since in this case $A_i^2 = A_i$, the bounds become

$$\rho(\mathbb{E}[A(t)])^2 \leq R \leq \rho(\mathbb{E}[A(t)]).$$

Then we compute, using [14] as above, $\rho(\mathbb{E}[A(t)]) = 1 - \frac{1}{N-1}\left(1 - \cos\frac{\pi}{N}\right)$, and we conclude. □

A similar analysis leads to a similar conclusion about T.

Theorem 7.8 *The algorithm T, with a randomized schedule* (7.6) *and $p_i = \frac{1}{N-1}$, converges almost surely to the centroidal Voronoi partition, and its rate of convergence is estimated as*

$$1 - \frac{1}{N-1}\left(1 - \cos\frac{\pi}{N}\right) \leq R_T \leq 1 - \frac{1}{N-1}\frac{3}{4}\left(1 - \cos\frac{\pi}{N}\right).$$

Then $R_T = 1 - \Theta(N^{-3})$ for $N \to \infty$.

Remark 7.9 The speed of convergence of T and T^∞ has the same order in N, but T is not faster than T^∞. It is also interesting to compare these results with the rate of convergence of the synchronous deterministic deployment algorithm R_{synch}, which degrades for $N \to \infty$ as $O(N^2)$. It is clear that the loss of performance is limited to one power of N, which is due to the pairwise communication schedule.

It is then interesting to compare the algorithms in terms of the required number of communication rounds needed to converge, the so called *total communication complexity* [61]. For a given algorithm whose rate of convergence is R, we define

$$\mathrm{TCC} = \frac{\#\text{ active edges per time step}}{1 - R}.$$

It comes out that, for N large enough,

$$\begin{aligned} &\mathrm{TCC}_{synch} = \frac{2N^2}{\pi^2}(N-1), \\ &2\frac{N^2}{\pi^2}(N-1) \leq \mathrm{TCC}_T \leq \frac{8}{3}\frac{N^2}{\pi^2}(N-1), \\ &\frac{N^2}{\pi^2}(N-1) \leq \mathrm{TCC}_{T^\infty} \leq 2\frac{N^2}{\pi^2}(N-1), \end{aligned}$$

and then

$$\mathrm{TCC}_{T^\infty} \leq \mathrm{TCC}_{synch} \leq \mathrm{TCC}_T\,.$$

We conclude that, from the point of view of the communication complexity, the gossip algorithms have the *same* order in N as the synchronous algorithm, and T^∞ improves the constants.

7.3 Robots motion protocols

An important reason of interest in the partitions based approach is the following. Since the actual position of the agents plays essentially no role, agents do not need to always stay at the centroid. We are going to show that this allows to consider agent models in which the agents have *very short communication radiuses.* In this section we sketch some robot motion protocols which assure the robots to communicate, so that pairwise communications are effective for the deployment task. In the natural application, agents have a continuous time dynamics in the domain: as a limit model we can consider the agents as point particles, and able to communicate when they collide. When a pair of agents collides and communicates, the algorithm makes one step. In the one dimensional case, it is easy to design movement strategies for the agents such that the requirements of the convergence theorems are met. We present two such different strategies.

(i) The simplest idea is the following one. Let us think to the agents as moving with a continuous time dynamics on the segment $[0, 1]$ at (common) constant speed c. Each agent k has a target region $[v_{k-1}, v_k]$. When two of them, say i and $i + 1$, collide, they communicate, so that the updated state of the system is $v^+ = T_i(v)$, and they invert their directions. Then the agent i will collide with agent $i + 1$ and $i - 1$ alternately. It is clear that in a (continuous) time equal to $2/c$ all agents have communicated with both neighbors: then this scheme satisfies Theorem 6.3. This movement strategy has two additional good properties. The first one is a natural robustness to agents failures. The second one lies in the fact that the agents can self-initialize their assigned regions in a very natural way. When agent i and $i + 1$ meet for the first time, say in position $x \in (0, 1)$, they set $v_i(0) = x$. At their second meeting, they will have already initialized $v_{i-1}(0)$ and $v_{i+1}(0)$ respectively.

    We conclude that this scheme is effective for deployment of agents in a one-dimensional environment, with arbitrary initial conditions. We can compare it with the algorithm in [83]. In that work, the authors design an algorithm for agents on a ring to achieve both deployment and synchronization of their dynamics. For the algorithm in [83], only a local convergence result is proven, whereas the algorithm we proposed above does converge for any initial condition, but does not guarantee synchronization of the dynamics.

