

Dissertation ETH Zurich No. 23845

**Inference and control for populations of systems:
from aggregative games to systems biology**

A thesis submitted to attain the degree of
Doctor of Sciences of ETH Zurich
(Dr. sc. ETH Zurich)

presented by

Francesca Parise

Dott.sa Magistrale, Università degli studi di Padova, Italy
born 10.10.1988 in Verona
citizen of Italy

accepted on the recommendation of

Prof. Dr. John Lygeros, examiner
Prof. Dr. John N. Tsitsiklis, co-examiner
Prof. Dr. Maria Elena Valcher, co-examiner

2016

ETH Zurich
IfA - Automatic Control Laboratory
ETL, Physikstrasse 3
8092 Zurich, Switzerland

© Francesca Parise, 2016
All rights reserved

DOI: 10.3929/ethz-a-010859584
ISBN: 978-3-906327-71-6

Printed in Switzerland by the ETH Druckzentrum

Acknowledgements

I first and foremost want to express my deepest gratitude to my PhD supervisor Prof. John Lygeros. Without any doubt the results achieved in this thesis would not have been possible without his constant advice and technical suggestions. In these last four years I feel I have grown not only from a scientific point of view, but also in confidence and independence. This is largely due to his trust in assigning me tasks that I thought were out of my reach and hence giving me the confidence to push my limits and keep growing. Most of all, I want to thank him for showing me by example that is possible to be an excellent researcher on many different fields and for encouraging me to pursue any topic I found interesting. From him I have learned that a curious and open mind makes not only for a great advisor, but also for a great scientist. I also want to thank Prof. Maria Elena Valcher not only for her careful scrutiny of this thesis but also for her continuous support throughout the years. I can safely say that, if it wasn't for her, I would never have ended up where I am now. Her course on Systems Analysis is what made me love control theory in the first place, it was her suggestion that made me apply to ETH and her enthusiasm and energy towards teaching and research that convinced me that academia is the right path for me. I truly believe that she sets the highest possible standard and I can only hope to be one day as good a teacher and advisor as her is to me. I would also like to express my sincere gratitude to Prof. John N. Tsitsiklis for accepting to be my co-examiner. I am extremely thankful for the time he spent reading my work and for the many constructive comments he gave me.

Besides the aforementioned, I want to thank the many professors that provided me with feedback and suggestions during the last four years: Prof. Giorgio Picci for trusting me with my first and only project outside academia, but most of all for always believing in my capabilities and giving me all the support I could hope for; Prof. Mustafa Khammash for allowing me to run experiments in his lab and for providing many valuable suggestions during our work together; Prof. Manfred Morari for making IfA an outstanding research environment and for his crucial and invaluable support and advice on my next career steps; Prof. Florian Dorfler for having me as teaching assistant for his course and Prof. Roy Smith and Maryam Kamgarpour for their feedback on my work.

If there is one thing I have learned during my PhD is that in research $1 + 1$ is usually much greater than $2!$ In fact, none of my achievements would have been possible without the contributions from the postdocs and PhD students with whom I have worked together at IfA. First of all, I would like to thank Jakob Ruesch for adopting me and showing me the

way in my first steps in biology and research in general. Coming from a very theoretical background, he helped me understand the challenges underlying biological systems. I still remember when in my first attempt at systems identification I told him: “why don’t you just do black box, measure every 5 seconds and be done” ... the revelation that a single measure can take minutes was eye opening to me! Very special thanks go to Marcello Colombino and Basilio Gentile that, apart from spending countless hours doing research with me, helped me in so many different ways throughout these years! I specifically want to thank Marcello for being my shoulder in the first part of my PhD (when we were not sure we were going to make it), for inspiring me with his brilliant shower ideas, for saving me through many of my accidents and in general for being one of the most generous and selfless person I know. There are so many ways Basilio influenced and helped me in the ten years I have known him, but I think the way we met says it all. It was after one of our first classes as undergrads. That day our extremely tough professor made a mistake at the blackboard which I (very boldly) pointed out. So Basilio at the end of the class decided to come to me and ask for my autograph! That made me feel so good... and that is Basilio’s approach to people and to life. He is the most positive, supportive and honest person I know and I am so very lucky to have him as my friend. I also want to thank all my other collaborators: Sergio Grammatico for teaching me meticulous mathematical thinking and writing, Andreas Miliadis-Argeitis for teaching me the beauty and the insights that control theorists can get from performing experiments themselves and Dario Paccagnan for sharing with me an amazing journey in the world of variational inequalities!

There are many other people in the lab that I want to thank. First of all my crazy but amazing office mates: Xiaojing (George) Zhang and Tony Wood. You made my days super fun and even though I didn’t want any of you in the first place (I like my office quiet) I am incredibly happy that you joined! We always joke that there aren’t three people in the world that are more diverse than us. I believe that is true but that is what made it so fun! Thanks again for always being there for me, whether I was in a deadline-cannot-talk or let’s-go-sunbathing-in-the-terrace mood... and most importantly for always opening the door for me! Special thanks also to Tobias Sutter, Stephan Huck and Peyman Mohajerin for showing me around in my very first year and giving me so many good suggestions while I was trying to decide my research topic; Giampaolo Torrisi and the rest of the italian gang for sharing many dinners and fun times together; Saverio Bolognani for many nice discussions and tips; Eva Ahbe, Chithrupa Ramesh and Sabrina Baumann for helping me bring a little bit of Christmas to IfA, Yvonne Stürz for being my room companion at many conferences and Markus Thirer and Martine Wassmer for all the administrative support.

Finally, I am incredibly grateful to my family and friends. A special thank goes to my parents for always supporting my decisions (even if that meant for me moving away and not calling as often as I should), but most of all for making me believe that anything

is possible if you set your mind to it; to my brother for opening the way for me first in Switzerland and now in Cambridge; to the girls at home for our reunions and for always making me feel like I never left and, last but not the least, to my fiancée Andrea for his love, support and contagious optimism. My pride forbids me to say that I would have not been able to do it without him, but I can safely say that it would have been much more difficult and way less fun! I am looking forward to spending the rest of my life with you and face all our next challenges together!

Abstract

The objective of this thesis is the analysis and eventually the control of two different types of populations of systems: populations of rational agents and populations of biological systems. The main challenge for both types of populations lies in their large scale, which makes traditional single-system modeling and control methods computationally intractable and architecturally unfeasible. To overcome these difficulties, we exploit global control signals that affect the whole population and are designed based on aggregate information only.

More in detail, in the first part of the thesis we consider the problem of coordinating multi-user systems. The main challenge for this type of populations is that most of today's multi-user networks are large scale, with lack of access to local information and consist of users with diverse requirements. Consequently, traditional distributed optimization methods, that focus on a unique control objective and typically assume that the users are cooperative, cannot be applied. On the other hand, a new distributed control paradigm, where the infrastructure is kept simple and the agents make their decisions independently ("selfishly") is emerging. This new framework motivates the interaction of control with game theory: control actions must in fact include economic notions related to markets and incentives. In the first part of the thesis, we propose a framework to implement this new paradigm in a scalable way by exploiting the aggregative nature of many multi-user systems of practical interest, as demand-response markets or social and traffic networks, to name just a few. Specifically, we derive control schemes, either decentralized or distributed, that can be used to coordinate very large populations of agents to a Nash equilibrium, that is, to a stable configuration where no agent has incentive to unilaterally change its strategy.

In the second part of the thesis we focus on populations of biological systems. Specifically, we consider populations of cells whose internal state is determined by a series of stochastic reactions whose rate can be influenced by an external control signal, applied to the whole population. The most prominent examples are gene expression reaction networks, which are responsible for the fundamental operations of any living organism, and whose rate can be influenced by means of light or concentration signals. In contrast to the previous part of the thesis, biochemical reaction networks are natural systems (instead of technological ones) and therefore are usually unknown and affected by high levels of noise. Our main concern is the analysis and modeling of such systems. Recent technological advancements have enabled researchers to observe populations with thou-

sands of cells and have accurate recordings of the internal state of large samples over time. These samples can be used to derive accurate statistics not only of the average state of the cells in the population but also of higher order moments (as for example the variance). As first result we propose a systematic procedure to infer a model of the unknown population by exploiting the available external control signals and the population measurements just described. Secondly, we use the identified model to study the range of behaviors that such controlled stochastic biochemical reaction networks can exhibit under different choices of the available control input. These are in fact the first two fundamental steps needed to understand these systems and consequently be able to use them, either as part of synthetic circuits or integrated in live organisms, for example for drug synthesis or targeted medicine. Our theoretical results are validated with in-vivo experiments on a light-inducible gene expression system.

Sommario

Questa tesi si occupa di tematiche riguardanti l'analisi e il controllo di due differenti tipi di popolazioni di sistemi: popolazioni di agenti razionali e popolazioni di sistemi biologici. In entrambi i casi le popolazioni considerate sono composte da un numero molto elevato di sistemi, pertanto tradizionali tecniche di controllo, basate su modelli specifici per ogni singolo sistema, non sono né computazionalmente né tecnologicamente implementabili. Per superare queste difficoltà, nella tesi viene suggerito l'utilizzo di un unico segnale di controllo che influenza l'intera popolazione ed è basato solo su informazioni aggregate.

Più nel dettaglio, la prima parte della tesi considera il coordinamento di sistemi multi-agente. Tali popolazioni sono solitamente caratterizzate da un elevato numero di agenti con obiettivi differenti e da una limitata disponibilità di informazioni locali, che sono per la maggior parte private. Pertanto tradizionali metodi di ottimizzazione distribuita, nei quali viene considerato un solo obiettivo e gli agenti sono tipicamente cooperativi, non possono essere utilizzati per il controllo di tali popolazioni. In contrapposizione sta emergendo un nuovo paradigma di controllo distribuito, nato dall'interazione tra la teoria dei giochi e del controllo, secondo il quale l'azione di controllo deve essere esercitata tramite incentivi e prezzi a cui i singoli agenti reagiscono indipendentemente (egoisticamente). La prima parte della tesi suggerisce come implementare questo nuovo paradigma, per grandi popolazioni, sfruttando la natura aggregativa della maggior parte dei sistemi multi-agente di interesse. Alcuni esempi di applicazione sono gestione della domanda in mercati energetici, controllo di reti di interazione sociale o direzione del traffico stradale. Nello specifico, il contributo di questa tesi è lo sviluppo di schemi di controllo, decentralizzato o distribuito, che possono essere utilizzati per coordinare gli agenti verso un equilibrio di Nash, ovvero verso una configurazione stabile di strategie, nella quale nessun agente è incentivato a modificare il suo comportamento dato il comportamento del resto della popolazione.

La seconda parte della tesi considera popolazioni di cellule il cui stato è determinato da una serie di reazioni chimiche stocastiche. Nello specifico, particolare attenzione è riservata a reti di reazioni la cui frequenza dipende da un segnale di controllo esterno, comune a tutta la popolazione. Come principale esempio il lettore può considerare le reti di espressione genica: queste reti sono alla base del funzionamento di ogni cellula e possono essere influenzate tramite stimoli luminosi o chimici. Contrariamente alla prima parte della tesi, i sistemi di reazioni biochimiche sono sistemi naturali (invece che tecnologici) le cui dinamiche sono perlopiù sconosciute e solitamente affette da rumore. Il principale

obbiettivo per tali popolazioni è pertanto l'analisi e l'identificazione di modelli matematici, che possono essere usati per descrivere tali sistemi. Grazie a recenti avanzamenti tecnologici è possibile osservare popolazioni composte da migliaia di cellule e ottenere di conseguenza statistiche accurate non solo del comportamento medio della popolazione, ma anche di momenti di ordine superiore (come ad esempio la varianza). Il primo obiettivo di questa parte della tesi è proporre una procedura per l'identificazione e la calibrazione sistematica di modelli matematici di reti biochimiche, sfruttando i segnali di controllo a disposizione e le misure di popolazione sopra descritte. Il secondo obiettivo è quello di utilizzare i modelli identificati per analizzare i possibili comportamenti della popolazione in risposta a diversi stimoli esterni. Questi sono infatti i primi due passi necessari per la comprensione delle reti biochimiche controllate e quindi per il loro utilizzo, sia come parte di circuiti sintetici sia integrati in organismi viventi, per esempio per la sintesi di nuovi farmaci o nello sviluppo di cure personalizzate. I risultati teorici di questa tesi sono affiancati da risultati sperimentali relativi ad una rete di espressione genica sintetica implementata in cellule di lievito.

Contents

Acknowledgments	i
Abstract	v
Contents	ix
Notation	xiii
Acronymns	xiii
Symbols	xiv
1 Overview	1
1.1 Outline and contributions	2
1.1.1 Part I: populations of rational agents	2
1.1.2 Part II: populations of biological systems	3
1.2 Publications	3
1.2.1 Part I: populations of rational agents	3
1.2.2 Part II: populations of biological systems	5
1.2.3 Other publications	5
I Populations of rational agents: Aggregative games	7
2 Introduction	9
2.1 The aggregative game framework	11
2.1.1 Average aggregative games (AAG)	12
2.1.2 Network aggregative games (NAG)	13
2.2 Problem statement	13
2.2.1 Myopic versus boundedly rational agents	15
2.3 Related work	17
3 Mathematical Preliminaries	23
3.1 Regularity properties	23
3.2 Fixed points	26

3.2.1	Existence and uniqueness	27
3.2.2	Fixed point iterations	27
3.3	Variational inequalities	30
3.3.1	Existence and uniqueness	35
3.3.2	Iterative algorithms	35
4	Myopic agents	43
4.1	Assumptions on the agents and optimal responses	44
4.1.1	The quadratic case	46
4.2	Average aggregative games: schemes with central operator	47
4.2.1	Structural assumptions and communication requirements	48
4.2.2	The aggregation mapping	49
4.2.3	A class of decentralized coordination schemes	50
4.2.4	An application of AAGs: PEV coordination	53
4.3	Network aggregative games: distributed schemes	57
4.3.1	Structural assumptions and communication requirements	58
4.3.2	The extended aggregation mapping (ν fixed)	58
4.3.3	A class of distributed coordination schemes	59
4.3.4	An application of NAGs: multidimensional opinion dynamics	62
4.4	Average aggregative games: distributed schemes	65
4.4.1	Structural assumptions and communication requirements	65
4.4.2	The extended aggregation mapping ($\nu \rightarrow \infty$)	65
4.4.3	An application of AAGs: hierarchical demand-response	67
4.5	Appendix	69
4.5.1	Proofs of the results stated in Section 4.1	69
4.5.2	Proofs of the results stated in Section 4.2	71
4.5.3	Proofs of the results stated in Section 4.3	73
4.5.4	Proofs of the results stated in Section 4.4	77
5	Boundedly rational agents	79
5.1	Relation among Nash equilibria and variational inequalities	79
5.1.1	Decentralized and distributed algorithms	81
5.1.2	Average aggregative games	82
5.2	Generalized Nash equilibria	83
5.2.1	Primal-dual reformulation	85
5.3	Average aggregative games: schemes with central operator	87
5.3.1	A decentralized scheme: quadratic case	88
5.3.2	A decentralized scheme: general case	89

5.4	Appendix	91
5.4.1	Proofs of the results stated in Section 5.1	91
5.4.2	Proofs of the results stated in Section 5.2	92
5.4.3	Proofs of the results stated in Section 5.3	93
II Populations of biological systems: Controlled stochastic biochemical reaction networks		95
6	Introduction	97
6.1	The stochastic biochemical reaction network framework	99
6.1.1	Solving the CME: finite state projection	102
6.1.2	Solving the CME: moment equations	104
6.2	Problem statement: controlled biochemical networks and population measurements	106
7	A systematic parameter inference procedure from population data	109
7.1	Parameter inference	109
7.1.1	The likelihood of sample moments	110
7.1.2	Bayesian parameter inference	111
7.2	Posterior predictive distributions of sample moments	112
7.3	Experiment design for parameter inference	113
7.3.1	The Fisher information	113
7.3.2	Designing an optimal experiment	114
7.4	Sequential experiment design for parameter inference	115
8	Reachability analysis	119
8.1	Problem formulation	119
8.2	Reachability tools	121
8.2.1	The reachable set	121
8.2.2	The hyperplane method for linear systems with bounded input	123
8.2.3	The hyperplane method for switched affine systems	127
8.3	Reachable set of networks with affine propensities	130
8.3.1	Linear moment equations	131
8.3.2	Switched affine moment equations	133
8.4	Reachable set of networks with generic reactions	136
8.4.1	An equivalent formulation of the controlled CME	136
8.4.2	The FSP reduction technique for the controlled CME	137
8.4.3	Relation between projected reachable sets	138

8.5	Appendix	139
8.5.1	Proofs of the results stated in Section 8.2	139
8.5.2	Proofs of the results stated in Section 8.4	141
9	From theory to practice: a light-inducible gene expression circuit	143
9.1	The system	144
9.2	The CME for heterogeneous populations	146
9.3	Systematic parameter inference procedure	147
9.4	Comparison with random and experience-based experiments	151
9.5	Regulating gene expression	153
9.6	Reachability analysis	154
III	Conclusions	157
10	Conclusions and possible extensions	159
10.1	Part I: populations of rational agents	159
10.1.1	Direct extensions	159
10.1.2	Further research directions	160
10.2	Part II: populations of biological systems	163
10.2.1	Direct extensions	163
10.2.2	Further research directions	165
	Bibliography	167
	Curriculum vitae	181

Notation

Acronymns

AA Average aggregative. 48

AAG Average aggregative game. 12

AMON Anti-monotone. 23

BR Best response. 11

CME Chemical master equation. 101

COC Co-coercive. 36

CON Contraction. 24

CQ Constraint qualification. 33

FNE Firmly non-expansive. 24

FSP Finite state projection. 102

GNE Generalized Nash equilibrium. 84

KKT Karush-Kuhn-Tucker conditions. 34

MILP Mixed integer linear program. 127

MON Monotone. 23

NA Network aggregative. 57

NAG Network aggregative game. 13

NE Nash equilibrium. 12

NEX Non-expansive. 24

PEV Plug-in electric vehicle. 53

SMON Strongly monotone. 23

SPC Strongly pseudo contractive. 24

VI Variational inequality. 30

Symbols

$\mathbb{1}_N, e_i \in \mathbb{R}^n$	vector of all ones, canonical vector with 1 in position i
$\mathbb{Z}[a, b]$	$\{z \in \mathbb{Z} \mid a \leq z \leq b\}$
\mathcal{H}_S	Hilbert space with scalar product $\langle x, y \rangle_S := x^\top S y$, $S \succ 0$
$\Pi_{\mathcal{X}}^S(y)$	projection operator, $\Pi_{\mathcal{X}}^S(y) := \arg \min_{x \in \mathcal{X}} \ y - x\ _S$
$f(x) = \mathcal{O}(g(x))$	$\lim_{x \rightarrow \infty} \frac{f(x)}{g(x)} = 0$
$\nabla_x f(x) \in \mathbb{R}^{n \times m}$	gradient of $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$
$f \in \mathcal{C}^1$	the function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuously differentiable
$[A]_{ij} = A_{ij}$,	element in position (i, j) of the matrix A
A'_{ij}	element in position (i, j) of the matrix A' , for $\nu \in \mathbb{N}$
$\Lambda(A)$	spectrum of the matrix A
$\sigma_{\min}(A), \sigma_{\max}(A)$	minimum, maximum singular value of the matrix A
$\lambda_{\min}(A), \lambda_{\max}(A)$	minimum, maximum eigenvalue of the matrix A
$\ A\ _S$	matrix norm, $\ A\ _S := \sup_{x \neq 0} \frac{\ Ax\ _S}{\ x\ _S}$
$A^{\frac{1}{2}}$	principal square root of a symmetric positive definite matrix A
$A \succ 0$	A is positive definite, not necessarily symmetric (same for \succeq, \prec, \preceq)
$A \otimes B$	Kronecker product of the matrices A, B
$[x]_+ \geq 0, [x]_- \geq 0$	positive and negative part of the real number $x \in \mathbb{R}$
$x \perp y, x, y \in \mathbb{R}^n$	$x_i y_i = 0$ for all $i \in \mathbb{Z}[1, n]$, where x_i is the i^{th} component of x
$\partial \mathcal{X}$	boundary of the set $\mathcal{X} \subset \mathbb{R}^n$
$\text{conv}(\mathcal{X})$	convex hull of the set $\mathcal{X} \subset \mathbb{R}^n$
$\text{VI}(\mathcal{X}, f)$	variational inequality with set \mathcal{X} and operator f
$\text{SOL}(\mathcal{X}, f)$	set of solutions of the $\text{VI}(\mathcal{X}, f)$
$\text{NE}(\mathcal{G})$	set of Nash equilibria of the game \mathcal{G}
$\text{GNE}(\mathcal{G})$	set of generalized Nash equilibria of the game \mathcal{G}
$\mathbb{E}[X], \mathbb{V}[X]$	expectation and variance of the random variable X
$\mathbb{V}[X, Y]$	covariance of the random variables X, Y