(ii) Scheme (i) allows the agents to possibly travel outside their regions. If, depending on the application, this shall be avoided, one can devise the following modification of the scheme (i). When an agent reaches the border of its region, it stops there until its neighbor comes to collide and communicate. Then, they adjourn their states as in the scheme (i), and move away. This implementation keeps the agents inside their regions, but is not intrinsically robust to agents failures.

The key point which allows the design of the above strategies is that the communication graph is time invariant and known to the agents. These facts are no longer true in higher dimensions, and this makes the design of agents strategies much harder. Such a problem will be the object of further research.

Chapter 8

Perspective on partitions optimization

The novel approach to coverage control which has been introduced in this thesis, and presented in [38], is based on considering algorithms evolving partitions, rather than robots positions. This naturally brings the need of dealing with a space of partitions, which has been a technical challenge. Moreover, the novel partitions-based approach, and the gossip paradigm, appear very promising of extensions and applications. We list here some of the most significant.

(i) *Algorithmical improvements.* The specific algorithm, and the convergence proof, that we have given in Chapter 6, leave the possibility of getting stuck in non-Voronoi partition, due to coincident centroids. A possible solution to this could be an algorithm in which, when two interacting regions share the same centroid, instead of keeping the regions unchanged as T^{δ} does, the union of the two regions is split following some other rule. This proposal seems to be able to avoid the cul-de-sacs due to coincident centroids, but its analysis looks difficult due to continuity issues.

(ii) *Extension to different objective functions.* We expect that the results can be extended, non trivially, to consider more general multi-center [11] function, including, e.g, Weber or disk-covering problems.

(iii) *Extension to non-convex environments.* Parts of our analysis hold with the weaker assumption of Q being compact: this encourages to apply the partition-based approach to non-convex environments. About the coverage of non-convex environments, much research is in progress [43], [77]: a natural extension would be based on computing distances in Q as the geodesic distances.

(iv) *Discrete partitioning.* We would like to apply these ideas to optimize partition of environments other than subsets of $\mathbb{R}^d$,for instance discrete objects like graphs. Besides the theoretical interest, this would be an important step towards an effective implementation on actual robots.

(v) *Ad-hoc motion planning.* The abstraction from the agents positions opens the possibility of designing motion strategies (see Chapter 7) for the agents, which allow them to communicate in spite of extremely short communication radiuses, and to perform other tasks (patrolling, target servicing) at the same time.

Bibliography

[1] A. Arsie, K. Savla, and E. Frazzoli. Efficient routing algorithms for multiple vehicles with no explicit communications. *IEEE Transactions on Automatic Control*, 2009. To appear.

[2] T. Aysal, M. Coates, and M. Rabbat. Distributed average consensus with dithered quantization. *IEEE Transactions on Signal Processing*, 56(10):4905–4918, 2008.

[3] T. C. Aysal, M. Coates, and M. Rabbat. Distributed average consensus using probabilistic quantization. In *IEEE Workshop on Statistical Signal Processing*, pages 640–644, Maddison, Wisconsin, Aug. 2007.

[4] L. Babai. Spectra of Cayley graphs. *J. Combin. Theory Ser. B*, 27(2):180–189, 1979.

[5] A. Bacciotti and L. Mazzi. An invariance principle for nonlinear switched systems. *Systems & Control Letters*, 54(11):1109–1119, 2005.

[6] G. Baldan and S. Zampieri. An efficient quantization algorithm for solving average-consensus problems. In *European Control Conference*, Budapest, Hungary, Aug. 2009.

[7] R. Börger. A non-Jordan-measurable regularly open subset of the unit interval. *Archiv der Mathematik*, 73(4):262–264, 1999.

[8] S. Boyd, P. Diaconis, and L. Xiao. Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4):667–689, Dec. 2004.

[9] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Transactions on Information Theory*, 52(6):2508–2530, 2006.

[10] R. W. Brockett and D. Liberzon. Quantized feedback stabilization of linear systems. *IEEE Transactions on Automatic Control*, 45(7):1279–1289, 2000.

[11] F. Bullo, J. Cortés, and S. Martínez. *Distributed Control of Robotic Networks*. Applied Mathematics Series. Princeton University Press, 2009. To appear. Available at http://www.coordinationbook.info.

[12] C. H. Caicedo-Nùñez and M. Žefran. Performing coverage on nonconvex domains. In *IEEE Conf. on Control Applications*, pages 1019–1024, San Antonio, TX, Sept. 2008.