Abbreviations

$$[x^1; \dots; x^N] := [x^1{}^\top, \dots, x^N{}^\top]^\top$$

x^{-i}	$:= [x^1; \dots; x^{i-1}; x^{i+1}; \dots; x^N]$
$x^{\mathcal{N}_i}$	$:= \{x^j \mid j \in \mathcal{N}_i\}$
A_s	$:= (A + A^\top)/2$
$\mathcal{X}_{1 \times N}$	$:= \mathcal{X}_1 \times \dots \times \mathcal{X}_N$
\mathcal{X}^{-i}	$:= \mathcal{X}_1 \times \dots \times \mathcal{X}_{i-1} \times \mathcal{X}_{i+1} \times \dots \times \mathcal{X}_N$
$\text{Conv}\{\mathcal{X}^i\}_{i=1}^N$	$:= \{x = \sum_{i=1}^N a^i x^i \mid \sum_{i=1}^N a^i = 1, a^i \geq 0, x^i \in \mathcal{X}^i \forall i \in \mathbb{Z}[1, N]\}$
$\frac{1}{N} \sum_{i=1}^N \mathcal{X}^i$	$:= \{x = \frac{1}{N} \sum_{i=1}^N x^i \mid x^i \in \mathcal{X}^i \forall i \in \mathbb{Z}[1, N]\}$
$A\mathcal{X}$	$:= \{y = Ax \mid x \in \mathcal{X}\}$
P_{ij}^ν	$:= [P^\nu]_{ij}$
\mathcal{P}_ν	$:= P^\nu \otimes I_n$
\mathcal{I}	$:= \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top \otimes I_n$

Reserved symbols

Part I

N	number of players
n	dimension of the strategy vector
$x^i \in \mathbb{R}^n$	strategy of agent i
$\tilde{J}^i(x^i, x^{-i})$	$:\mathbb{R}^{Nn} \rightarrow \mathbb{R}$, cost function of player $i \in \mathbb{Z}[1, N]$ in a generic game
$J^i(x^i, \sigma^i(x))$	$:\mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ cost function of player i in an AAG
$J^i(x^i, z^i)$	$:\mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ cost function of player i for an external signal z^i
P	adjacency matrix of the network
$\sigma^i(x)$	aggregator function
$\sigma_\nu^i(x)$	aggregator function for NAGs, $\sigma_\nu^i(x) := \sum_j P_{ij}^\nu x^j$
$\bar{\sigma}(x)$	aggregator function for AAGs, $\bar{\sigma}(x) := \frac{1}{N} \sum_j x^j$
$\mathcal{X}^i \subset \mathbb{R}^n$	local constraint set for the strategy of player $i \in \mathbb{Z}[1, N]$
$\tilde{\mathcal{X}} \subset \mathbb{R}^n$	superset of all the local sets \mathcal{X}^i for fixed size N
$\mathcal{X} \subset \mathbb{R}^n$	superset of all the local sets \mathcal{X}^i for any population size N
$\mathcal{C} \subset \mathbb{R}^{Nn}$	global constraint coupling the strategies of all players
$\mathcal{Q} \subset \mathbb{R}^{Nn}$	$:= \mathcal{X}_{1 \times N} \cap \mathcal{C}$
$\mathcal{Q}^i(x^{-i}) \subset \mathbb{R}^n$	$:= \{x^i \in \mathcal{X}^i \mid x \in \mathcal{C}\}$
$\Phi^{\text{P-B}}, \Phi^{\text{K}}, \Phi_{(k)}^{\text{M}}$	Picard-bannach (3.1), Krasnoselskij (3.2) and Mann (3.3) iteration
$x_{\text{br}}^i(x^{-i}) \in \mathbb{R}^n$	best response of player $i \in \mathbb{Z}[1, N]$, as defined in (2.1)
$x^{i\star}(z^i) \in \mathbb{R}^n$	optimal response of player $i \in \mathbb{Z}[1, N]$, as defined in (2.3)
$\mathbf{z} \in \mathbb{R}^{Nn}$	$:= [z^1; \dots; z^N]$
$\mathbf{x}(\mathbf{z}) \in \mathbb{R}^{Nn}$	vector of optimal responses as defined in (4.20)
$\mathcal{A} : \mathbb{R}^n \rightarrow \mathbb{R}^n$	aggregation mapping for AAGs, $\mathcal{A}(z) := \frac{1}{N} \sum_{i=1}^N x^{i\star}(z)$
$\mathcal{A}_\nu : \mathbb{R}^{Nn} \rightarrow \mathbb{R}^{Nn}$	extended aggregation mapping for NAGs, $\mathcal{A}_\nu(\mathbf{z}) := \mathcal{P}_\nu \mathbf{x}(\mathbf{z})$
$\mathcal{I} : \mathbb{R}^{Nn} \rightarrow \mathbb{R}^{Nn}$	extended aggregation mapping for AAGs, $\mathcal{I}(\mathbf{z}) := \mathcal{I} \mathbf{x}(\mathbf{z})$

Part II

S, R	number of species, reactions
Z_s, z_s	amount of species $s \in \mathbb{Z}[1, S]$, random variable and realization
Z, z	$:= [Z_1; \dots; Z_S], := [z_1; \dots; z_S]$
$\nu_r, \alpha_r(\theta_r, z), \theta_r$	stoichiometric vector, propensity and rate of reaction $r \in \mathbb{Z}[1, R]$
$p(t, z) = p(t, z \mid \theta, e)$	$:= \mathbb{P}[Z(t) = z]$ for experiment e with parameter θ
$P_i(t)$	$:= p(t, z^i)$ where z^i is the i^{th} possible realization of $Z(t)$
$\tilde{x}_\infty, x_\infty$	uncentered and centered moments of $\mathbb{P}[Z]$
$x_{\leq l}$	centered moments of $\mathbb{P}[Z]$ up to order l
x_i	mean of species $i \in \mathbb{Z}[1, S]$, $x_i := \mathbb{E}[Z_i]$
x_{ij}	covariance of species $i, j \in \mathbb{Z}[1, S]$, $x_{ij} := \mathbb{E}[(Z_i - x_i)(Z_j - x_j)]$
t_h	h^{th} measurement time, $h \in \mathbb{Z}[1, H]$
$\sigma_m \in \Sigma_m$	m^{th} external signal, $m \in \mathbb{Z}[1, M]$
\mathcal{E}	set of possible experiments $e \in \mathcal{E}$
$\hat{\mu}_s^e(t_h), \hat{\mu}_{s^2}^e(t_h)$	sample mean and variance for species s in exp. e at time t_h
$\hat{\mu}^e(t_h)$	$:= [\hat{\mu}_s^e(t_h), \hat{\mu}_{s^2}^e(t_h)]^\top$ vector of sample moments
$\mathcal{Y}^e, \mathcal{D}^e$	original dataset and dataset of sample moments for exp. e
$\hat{\theta}^{\text{ML}}(\mathcal{D}^e), \hat{\theta}^{\text{MAP}}(\mathcal{D}^e)$	maximum likelihood and maximum a posteriori estimators
$L(\mathcal{D}^e \mid \theta)$	likelihood of \mathcal{D}^e given the parameter θ
$I(\theta, e)$	Fisher information matrix for $e \in \mathcal{E}$ given the parameter θ
$p(\theta), p(\theta \mid \mathcal{D}^e)$	parameter prior and posterior distribution
$p^{\text{pred}}(\hat{\mu}_t^v \mid \mathcal{D}^e, v)$	posterior predictive distribution for $v \in \mathcal{E}$ given $p(\theta \mid \mathcal{D}^e)$
$p(\hat{\mu}^e(t_h) \mid \theta, e)$	probability distribution of the sample moments according to θ
$\mu_{t_h \theta}^e, \Sigma_{t_h \theta}^e$	mean and variance of $p(\hat{\mu}^e(t_h) \mid \theta, e)$
$\mathcal{R}(x_0)$	infinite-time reachable set from $x(0) = x_0$
$\mathcal{R}_T(x_0)$	reachable set from $x(0) = x_0$ at time T
$\mathcal{R}_T^C(x_0)$	output reachable set from $x(0) = x_0$ at time T
$\mathcal{R}_T^{\text{in}}(x_0), \mathcal{R}_T^{\text{out}}(x_0)$	inner and outer approximation of $\mathcal{R}_T(x_0)$
$\mathcal{R}_T^{C,\text{in}}(x_0), \mathcal{R}_T^{C,\text{out}}(x_0)$	inner and outer approximation of $\mathcal{R}_T^C(x_0)$

Overview

The last decade has witnessed a shift in research interest from the analysis of single systems, usually described in great detail, to the study of large populations of interconnected systems. This change of perspective is motivated by the emergence of a vast range of applications where complex behaviors can only be explained as the aggregate emergent behavior of a collection of many simple systems. Within this thesis we focus on two types of such populations.

In the first part of the thesis, we consider *populations of rational agents* satisfying two fundamental assumptions. Firstly, we suppose that the agents are egoistic in the sense that their objective is to select, among a set of feasible strategies, the one that maximizes their profit. Secondly, we assume that the agents have a collective influence on each other, so that the well being of an agent (and consequently its best strategy) depends on an aggregate of the strategies of the rest of the population (e.g., the average).

In the second part of the thesis, we consider *populations of biological systems*. In this case, we assume that the individual systems are cells whose internal state is governed by the same network of stochastic reactions. Each cell can therefore be seen as an independent realization of a common, usually unknown, stochastic process.

The main challenge in the analysis of both types of populations is their size, which could be of the order of thousands of systems making traditional single-system approaches unfeasible. Our objective is thus to provide a mathematical framework that can be used to concisely and tractably describe and eventually control these large populations of systems. The fundamental assumption that we make, to guarantee scalability of the proposed methods, is that we cannot design personalized control signals for each system in the population. Instead, we rely on global control signals, that are designed based on aggregate information only and affect the whole population. These global signals could be common incentives broadcast by a central operator, in the case of rational agents, or external signals applied to the environment, in the case of biological systems. In the case of population of rational agents we additionally investigate distributed control algorithms that require only local communications among the agents over a sparse network.

1.1 Outline and contributions

In the following we report a general overview of the material presented in the thesis. We refer to Chapter 2 and Chapter 6 for a detailed introduction, literature review and outline specific to populations of rational agents and biological systems, respectively.

1.1.1 Part I: populations of rational agents

To describe populations of rational agents satisfying the two assumptions mentioned above, we use the setting of *aggregative games* which we introduce in Chapter 2. Specifically, we focus on two classes of aggregative games that appear often in applications: *average aggregative games* in which the cost function of each agent depends on the average of the strategies of the whole population and *network aggregative games* in which the cost function of each agent depends on a convex combination of the strategies of its neighbors in a given network. Our main objective, for both types of games, is to propose scalable algorithms that the agents can use to update their strategies, in response to the strategies of the rest of the population, and that guarantee convergence to a Nash equilibrium (see Definition 2.1). We refer to this task as the coordination problem. In Chapter 3 we review the mathematical tools needed to derive our theoretical results. All the statements contained therein are known except for the technical result detailed in Theorem 3.3.12. The main contributions of the first part of the thesis are in Chapter 4 and 5:

1. In Chapter 4, we suggest different types of algorithms (both decentralized and distributed) that solve the coordination problem for populations of *myopic agents*, that is, agents that update their strategies by solving the profit-maximization problem at every step. To the best of our knowledge these are the only algorithms available in the literature to coordinate populations of myopic agents playing aggregative games, with multidimensional and heterogeneously constrained strategies.
2. In Chapter 5, we review different types of algorithms that have been proposed in the literature to solve the coordination problem for populations of *boundedly rational agents*, that is, agents that update their strategies by taking a gradient step of the profit-maximization problem instead of computing the maximizer. We then propose an extension of these algorithms that can be used to solve the coordination problem for populations of boundedly rational agents in the presence of coupling constraints (i.e. for generalized Nash problems as in Definition 5.2).

1.1.2 Part II: populations of biological systems

In the second part of the thesis we focus on populations of identical *stochastic biochemical reaction networks*. In this case, the population behavior can be described by using the so-called chemical master equation (CME). This is an infinite dimensional ordinary differential equation that describes the evolution of the probability that any cell in the population is in a given state over time. Solving the CME is in general very difficult. Consequently, we review in Chapter 6 two approaches that have been proposed in the literature to approximate it: the finite state projection method and the moments equations. The main contributions of the second part of the thesis are as follows:

1. In Chapter 8, we characterize the range of possible behaviors of stochastic biochemical reaction networks when the rate of (some of) the reactions can be controlled by an external signal. To this end, we propose to apply the tools of reachability analysis, typical of control theory, to the system of moments equations approximating the CME. From a theoretical perspective, we extend the hyperplane method, originally introduced in [GK91] for linear systems, to switched affine systems. We also extend the finite state projection algorithm proposed in [MK06] to the case of controlled reaction networks.
2. In Chapter 7, we review some results on parameter inference and experiment design and we combine them to propose an iterative characterization procedure to systematically identify biochemical reaction networks from population data. In Chapter 9 we demonstrate the utility of the proposed procedure by applying it to an in vivo light inducible circuit implemented in yeast. To the best of our knowledge this is the first study where parameter inference and experiment design have been systematically applied to characterize a stochastic biochemical network from population data.

1.2 Publications

The research of this thesis is interdisciplinary and was performed in close collaboration with other theorists and experimentalists. In the following, we provide a list of publications, underlying the results of this work. The symbol * denotes equal contribution.

1.2.1 Part I: populations of rational agents

The decentralized coordination algorithm for myopic agents proposed in Section 4.2 and its application to plug-in electric vehicles were developed in collaboration with M. Colombino, S. Grammatico and J. Lygeros and were published in:

- F. Parise, M. Colombino, S. Grammatico and J. Lygeros. “Mean field constrained charging policy for large populations of plug-in electric vehicles”, in *Proceedings of the IEEE Conference on Decision and Control (CDC)*, Los Angeles, CA, USA, Dec 2014, pp. 5101–5106. [PCGL14]
- F. Parise, S. Grammatico and J. Lygeros. “On constrained mean field control for large populations of heterogeneous agents: Decentralized convergence to Nash equilibria”, in *Proceedings of the European Control Conference (ECC)*, Linz, Austria, Jun 2015, pp. 3316–3321. [PGCL15]
- S. Grammatico*, F. Parise*, M. Colombino* and J. Lygeros. “Decentralized convergence to Nash equilibria in constrained deterministic mean field control”. *IEEE Transaction on Automatic Control*, accepted as full paper, to appear Nov. 2016. Arxiv: <http://arxiv.org/abs/1410.4421>. [GPCL16]

The distributed coordination algorithms for myopic agents proposed in Section 4.3 and 4.4 and their applications to opinion dynamics and hierarchical demand response were developed in collaboration with S. Grammatico, B. Gentile and J. Lygeros and were published in:

- F. Parise, S. Grammatico, B. Gentile and J. Lygeros. “Network aggregative games and distributed mean field control via consensus theory”. Arxiv: <http://arxiv.org/abs/1506.07719>. [PGGL15b]
- F. Parise, B. Gentile, S. Grammatico and J. Lygeros. “Network aggregative games: Distributed convergence to Nash equilibria”, in *Proceedings of the IEEE Conference on Decision and Control (CDC)*, Osaka, Japan, Dec 2015, pp. 2295–2300. [PGGL15a]

The decentralized coordination algorithm for boundedly rational agents with coupling constraints proposed in Section 5.3 was developed in collaboration with B. Gentile, D. Paccagnan, M. Kamgarpour and J. Lygeros and was published in:

- D. Paccagnan*, B. Gentile*, F. Parise*, M. Kamgarpour and J. Lygeros. “Distributed computation of generalized Nash equilibria in aggregative games with coupling constraints”. Accepted to *IEEE Conference on Decision and Control (CDC)*, Las Vegas, CA, USA, Dec 2016. [PPG+16]
- B. Gentile*, F. Parise*, D. Paccagnan*, M. Kamgarpour and J. Lygeros. “Nash and Wardrop equilibria in aggregative games with coupling constraints”. ArXiv: <https://arxiv.org/abs/1702.08789>. [GPP+17]

1.2.2 Part II: populations of biological systems

The methods presented in Chapter 8 were developed in collaboration with M.E. Valcher and J. Lygeros and were published in a series of papers:

- F. Parise, M.E. Valcher and J. Lygeros. “On the reachable set of the controlled gene expression system”, in *Proceedings of the IEEE Conference on Decision and Control (CDC)*, Los Angeles, CA, USA, Dec 2014, pp. 4597–4604. [PVL14]
- F. Parise, M.E. Valcher and J. Lygeros. “On the use of hyperplane methods to compute the reachable set of controlled stochastic biochemical reaction networks”, in *Proceedings of the IEEE Conference on Decision and Control (CDC)*, Osaka, Japan, Dec 2015, pp. 1259–1264. [PVL15]
- F. Parise, M.E. Valcher and J. Lygeros. “Reachability analysis for switched affine autonomous systems and its application to controlled stochastic biochemical reaction networks”. Accepted to *IEEE Conference on Decision and Control (CDC)*, Las Vegas, CA, USA, Dec 2016. [PVL16]

The characterization procedure proposed in Chapter 7 and the experimental results described in Chapter 9 were carried out in collaboration with J. Ruess, A. Miliadis-Argeitis, M. Khammash and J. Lygeros and were published in:

- J. Ruess*, F. Parise*, A. Miliadis-Argeitis, M. Khammash and J. Lygeros. “Iterative experiment design guides the characterization of a light-inducible gene expression circuit”. *Proceedings of the National Academy of Sciences of the USA (PNAS)*, vol. 112, no. 26, pp. 8148–8153, 2015. [RPMA+15]

1.2.3 Other publications

The following articles were published by the author during her doctoral studies but are not featured in this dissertation.

Further developments and applications of the presented methods

- S. Grammatico, B. Gentile, F. Parise and J. Lygeros. “A mean field control approach for demand side management of large populations of thermostatically controlled loads”, in *Proceedings of the European Control Conference (ECC)*, Linz, Austria, Jun 2015, pp. 3548–3553. [GGPL15]
- S. Grammatico, F. Parise and J. Lygeros. “Constrained linear quadratic deterministic mean field control: Decentralized convergence to Nash equilibria in large

populations of heterogeneous agents”, in *Proceedings of the IEEE Conference on Decision and Control (CDC)*, Osaka, Japan, Dec 2015, pp. 4412–4417. [GPL15]

- L. Möller, B. Gentile, F. Parise, S. Grammatico and J. Lygeros. “Constrained deterministic leader-follower mean field control”, in *Proceedings of the American Control Conference (ACC)*, Boston, MA, USA, Jul 2016. [MGP⁺16]

Identification of biological and ecological systems

- F. Parise, J. Lygeros and J. Ruess. “Bayesian inference for stochastic individual-based models of ecological systems: A pest control simulation study”. *Frontiers in Environmental Science*, vol. 3, no. 42, 2015. [PLR15]
- F. Parise, J. Ruess and J. Lygeros. “Grey-box techniques for the identification of a controlled gene expression model”, in *Proceedings of the European Control Conference (ECC)*, Strasbourg, France, Jun 2014, pp. 1498–1503. [PRL14]

Identification of a model for the tide in Venice

- F. Parise and G. Picci. “System identification for tide prediction in the Venetian lagoon”, in *Proceedings of the European Control Conference (ECC)*, Zurich, Switzerland, Jul 2013, pp. 2994–2999. [PP13]
- F. Parise and G. Picci. “Identification of high tide models in the Venetian lagoon: Variable selection and G-LASSO”, in *Proceedings of the IFAC World Congress*, Cape Town, South Africa, Aug 2014, pp. 10385–10390. [PP14]

Part I

Populations of rational agents: Aggregative games

Introduction

In the first part of the thesis we focus on populations of rational agents that are egoistic, profit maximizing and have a collective influence on each other. Some examples of applications are as follows.