[13] G. C. Calafiore and F. Abrate. Distributed linear estimation over sensor networks. *International Journal of Control*, 2008. Accepted.

[14] R. Carli and F. Bullo. Quantized coordination algorithms for rendezvous and deployment. *SIAM Journal on Control and Optimization*, Dec. 2007. Accepted.

[15] R. Carli, A. Chiuso, L. Schenato, and S. Zampieri. Distributed kalman filtering based on consensus strategies. *IEEE Journal on Selected Areas in Communications*, 26(4):622–633, 2008.

[16] R. Carli, A. Chiuso, L. Schenato, and S. Zampieri. A PI consensus controller for networked clocks synchronization. In *IFAC World Congress*, pages 10289–10294, Seoul, Korea, July 2008.

[17] R. Carli, G. Como, P. Frasca, and F. Garin. Average consensus on digital noisy networks. In *IFAC Workshop on Estimation and Control of Networked Systems*, 2009. Submitted.

[18] 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.

[19] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri. Efficient quantization techniques in consensus algorithms. In *NeCST Workshop*, Nancy, France, June 2007.

[20] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri. A probabilistic analysis of the average consensus algorithm with quantized communication. In *IFAC World Congress*, Seoul, Korea, July 2008.

[21] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri. The quantization error in the average consensus problem. In *Mediterranean Conference on Control and Automation*, Ajaccio, France, June 2008.

[22] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri. Efficient quantization for average consensus. arXiv:0903.1337, 2009.

[23] R. Carli, F. Fagnani, P. Frasca, and S. Zampieri. Gossip consensus algorithms with quantized communication. *Automatica*, 2009. Submitted.

[24] R. Carli, F. Fagnani, A. Speranzon, and S. Zampieri. Communication constraints in the average consensus problem. *Automatica*, 44(3):671–684, 2008.

[25] F. R. K. Chung. *Spectral graph theory*, volume 92 of *CBMS Regional Conference Series in Mathematics*. Conference Board of the Mathematical Sciences, Washington, DC, 1997.

[26] R. Cogburn. The ergodic theory of Markov chains in random environments. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 66(1):109–128, 1984.

[27] J. Cortés, S. Martínez, and F. Bullo. Spatially-distributed coverage optimization and control with limited-range interactions. *ESAIM. Control, Optimisation & Calculus of Variations*, 11:691–719, 2005.

[28] J. Cortés, S. Martínez, and F. Bullo. Robust rendezvous for mobile autonomous agents via proximity graphs in arbitrary dimensions. *IEEE Transactions on Automatic Control*, 51(8):1289–1298, 2006.

[29] 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.

[30] V. Coscia, M. Delitala, and P. 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.

[31] D. Cvetovic, M. Doob, and H. Sachs. *Spectra of graphs*. Academic Press, 1980.

[32] G. Cybenko. Dynamic load balancing for distributed memory multiprocessors. *Journal of Parallel and Distributed Computing*, 7(2):279–301, 1989.

[33] P. J. Davis. *Circulant Matrices*. John Wiley, 1979. A Wiley-Interscience Publication, Pure and Applied Mathematics.

[34] J. C. Delvenne, R. Carli, and S. Zampieri. Fast strategies in the consensus problem. In *NeCST Workshop*, Nancy, France, 2007.

[35] R. Elsässer, B. Monien, and S. Schamberger. Distributing unit size workload packages in heterogeneous networks. *J. Graph Algorithms Appl.*, 10(1):51–68, 2006.

[36] F. Fagnani and S. Zampieri. Randomized consensus algorithms over large scale networks. *IEEE Journal on Selected Areas in Communications*, 26(4):634–649, 2008.

[37] F. Fagnani and S. Zampieri. Average consensus with packet drop communication. *SIAM Journal on Control and Optimization*, 2009. to appear.

[38] P. Frasca, R. Carli, and F. Bullo. Multiagent coverage algorithms with gossip communication: control systems on the space of partitions. In *American Control Conference*, St. Louis, MO, June 2009. To appear.

[39] P. Frasca, R. Carli, and F. Bullo. Multiagent coverage algorithms with gossip communication: control systems on the space of partitions, Mar. 2009. Available at `http://arXiv.org/abs/0903.3642`.

[40] P. Frasca, R. Carli, F. Fagnani, and S. Zampieri. Average consensus by gossip algorithms with quantized communication. In *IEEE Conf. on Decision and Control*, pages 4831 – 4836, Cancún, México, Dec. 2008.