- *Energy markets*: as a result of the introduction of renewable and distributed power generation, energy supply is becoming more and more unpredictable. To compensate for such uncertainty and to guarantee the required balance between energy demand and production, it has been suggested to regulate the energy consumption by shifting the demand of flexible loads (demand response) or by adjusting the price of the energy based on the instantaneous energy demand (dynamic pricing). Such real-time coupling of supply and demand requires the development of new control architectures that are able to coordinate the energy needs of large populations of consumers, by means of a distributed and hierarchical infrastructure [MCH13, CLLV14, BB14, KMC12].
- *Economics*: as exemplified by the recent financial crisis, the fit net of interlinkages that connects major and minor financial institutions makes the current economy highly complex and to a large extent unpredictable. Local interconnections among the agents allow shocks to propagate and amplify, potentially affecting the whole economic system. As stated in a number of investigations [AG00, FPR00, AOTS15], the study of distributed models that can describe, and eventually predict, these highly complex phenomena is of paramount importance for the safe regulation of the economic system. The use of game theoretical models to describe competitive behavior in economics has a long standing tradition, starting from the renowned Cournot model.
- *Opinion dynamics*: with the establishment of social media as a major communication channel, the study of opinion formation in social networks is becoming more and more relevant for interpreting crowd behavior and social trends. Since the seminal work of De Groot [DeG74], many researchers have derived models that can explain how phenomena as consensus, polarization or fragmentation can emerge when people are repeatedly updating their opinions by communicating one

with each other [HK02]. These models can be used to study the effect of rumors [AO11, For05], advertisement [MJB14] or innovation diffusion mechanisms [AOY11], and can have substantial repercussions on strategies for product marketing, political campaigning or online advertising pricing.

- *Traffic*: interactions among agents may arise from the use of a common resource or infrastructure. The classical example is a road transportation network in which the travel time of each agent, from its origin to destination, depends on the number of players that are on the road at the same time. These type of models can be mathematically described by congestion games [Ros73]. While the learning dynamics and equilibria properties of these systems have been characterized, the control problem of how to intervene on the system in order to reduce the price of anarchy (e.g. tragedy of the commons) is still a challenging question [Rou05, JT04].
- *Medicine*: It is well known that cells influence one with each other, both via short and long range interactions [WB05]. Modeling and possibly controlling these networks of interactions could have a huge effect on the way traditional medicine is understood. In [DPK⁺15], for example, it is shown how the quorum sensing mechanism, used by cells to sense the population density, can be employed to detect liver cancer at very early stages. In a much larger scale, models of human interaction networks can also be used to study the spread of diseases, thus opening the possibilities for the design of more efficient prevention and vaccination campaigns [NPP16].

Even though quite disparate, all the applications previously described have two main features in common:

- **Property 1:** they involve very *large* populations of agents that have a collective influence one on each other and are *egoistic* in the sense that each agent aims at achieving a desired objective, regardless of the interest of the other agents (e.g., minimize its energy bill, select the best opinion, maximize the profit or the cell growth, reduce infection or travel time, etc.).
- **Property 2:** the state of each agent depends on the *aggregate* effect of the state of the other players, rather than on one-to-one interactions. For example, in energy markets the price is determined by the overall energy consumption, in traffic the congestion level depends on the total number of cars and in economics the well-being of a firm depends on the overall state of the interacting firms.

The main objective of the first part of the thesis is to present a framework that can be used to describe and eventually control these types of systems. Specifically, we aim at exploiting the aggregative nature of these applications to derive coordination algorithms that are distributed and scalable, allowing for the analysis and control of very large populations.

2.1 The aggregative game framework

The standard setting to describe populations of agents that, as detailed in Property 1, influence each other and are profit-maximizing is the one of non-cooperative game theory. Among the vast class of non-cooperative games, according to Property 2, we focus on games played by a population of $N \in \mathbb{N}$ agents, that satisfy the following fundamental assumption.

Aggregative game: each player $i \in \mathbb{Z}[1, N]$ is not subject to one-to-one interactions, but is influenced only by an aggregate σ^i of the other players strategies (i.e., the aggregate strategies of a possibly agent-dependent subset of the population).

Mathematically, an *aggregative game* is a game \mathcal{G} in which each agent aims at selecting a strategy/action $x^i \in \mathbb{R}^n$ that belongs to the set of admissible strategies $\mathcal{X}^i \subset \mathbb{R}^n$ and minimizes a given cost function $J^i(\cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$, according to the following optimization problem

$$\mathcal{G} = \begin{cases} x_{\text{br}}^i(x^{-i}) := \arg \min_{x^i \in \mathbb{R}^n} J^i(x^i, \sigma^i(x)) \\ \text{s.t. } x^i \in \mathcal{X}^i \end{cases} \quad \forall i = 1, \dots, N, \quad (2.1)$$

where $x := [x^1; \dots; x^N] \in \mathbb{R}^{Nn}$ is the vector composed by the strategies of the whole population and $x^{-i} := [x^1; \dots; x^{i-1}; x^{i+1}; \dots; x^N] \in \mathbb{R}^{(N-1)n}$ is the vector of all the strategies except for x^i . The key feature of aggregative games is that the cost of each player is a function of 2 arguments, $J^i(s_1, s_2)$, where s_1 is always the strategy of player i , $s_1 = x^i$, and s_2 is a, possibly agent-specific, function of the strategies of the whole population $s_2 = \sigma^i(x)$, which could depend on x^i as well. In generic games, on the other hand, the cost function of player i is a function of N arguments, that we denote as $\tilde{J}^i(x^1, \dots, x^N)$ or for brevity $\tilde{J}^i(x^i, x^{-i})$, with slight abuse of notation¹. Note that a strategy is an n -dimensional vector with real entries, that is $x^i \in \mathbb{R}^n$. The best strategy that a player i can play, when the strategies of the other players are fixed to x^{-i} , is called in game theory the *best response* (BR) and is here denoted by $x_{\text{br}}^i(x^{-i})$. Note that in aggregative games, the best response depends on the strategies x^{-i} of the other players only via the aggregate quantity $\sigma^i(x)$. For example, in dynamic pricing, the electricity cost is proportional to the overall energy demand, so that $\sigma^i(x) = \sum_{j=1}^N x^j =: \sigma(x)$ for

¹The notation $\tilde{J}^i(x^i, x^{-i}) : \mathbb{R}^n \times \mathbb{R}^{(N-1)n} \mapsto \mathbb{R}$ should not be confused with $J^i(x^i, s_2) : \mathbb{R}^n \times \mathbb{R}^n \mapsto \mathbb{R}$, in fact $x^{-i} \in \mathbb{R}^{(N-1)n}$ while $s_2 \in \mathbb{R}^n$. In the thesis, the symbol \tilde{J} is used to denote the cost function of a generic game, while J is reserved for aggregative games, that is for games with the structure given in (2.1). We finally stress the fact that the cost function of player i is thus obtained as the composition of the function $J(s_1, s_2) : \mathbb{R}^n \times \mathbb{R}^n \mapsto \mathbb{R}$ with $s_1 = x^i \in \mathbb{R}^n$ and $s_2 = \sigma^i(x) : \mathbb{R}^{Nn} \rightarrow \mathbb{R}^n$.

all $i \in \mathbb{Z}[1, N]$. In the following, we refer to the mapping $\sigma^i(\cdot) : \mathbb{R}^{Nn} \rightarrow \mathbb{R}^n$ as the *aggregator function*.

A set of strategies in which each agent is playing the best response to the strategies of the other players is called a Nash equilibrium and has the fundamental property that no agent can benefit from unilaterally changing its strategy.

Definition 2.1 (Nash equilibrium (NE)). *Given N cost functions $\{J^i : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}\}_{i=1}^N$, N aggregator functions $\{\sigma^i : \mathbb{R}^{Nn} \rightarrow \mathbb{R}^n\}_{i=1}^N$ and $\varepsilon > 0$, a set of strategies $\{\bar{x}^i \in \mathcal{X}^i\}_{i=1}^N$ is an ε -Nash equilibrium for (2.1) if, for all $i \in \mathbb{Z}[1, N]$, it holds*

$$J^i(\bar{x}^i, \sigma^i(\bar{x})) \leq \min_{x^i \in \mathcal{X}^i} J^i(x^i, \sigma^i([\bar{x}^1; \dots; x^i; \dots; \bar{x}^N])) + \varepsilon.$$

If the above inequality holds for $\varepsilon = 0$ then $\{\bar{x}^i \in \mathcal{X}^i\}_{i=1}^N$ is a Nash equilibrium.

Under mild assumptions on the constraint sets, cost and aggregator functions, it is possible to show that the aggregative game (2.1) admits at least a Nash equilibrium.

Proposition 2.1.1 (Existence of a Nash equilibrium). *Assume that for each agent $i \in \mathbb{Z}[1, N]$ the constraint set \mathcal{X}^i is convex and compact and the cost function $x^i \mapsto J^i(x^i, \sigma^i(x))$ is convex and continuously differentiable in x^i for each x^{-i} . Assume that $f^i(x) := \nabla_{x^i} J^i(x^i, \sigma^i(x))$ is continuous in x . Then the aggregative game in (2.1) admits at least one Nash equilibrium.*

Proof. This is a sub case of the more general Proposition 5.1.2. We report here a proof for completeness. Let us start with an equivalent characterization of the Nash equilibria. A set of strategies $\{\bar{x}^i\}_{i=1}^N$ is a Nash equilibrium if and only if, for each $i \in \mathbb{Z}[1, N]$, $\bar{x}^i \in \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, \sigma^i([\bar{x}^1; \dots; x^i; \dots; \bar{x}^N]))$. Since for each $i \in \mathbb{Z}[1, N]$ and for each x^{-i} the function $x^i \rightarrow J^i(x^i, \sigma^i(x))$ is convex, by the minimum principle [SPFP10, Eq. (5) and (15)] we can equivalently state that $\{\bar{x}^i\}_{i=1}^N$ is a Nash equilibrium if and only if $\bar{x}^i = \Pi_{\mathcal{X}^i}(\bar{x}^i - f^i(\bar{x}))$ for each $i \in \mathbb{Z}[1, N]$ or, in compact form, if and only if $\bar{x} = \Pi_{\mathcal{X}^1 \times \dots \times \mathcal{X}^N}(\bar{x} - f(\bar{x}))$, where $\Pi_{\mathcal{Y}}(\cdot)$ denotes the projection operator onto \mathcal{Y} . The projection operator onto a convex and compact set is continuous. Therefore the function $x \rightarrow \Pi_{\mathcal{X}^1 \times \dots \times \mathcal{X}^N}(x - f(x))$ is a composition of continuous functions and it is itself continuous. Moreover, it takes value in the convex and compact set $\mathcal{X}^1 \times \dots \times \mathcal{X}^N$. Consequently, by [Sma74, Theorem 4.1.5(b)] (see also Section 3.2.1) it admits at least a fixed point \bar{x} which is then a Nash equilibrium. \square

In the rest of the thesis we focus on the two following types of aggregative games.

2.1.1 Average aggregative games (AAG)

In an *average aggregative game (AAG)* each agent has the same aggregator function, which coincides with the average (or the sum) of the strategies of the population. Mathematically, $\sigma^i(x) = \frac{1}{N} \sum_{i=1}^N x^i =: \bar{\sigma}(x)$. Examples of average aggregative games arise for

instance in demand-response applications, where the agents are coupled via the cost of electricity or in traffic applications, where the travel time depends on the total amount of cars driving on a given street.

2.1.2 Network aggregative games (NAG)

Network aggregative games (NAGs) are an extension of average aggregative games to games in which the agents interact through a network (Fig. 2.1). Contrary to AAGs, in NAGs each agent has a different aggregator function that is obtained as an agent specific convex combination of the players strategies, $\sigma^i(x) := \sum_{j=1}^N P_{ij}x^j$ where $\sum_{j=1}^N P_{ij} = 1$. The agent-specific coefficient $P_{ij} \in [0, 1]$ denotes the relevance of the strategy of agent j to agent i , $P_{ij} = 0$ implying no influence. The matrix $P \in \mathbb{R}^{N \times N}$, whose components are the weights P_{ij} , is known in graph theory as the weighted adjacency matrix of the network. The set of agents $\mathcal{N}^i = \{j \in \mathbb{Z}[1, N] \mid P_{ij} \neq 0\}$ is the set of neighbors of agent i . Without loss of generality we assume that the weights are normalized so that $\sum_{j=1}^N P_{ij} = 1$.

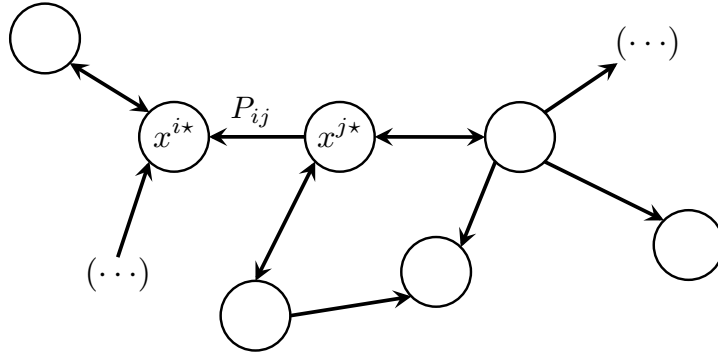


Figure 2.1: Scheme of a network aggregative game.

The framework of NAGs can be used to model for example opinion dynamics, economics or resource allocation problems.

2.2 Problem statement

One of the main challenges in aggregative games, and games in general, is to characterize and eventually control the evolution of the players strategies when the game is repeated iteratively, that is, when the agents are allowed to iteratively modify their strategies in response to the current population state. Algorithm 1 illustrates the general algorithmic structure that we consider in the thesis: at each iteration k the agents synchronously

update their belief $z_{(k)}^i$ on the aggregate state $\sigma_{(k)}^i$, by either communicating with their neighbors or with a central operator, and consequently update their strategy $x_{(k+1)}^i$.

Algorithm 1: Coordination Scheme in Aggregative Games

Initialization: Set $k = 0$. Each agent i has initial state $x_{(0)}^i$.

Iterate:

Communication

Each agent receives an aggregate information, $a_{(k)}^i$, on the state of the population either by communicating with the neighbors or with a central operator;

Estimation

Based on the received information each agent produces a local estimate

$$z_{(k)}^i = \Phi_{(k)} \left(a_{(k)}^i, z_{(k-1)}^i \right) \text{ of the aggregate state } \sigma^i(x_{(k)});$$

Strategy update

$$\text{Each agent updates its strategy: } x_{(k+1)}^i = \Gamma_{(k)} \left(x_{(k)}^i, z_{(k)}^i \right).$$

A fundamental property of Algorithm 1 is that the strategy updates are executed locally by each agent, based only on the most recent information received from the neighbors or from the central operator. This is a mandatory requirement for all the applications described above: for privacy reasons agents may not be willing to share their local constraints with a central coordinator, for security reasons distributed algorithms are more reliable than centralized solutions and finally, for computational reasons, they are preferable in applications involving very large populations.

The scheme summarized in Algorithm 1 can be interpreted in two different ways. The iterations k could be *repetitions* of the game over time². For example, in opinion dynamics the agents change their opinion over time by communicating more than once with their friends or, in traffic applications, agents learn their best route from home to work by repeating the coordination game day after day. On the other hand, in technological applications these iterations could be implemented before the game begins to *coordinate* the agents to a set of strategies of interest. The game is then played once and every agent plays only the strategy obtained at convergence of the algorithm. For example, in demand response, Algorithm 1 could be implemented remotely to coordinate the charging schedule of a fleet of electric vehicles for the following day.

²This concept should not be confused with the game theoretical concept of *repeated* games. Here, in fact, the objective of each agent is to find a Nash equilibrium of the single stage game given in (2.1). In other words, we assume that the agents do not take into account the impact of their current actions on the future actions of other players.

In both cases, our objective is to solve the following coordination problem.

Coordination problem: Given the game primitives and given the information structure, select the estimates and strategies update laws $(\Phi_{(k)}$ and $\Gamma_{(k)})$ such that the set of strategies $\{x_{(k)}^i\}_{i=1}^N$ converges to a Nash equilibrium.

2.2.1 Myopic versus boundedly rational agents

According to the set of admissible strategies update laws $\Gamma_{(k)}$ we distinguish two classes of coordination problems.

- **Case of myopic agents:** in many games the type of strategy update law used by the agents is fixed a priori (i.e. cannot be decided by the system operator). Consequently, the only degree of freedom that can be exploited to solve the coordination problem is the design of the estimate update law. To model such cases, we consider a scenario in which the agents are myopic and they select as next strategy, at every step of Algorithm 1, the one that minimizes their cost, given their current estimate of the state of the population. This type of agent behaviour is referred to as “best response” and has been hypothesized to be descriptive of the real agent behaviour in many studies³ [GM91, Mat92, HCW98]. Mathematically, the best response corresponds to the strategy update law

$$x_{(k+1)}^i := \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, \sigma^i(x_{(k)})). \quad (2.2)$$

Note that to evaluate (2.2) the agents need to know the real aggregate state $\sigma^i(x_{(k)})$. Since such requirement is unrealistic in large population games, we assume instead that the agents use the alternative strategy update law

$$x_{(k+1)}^i := \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, z_{(k)}^i), \quad (2.3)$$

³As mentioned in the previous section Algorithm 1 can be interpreted in two different ways. In the first case the iterations k are repetitions of the game over time: in this case the assumption of myopic agents is natural in the sense that agents have an incentive to best respond at every step, since their action at each step is actually implemented. In the second case, on the other hand, the iterations k are implemented before the game begins to coordinate the agents strategies. In this case there is less incentive for an agent to play the best response at each iteration, as the only action that is actually implemented is the one reached at convergence. Nonetheless, best response behaviour could be induced if the algorithm is stopped by the central operator after a random number of iterations. In other words, if the agents do not know when the iterations will stop (and hence when is that the strategy that they report will actually be implemented) then they have an incentive in correctly reporting their best response at every step.

which we refer to as the *optimal response* mapping (see also Section 4.1). Note that in (2.3) the agents evaluate their cost function using as second argument an estimate $z_{(k)}^i \in \mathbb{R}^n$ instead of the real aggregate state $\sigma^i(x_{(k)})$, as this is in general unknown. The design of good estimate update laws is the focus of Chapter 4. In the case with central operator, we are going to impose that such estimates are the same for every agent, so that $z_{(k)}^i = z_{(k)} \in \mathbb{R}^n$ for all $i \in \mathbb{Z}[1, N]$. Consequently, the signal $z_{(k)}$ can be seen as an incentive broadcast by the central operator and the proposed coordination scheme can be thought of as an incentive design scheme.

- **Case of boundedly rational agents:** For the purpose of solving the coordination problem, fixing the update law to the optimal response given in (2.3) is quite restrictive. Moreover, in applications where the agents have limited computation capabilities, solving an optimization problem at each algorithmic step may not be feasible. Consequently, in Chapter 5, we relax this constraint and consider coordination problems where the choice of the strategy update law is normative instead of descriptive (i.e., the strategy update law is not decided by the agents but by a central coordinator). We note that nonetheless, the global objective of the central operator is still only that of ensuring convergence to a Nash equilibrium. In other words, the central operator is a benevolent controller that does not seek its own interests nor some social interest (e.g., the central operator could be a regulator in an economic market)⁴. In the following, we consider gradient steps updates of the type

$$x_{(k+1)}^i := \Pi_{\mathcal{X}^i}(x_{(k)}^i - \alpha_k \nabla_{x^i} J^i(x_{(k)}^i, z_{(k)}^i)). \quad (2.4)$$

These update laws usually depend on a parameter α_k (i.e. the step length) which may need to be centrally coordinated and in general result in a suboptimal choice for the updated strategy $x_{(k+1)}^i$. For this reason we refer to agents using these update laws as boundedly rational. It is important to underline that, even though during the iterations of the algorithm the updates may be suboptimal, due to the imposed suboptimal update law, we devise schemes that guarantee convergence to a Nash equilibrium. Therefore, even though the agents do not have the authority of selecting their updates they are still guaranteed to reach at convergence the strategy that minimizes their cost function, given the strategies of the other players⁵. Consequently, for applications where only the strategy at convergence is implemented (as in demand response schemes), both the myopic and boundedly

⁴We refer to Section 10.1.2 for a discussion on the case when the central operator seeks a social optimum instead.

⁵Note that the setting of boundedly rational agents is different from the one of cooperative agents. While in both case the central operator can design the update laws of the agents, in the latter the objective of the central operator is to steer the agents to a configuration that minimizes a social cost function instead of steering the agents to a Nash equilibrium. This means that in the cooperative setup the performance at convergence could be arbitrarily unfavorable for a single agent. This is not the

rational schemes can be applied. For the case of bounded rationality, a number of different gradient schemes guaranteeing convergence to Nash equilibria in AAGs, either using a central operator or local communications, have been proposed in the literature [CLLV14, KNS12, PKL16, KNS16]. In Chapter 5 we focus on a generalization of these methods to aggregative games with global constraints coupling the strategies of the agents.

A scheme of the contributions of Part I of the thesis is given in the following table.