[41] P. Frasca, R. Carli, F. Fagnani, and S. Zampieri. Average consensus on networks with quantized communication. *International Journal of Robust and Nonlinear Control*, 2009. In press.

[42] P. Frasca, P. Mason, and B. Piccoli. Detection of gaussian signals via hexagonal sensor networks. In *IEEE Conf. on Decision and Control*, 2009. Submitted.

[43] A. Ganguli, J. Cortés, and F. Bullo. Distributed coverage of nonconvex environments. In V. Saligrama, editor, *Networked Sensing Information and Control*, Lecture Notes in Control and Information Sciences, pages 289–305. Springer, 2007.

[44] F. R. Gantmacher. *The theory of matrices*. Chelsea, New York, 1959.

[45] F. Garin and E. Lovisari. Personal communication, Aug. 2008.

[46] F. Garin and S. Zampieri. Performance of consensus algorithms in large-scale distributed estimation. In *European Control Conference*, Budapest, Hungary, Aug. 2009.

[47] R. Graham and J. Cortés. Asymptotic optimality of multicenter Voronoi configurations for random field estimation. *IEEE Transactions on Automatic Control*, 54(1):153–158, 2009.

[48] H. Groemer. On the symmetric difference metric for convex bodies. *Beiträge zur Algebra und Geometrie*, 41(1):107–114, 2000.

[49] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.

[50] J. M. Hendrickx. *Graphs and Networks for the Analysis of Autonomous Agent Systems*. PhD thesis, Université catholique de Louvain, 2008.

[51] J. Hespanha, D. Liberzon, D. Angeli, and E. D. Sontag. Nonlinear norm-observability notions and stability of switched systems. *IEEE Transactions on Automatic Control*, 50(2):154–168, 2005.

[52] I. I. Hussein and D. M. Stipanoviĉ. Effective coverage control for mobile sensor networks with guaranteed collision avoidance. *IEEE Transactions on Control Systems Technology*, 15(4):642–657, 2007.

[53] A. Jadbabaie, J. Lin, and A. S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, 2003.

[54] S. Kar and J. M. F. Moura. Distributed consensus algorithms in sensor networks: Quantized data. Dec. 2007. Submitted.

[55] A. Kashyap, T. Başar, and R. Srikant. Quantized consensus. *Automatica*, 43(7):1192–1203, 2007.

[56] J. Lavaei and R. M. Murray. On quantized consensus by means of gossip algorithm – Part I: Convergence proof. In *American Control Conference*, St. Louis, MO, July 2009.

[57] J. Lavaei and R. M. Murray. On quantized consensus by means of gossip algorithm – Part II: Convergence time. In *American Control Conference*, St. Louis, MO, July 2009.

[58] Q. Li and D. Rus. Global clock syncronization in sensor networks. *IEEE Transactions on Computers*, 55(2):214– 226, 2006.

[59] S. P. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, 1982.

[60] D. G. Luenberger. *Linear and Nonlinear Programming*. Addison-Wesley, Reading, MA, 2 edition, 1984.

[61] S. Martínez, F. Bullo, J. Cortés, and E. Frazzoli. On synchronous robotic networks – Part I: Models, tasks and complexity. *IEEE Transactions on Automatic Control*, 52(12):2199–2213, 2007.

[62] S. Martínez, F. Bullo, J. Cortés, and E. Frazzoli. On synchronous robotic networks – Part II: Time complexity of rendezvous and deployment algorithms. *IEEE Transactions on Automatic Control*, 52(12):2214–2226, 2007.

[63] S. Martínez, J. Cortés, and F. Bullo. Motion coordination with distributed information. *IEEE Control Systems Magazine*, 27(4):75–88, 2007.

[64] A. Montanari and A. Saberi. Convergence to equilibrum in local interaction games and Ising models. Online available.

[65] L. Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50(2):169–182, 2005.

[66] D. Mosk-Aoyama and D. Shah. Fast distributed algorithms for computing separable functions. *IEEE Transactions on Information Theory*, 54(7):2997–3007, 2008.

[67] S. B. Nadler. *Hyperspaces of sets*, volume 49 of *Monographs and Textbooks in Pure and Applied Mathematics*. Marcel Dekker, New York, 1978.

[68] G. Nair, F. Fagnani, S. Zampieri, and R. Evans. Feedback control under data rate constraints: An overview. *Proceedings of the IEEE*, 95(1):108–137, Jan. 2007.