<i>Game structure</i>	<i>Communication structure</i>	<i>Myopic agents (Chapter 4)</i>	<i>Boundedly rational agents with coupling constraints (Chapter 5)</i>
AAG	central operator	Sec. 4.2	Sec. 5.3
NAG	distributed	Sec. 4.3	(Sec. 10.1.1)
AAG	distributed	Sec. 4.4	(Sec. 10.1.1)

Table 2.1: Scheme of the contributions in Part I. An extension of the decentralized algorithm proposed in Section 5.3, for AAGs with coupling constraints, to distributed NAGs and AAGs is suggested in the concluding Section 10.1.1.

2.3 Related work

Before delving in the technical details of Chapter 4 and 5, we discuss in this section some connections of the class of aggregative games considered in this thesis with other games typically found in the literature. To this end, we start by summarizing the most important features of the framework considered in this thesis.

1. *large populations*: we are interested in algorithms that are computationally efficient for very large populations;
2. *aggregative nature*: the cost function of each agents depends on the strategy of the other players via the aggregator function, which could be agent-dependent;
3. *multidimensional and constrained strategies*: the strategy of each player is a multidimensional vector that must satisfy some local constraints;
4. *distributed solutions*: we are interested in algorithms that can be implemented distributedly, so that local constraints and cost functions remain private.

case here, where the agents are always guaranteed to converge to a strategy that minimizes their cost function, given the strategies of the other players. Part of the reasons motivating our work is that such type of guarantees might make the agents more willing to participate (i.e. allow the central operator to design their strategy update law) in coordination schemes with guaranteed convergence to a Nash equilibrium instead of a social optimum.

Aggregative games

Aggregative games were originally proposed in the economic literature [Jen10, Kuk04] to model games in which each agent has the same aggregator function. In this community, particular attention has been devoted to the study of the simultaneous best response dynamics, which are the learning dynamics of a population of myopic agents that select at each iteration the best strategy (i.e. the BR) given the current population state. Mathematically, $x_{(k+1)}^i = \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, \sigma([x_{(k)}^1; \dots; x^i; \dots; x_{(k)}^N]))$. Results have been derived mainly for the case of scalar strategies or scalar valued aggregator functions $\sigma : \mathbb{R}^{Nn} \rightarrow \mathbb{R}$, under suitable monotonicity assumptions [Jen10, CH12].

Average aggregative games, on the other hand, have recently attracted attention in the control community, mainly motivated by demand-response coordination problems. Limited to the case of boundedly rational agents, distributed algorithms have been suggested in [CLLV14, KNS12] and more recently in [PKL16, KNS16].

Network games

The main difference between the standard literature on *network games* [Gal10] and NAGs is that, in the latter, the interaction among the players is not one-to-one but has an aggregative structure. In this respect, the models that are most related to the work proposed here are described in [JZ14]. Therein, results are derived for a scalar version of the model presented in Section 4.3. Differently from the standard literature on network games, we focus on games where the strategies are multidimensional (so that in most cases it is not possible to apply the ordering arguments in [JZ14]) and constrained (specific types of constraint sets are considered in [BK15]).

Population games

Population games [San10] describe games played among a finite number of populations, each containing an infinite number of homogeneous players, with discrete action sets. Here, on the other hand, we focus on a single population with a large but finite number of players that have continuous strategies sets.

Mean field games

AAGs have some features in common with the class of *mean field games*. The latter are stochastic dynamic games where the cost of each agent depends either on the average [HCM07] or on the whole distribution of the strategies in the population [LL07, BFY13]. Differently from our setup, in mean field games it is typically assumed that the agents are homogeneous [LL07] or that they have access to the parameter distribution of the population [HCM07, BP13] (in [KC10] the authors suggest a method to estimate this dis-

tribution from neighboring data). Moreover, the strategy set is typically unconstrained. Results are derived in the limit of infinite number of players by exploiting the so-called mean field approximation. We note that results on mean field games cannot be applied in our setting as we consider deterministic games with heterogeneous players that do not have access to the parameter distribution of the population and whose strategies are constrained by private convex sets.

Monotone games

Under some technical assumptions, certain classes of aggregative games can be seen as a specific case of *monotone games*, which are games whose Nash equilibria can be described as the solutions of a monotone variational inequality [FP03]. The literature on monotone games is vast, see e.g., [BT97, Pav07, FK07, SPFP10, ZF16]. Most of the algorithms proposed for monotone games however use variants of the gradient update law given in (2.4) and are therefore suited only for populations of boundedly rational agents.

Distributed optimization and potential games

Most of the algorithms used in this thesis, as well as in the previously cited references, are algorithms that are well known in the fields of fixed point iterations, distributed optimization or variational inequality problems [BT97, FP03, Ber07, BC10]. In this thesis, however, they are used to distributedly coordinate the agents to a Nash equilibrium (which can be seen as a “minimum” of N coupled optimization problems) instead of coordinating them to a minimum of a single cost function (as typically done in distributed optimization). This difference vanishes in the case of potential games [MS96].

Definition 2.2 (Potential game). *Consider a game \mathcal{G}_p composed by N agents with generic cost function $\tilde{J}^i(x^i, x^{-i}) \in \mathcal{C}^1$ and strategy $x^i \in \mathcal{X}^i \subset \mathbb{R}^n$. The game \mathcal{G}_p is termed potential if there exists a potential function $S(x) : \mathbb{R}^{Nn} \rightarrow \mathbb{R}$, $S(x) \in \mathcal{C}^1$, such that for all $i \in \mathbb{Z}[1, N]$, $x^{-i} \in \mathcal{X}^{-i}$ and $z^i, y^i \in \mathcal{X}^i$ it holds*

$$\tilde{J}^i(z^i, x^{-i}) - \tilde{J}^i(y^i, x^{-i}) = S([x^1; \dots; x^{i-1}; z^i; x^{i+1} \dots; x^N]) - S([x^1; \dots; x^{i-1}; y^i; x^{i+1} \dots; x^N])$$

or, equivalently, for all $i \in \mathbb{Z}[1, N]$ and $x^{-i} \in \mathcal{X}^{-i}$

$$\nabla_{x^i} \tilde{J}^i(x^i, x^{-i}) = \nabla_{x^i} S(x).$$

An immediate consequence of this definition is in fact that

$$x_{\text{br}}^i(x^{-i}) := \arg \min_{x^i \in \mathcal{X}^i} \tilde{J}^i(x^i, x^{-i}) = \arg \min_{x^i \in \mathcal{X}^i} S(x),$$

therefore, for each agent, minimizing the cost function is equivalent to minimizing the global potential function $S(x)$. Consequently, each global minimizer of $S(x)$ is a Nash

equilibrium [MAS09, Section II.A]. For these games, standard distributed optimization tools can be used to find a global minimizer of $S(x)$ and thus a Nash equilibrium. For example, one can show that the *sequential BR* dynamics, which are the dynamics obtained when every agent sequentially updates its strategy as its BR to the strategies of the other players, converge under mild assumptions.

Sequential best response dynamics

$$x_{(k+1)}^i = \arg \min_{x^i \in \mathcal{X}^i} \tilde{J}^i(x_{(k+1)}^1, \dots, x_{(k+1)}^{i-1}, x^i, x_{(k)}^{i+1}, \dots, x_{(k)}^N) \quad \forall i \in \mathbb{Z}[1, N] \quad (2.5)$$

Intuitively, the sequential BR dynamics converge, in the case of potential games, because each step leads to a decrease in the potential function $S(x)$ towards its global minimum, which is a Nash equilibrium. To see this, one can note that each step can be rewritten as $x_{(k+1)}^i = \arg \min_{x^i \in \mathcal{X}^i} S(x_{(k+1)}^1, \dots, x_{(k+1)}^{i-1}, x^i, x_{(k)}^{i+1}, \dots, x_{(k)}^N)$. More formally, one way to prove this statement is by noticing that, in the case of potential games, the sequential best response dynamics coincide with the nonlinear Gauss-Seidel algorithm to minimize $S(x)$ over $\mathcal{X}_{1 \times N} := \mathcal{X}^1 \times \dots \times \mathcal{X}^N$, see for example [BT97, Chapter 3, Eq. (3.13)]. The conclusion then follows from [BT97, Chapter 3, Proposition 3.9].

Proposition 2.3.1 (Convergence of the sequential best response dynamics). *Suppose that $S(x) : \mathbb{R}^{Nn} \rightarrow \mathbb{R}$ is \mathcal{C}^1 and convex and that the sets \mathcal{X}^i are non-empty, compact and convex for all $i \in \mathbb{Z}[1, N]$. Furthermore suppose that for each $i \in \mathbb{Z}[1, N]$, the cost function $\tilde{J}^i(x^i, x^{-i}) : \mathbb{R}^n \rightarrow \mathbb{R}$ is strictly convex in x^i for each $x^{-i} \in \mathcal{X}^{-i}$. For any initial condition $x_{(0)}$ let $x_{(k)}$ be the sequence generated by the sequential best response dynamics. Then every limit point of $x_{(k)}$ minimizes $S(x)$ over $\mathcal{X}_{1 \times N}$ and is therefore a Nash equilibrium of the potential game \mathcal{G}_p . If additionally $S(x)$ is strictly convex then $x_{(k)}$ converges to the unique minimizer of $S(x)$.*

We note that the sequential BR dynamics do not comply with the structure of Algorithm 1, since the agents are required to update their strategies one at the time in a sequential order. For large population games one would prefer a scheme where all the agents update their strategies *simultaneously* (which in the case of potential games is equivalent to the nonlinear Jacobi algorithm minimizing $S(x)$, see [BT97, Chapter 3, Eq. (3.12)]).

Simultaneous best response dynamics

$$x_{(k+1)}^i = \arg \min_{x^i \in \mathcal{X}^i} \tilde{J}^i(x_{(k)}^1, \dots, x_{(k)}^{i-1}, x^i, x_{(k)}^{i+1}, \dots, x_{(k)}^N) \quad \forall i \in \mathbb{Z}[1, N] \quad (2.6)$$

Unfortunately, the potential game structure is not sufficient to guarantee convergence of the simultaneous best response. The main reason is that, while each step of the sequential best response is guaranteed to decrease the value of the potential function $S(x)$, the same is not true for a step of the simultaneous best response dynamics. Convergence of the latter has been proven in [BT97, Chapter 3, Proposition 3.10], under the assumption that there exists $\gamma > 0$ such that the mapping $R(x) = x - \gamma \nabla S(x)$ is a block contraction (see Definition 5.1). This result can be generalized to games that are not potential under a similar block contraction condition (see [BT97, Section 3.5.6] and Section 5.1.1). The main contribution of the first part of the thesis is to propose alternative decentralized and distributed schemes that guarantee convergence to an almost Nash equilibrium, by using simultaneous updates, for generic aggregative games (i.e., not necessarily potential) and for mappings that are not necessarily block-contractions.