[69] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis. On distributed averaging algorithms and quantization effects. Technical Report 2778, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Nov. 2007.

[70] A. Nedic and A. Ozdaglar. Distributed Subgradient Methods for Multi-Agent Optimization. *IEEE Transactions on Automatic Control*, 54(1):48–61, 2009.

[71] A. Nedic, A. Ozdaglar, and P. Parrilo. Constrained consensus and optimization in multi-agent networks. Technical Report 2779, Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Dec. 2008.

[72] J. R. Norris. *Markov Chains*. Cambridge University Press, Cambridge, UK, 1997.

[73] R. Olfati-Saber, J. A. Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.

[74] M. Pavone, E. Frazzoli, and F. Bullo. Distributed policies for equitable partitioning: Theory and applications. In *IEEE Conf. on Decision and Control*, pages 4191–4197, Cancún, México, Dec. 2008.

[75] M. Pease, R. Shostak, and L. Lamport. Reaching agreement in the presence of faults. *Journal of the Association for Computing Machinery*, 27(2):228–234, 1980.

[76] M. Penrose. *Random Geometric Graphs*. Oxford Studies in Probability. Oxford University Press, 2003.

[77] L. Pimenta, V. Kumar, R. Mesquita, and G. Pereira. Sensing and coverage for a network of heterogeneous robots. In *IEEE Conf. on Decision and Control*, pages 3947–3952, Cancún, México, Dec. 2008.

[78] L. C. A. Pimenta, V. Kumar, R. C. Mesquita, and G. A. S. Pereira. Sensing and coverage for a network of heterogeneous robots. In *IEEE Conf. on Decision and Control*, pages 3947–3952, Cancún, México, Dec. 2008.

[79] M. Rabbat and R. Nowak. Distributed optimization in sensor networks. In *Symposium on Information Processing of Sensor Networks (IPSN)*, pages 20–27, Berkeley, USA, Apr. 2004.

[80] R. Rajagopal and M. J. Wainwright. Network-based consensus averaging with general noisy channels. *CoRR*, abs/0805.0438, 2008. on-line available.

[81] W. Ren and R. W. Beard. *Distributed Consensus in Multi-Vehicle Cooperative Control. Theory and Applications*. Springer, New York, 2007.

[82] R. G. Sanfelice, R. Goebel, and A. R. Teel. Invariance principles for hybrid systems with connections to detectability and asymptotic stability. *IEEE Transactions on Automatic Control*, 52(12):2282–2297, 2007.

[83] S. Susca, F. Bullo, and S. Martínez. Synchronization of beads on a ring. In *IEEE Conf. on Decision and Control*, pages 4845–4850, New Orleans, LA, Dec. 2007.

[84] S. Tatikonda and S. Mitter. Control under communication constraints. *IEEE Transactions on Automatic Control*, 49(7):1056–1068, 2004.

[85] J. N. Tsitsiklis. *Problems in Decentralized Decision Making and Computation*. PhD thesis, Massachusetts Institute of Technology, Nov. 1984.

[86] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6-7):1226–1229, 1995.

[87] J. J. Xiao and Z. Q. Lou. Decentralized estimation in an inhomogeneous sensing enviromnent. *IEEE Transactions on Information Theory*, 51(10):3564–3575, 2005.

[88] L. Xiao, S. Boyd, and S.-J. Kim. Distributed average consensus with least-mean-square deviation. *Journal of Parallel and Distributed Computing*, 67(1):33–46, 2007.

[89] L. Xiao, S. Boyd, and S. Lall. A scheme for robust distributed sensor fusion based on average consensus. In *Symposium on Information Processing of Sensor Networks (IPSN)*, pages 63–70, Los Angeles, CA, Apr. 2005.

[90] L. Xiao, S. Boyd, and S. Lall. A space-time diffusion scheme for peer-to-peer least-square estimation. In *Symposium on Information Processing of Sensor Networks (IPSN)*, pages 168–176, Nashville, TN, Apr. 2006.

[91] P. Young. The evolution of conventions. *Econometrica*, 61:57–84, 1993.

[92] M. Zhong and C. G. Cassandras. Distributed coverage control in sensor network environments with polygonal obstacles. In *IFAC World Congress*, pages 4162–4167, Seoul, Korea, July 2008.

[93] M. Zhu and S. Martínez. On the convergence time of asynchronous distributed quantized averaging algorithms. *IEEE Transactions on Automatic Control*, 2008. Submitted.