A detailed review on potential games is given in [San10, DHZ06]. We finally note that the important class of *congestion games* is a subclass of potential games [MS96].

Mathematical Preliminaries

In this chapter we present mathematical tools and results that are used in Chapter 4 and Chapter 5. Throughout the thesis we assume to work in a finite dimensional Hilbert space \mathcal{H}_S , defined by the scalar product $\langle x, y \rangle_S := x^\top S y$, with $S = S^\top \succ 0$. Most of the definitions and propositions can be extended to more general spaces, see for example [Ber07] or [Chi09]. All the sets are considered to be non-empty.

3.1 Regularity properties

In the next definition we summarize some regularity properties for single valued operators.

Definition 3.1 (Monotonicity properties [FP03, Definition 2.3.1]). *Consider the Hilbert space \mathcal{H}_S defined by the matrix $S = S^\top \in \mathbb{R}^{n \times n}$, $S \succ 0$. A mapping $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is*

1. *Strongly monotone (SMON) in \mathcal{H}_S if $\exists \epsilon > 0$ such that*

$$(f(x) - f(y))^\top S (x - y) \geq \epsilon \|x - y\|_S^2, \quad \forall x, y \in \mathbb{R}^n;$$

2. *Monotone (MON) in \mathcal{H}_S if*

$$(f(x) - f(y))^\top S (x - y) \geq 0, \quad \forall x, y \in \mathbb{R}^n;$$

3. *Anti-monotone (AMON) in \mathcal{H}_S if $-f(\cdot)$ is MON in \mathcal{H}_S . □*

Note that in the more general Banach spaces the terminology (strongly) accretive is used instead of (strongly) monotone (e.g. [Ber07, Definition 1.14]). Here we always use the latter since we work in Hilbert spaces. Intuitively, the concepts of MON and SMON operators generalize the concepts of non-decreasing and increasing functions from \mathbb{R} to \mathbb{R}^n .

Definition 3.2 (Regularity properties). *Consider the Hilbert space \mathcal{H}_S defined by the matrix $S = S^\top \in \mathbb{R}^{n \times n}$, $S \succ 0$. A mapping $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is*

1. Lipschitz [Ber07, Definition 1.6] in \mathcal{H}_S if $\exists L > 0$ such that

$$\|f(x) - f(y)\|_S \leq L \|x - y\|_S, \quad \forall x, y \in \mathbb{R}^n;$$

2. a Contraction (CON) [Ber07, Definition 1.6] in \mathcal{H}_S if $\exists \delta \in [0, 1)$ such that

$$\|f(x) - f(y)\|_S \leq \delta \|x - y\|_S, \quad \forall x, y \in \mathbb{R}^n;$$

3. Non-expansive (NEX) [BC10, Definition 4.1 (ii)] in \mathcal{H}_S if

$$\|f(x) - f(y)\|_S \leq \|x - y\|_S, \quad \forall x, y \in \mathbb{R}^n;$$

4. Firmly non-expansive (FNE) [BC10, Definition 4.1 (i)] in \mathcal{H}_S if

$$\|f(x) - f(y)\|_S^2 \leq \|x - y\|_S^2 - \|f(x) - f(y) - (x - y)\|_S^2, \quad \forall x, y \in \mathbb{R}^n;$$

5. Strongly pseudo contractive (SPC) [Ber07, Definition 1.13(a)] in \mathcal{H}_S if $Id - f$ is SMON in \mathcal{H}_S , that is,

$$(x - y - f(x) + f(y))^\top S(x - y) \geq \epsilon \|x - y\|_S^2, \quad \forall x, y \in \mathbb{R}^n.$$

□

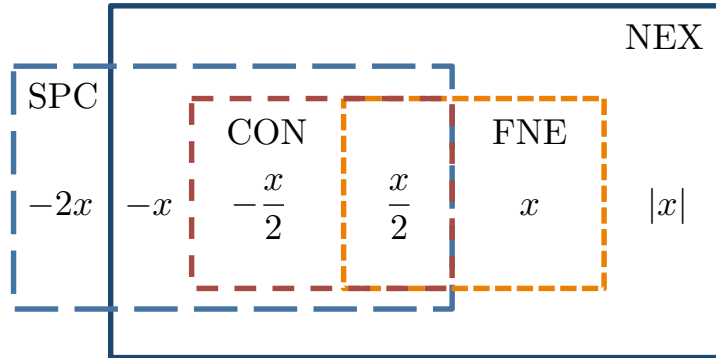


Figure 3.1: Relation among regularity properties and simple scalar examples. We note that FNE+SPC implies CON and that the concept of SPC mapping, as given in Definition 3.2, is different from the one of pseudo contraction defined in [BT97, Section 3.1]. For example $f(x) = -2x$ is SPC but is not a pseudo contraction and the pseudo contraction given in [BT97, Figure 3.1.1 (b)] is not SPC.

Note that the regularity properties of a mapping are strictly related to the considered Hilbert space. Therefore, the choice of a proper Hilbert space (i.e., norm) is critical.

Figure 3.1 illustrates the relations among the properties given in Definition 3.2. Contractiveness is a powerful property, frequently used in the game theoretical literature, however it usually leads to quite restrictive assumptions on the game primitives. This is the reason why, in the following, we consider different generalizations of this property. The simplest generalization is the class of **NEX** mappings. Among **NEX** mappings of particular interest is the class of **FNE** mappings, that includes for example the metric projection onto a closed convex set $\Pi_{\mathcal{X}} : \mathbb{R}^n \rightarrow \mathcal{X} \subseteq \mathbb{R}^n$ [BC10, Proposition 4.8]. Intuitively, **FNE** mappings are mappings that are **NEX** along certain directions and **CON** along others. A second different generalization of **CON** mappings is that of **SPC** mappings, which (contrary to **NEX** and hence **CON** and **FNE** mappings) are not necessarily continuous. **SPC** mappings are interesting because computing their fixed points is equivalent to computing the zeros of a **SMON** mapping (since $x = f(x)$ if and only if $x - f(x) = 0$ and $\mathcal{I} - f$ is **SMON** by definition) and many results have been derived in the literature for the latter problem.

In the technical proofs we make use of the following equivalent characterizations of **FNE** and **MON** mappings.

Lemma 3.1.1 ([BC10, Proposition 4.2]). *A mapping $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is **FNE** in \mathcal{H}_S if and only if $\|f(x) - f(y)\|_S^2 \leq (x - y)^\top S (f(x) - f(y))$, $\forall x, y \in \mathbb{R}^n$. \square*

Lemma 3.1.2 ([GPCL16, Lemma 1]). *If $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is **MON** and $g : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is **SMON** in \mathcal{H}_S , then $f + g$ is **SMON** in \mathcal{H}_S . \square*

Lemma 3.1.3. *Consider a mapping $f(y) : \mathbb{R}^{n_1+n_2} \rightarrow \mathbb{R}^m$ where $y = \begin{bmatrix} x \\ z \end{bmatrix}$, $x \in \mathbb{R}^{n_1}$, $z \in \mathbb{R}^{n_2}$. The mapping f is Lipschitz if and only if there exists $\bar{L} > 0$ such that $\|f(y_1) - f(y_2)\|_{I_m} \leq \bar{L}(\|x_1 - x_2\|_{I_{n_1}} + \|z_1 - z_2\|_{I_{n_2}})$ for all $y_1, y_2 \in \mathbb{R}^{n_1+n_2}$.*

Proof. Given the equivalence of norms in finite spaces, a mapping f is Lipschitz if and only if it is Lipschitz in the Euclidean norm. We therefore prove the statement for $S = I$, without loss of generality. If f is Lipschitz, then there exists $L > 0$ such that for all $y_1, y_2 \in \mathbb{R}^{n_1+n_2}$:

$$\|f(y_1) - f(y_2)\| \leq L\|y_1 - y_2\| = L\sqrt{\|x_1 - x_2\|^2 + \|z_1 - z_2\|^2} \leq L(\|x_1 - x_2\| + \|z_1 - z_2\|).$$

If there exists $\bar{L} > 0$ such that $\|f(y_1) - f(y_2)\| \leq \bar{L}(\|x_1 - x_2\| + \|z_1 - z_2\|)$ then

$$\|f(y_1) - f(y_2)\| \leq \bar{L}(\|y_1 - y_2\| + \|y_1 - y_2\|) \leq 2\bar{L}\|y_1 - y_2\|.$$

\square

Regularity of affine mappings

We next present necessary and sufficient conditions to characterize the regularity properties in Definitions 3.1 and 3.2 in the case of affine mappings.

Lemma 3.1.4 (Regularity of affine mappings). *Given an affine mapping $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ defined as $f(x) := Ax + b$, for some $A \in \mathbb{R}^{n \times n}$ and $b \in \mathbb{R}^n$, the following statements hold:*

$$\begin{aligned} f \text{ is } \text{NEX (CON)} \quad & \text{in } \mathcal{H}_S \iff A^\top SA - S \preccurlyeq (\prec) 0 \iff \|A\|_S \leq 1 \quad (\leq \delta) \\ f \text{ is } \text{FNE} \quad & \text{in } \mathcal{H}_S \iff 2A^\top SA \preccurlyeq A^\top S + SA \\ f \text{ is } (S)\text{MON} \quad & \text{in } \mathcal{H}_S \iff A^\top S + SA \succcurlyeq (\succ) 0 \end{aligned}$$

Proof. This is a classical result. The proof, which would be identical in the case of linear mappings, is reported for completeness. The mapping f is a **CON** in \mathcal{H}_S if and only if there exists $\delta \in [0, 1)$ such that $\|Ar - As\|_S \leq \delta\|r - s\|_S \forall r, s \Leftrightarrow \|A(r - s)\|_S \leq \delta\|r - s\|_S \forall r, s \Leftrightarrow \|Ax\|_S \leq \delta\|x\|_S \forall x \Leftrightarrow \|A\|_S \leq \delta \Leftrightarrow \|Ax\|_S^2 \leq \delta^2\|x\|_S^2 \forall x \Leftrightarrow x^\top A^\top SAx \leq \delta^2 x^\top Sx \forall x \Leftrightarrow x^\top (A^\top SA - S)x \leq (\delta^2 - 1)x^\top Sx \Leftrightarrow x^\top (A^\top SA - S)x \prec 0$. The equivalence for **NEX** can be proven in a similar fashion, by setting $\delta = 1$. The mapping f is **FNE** in \mathcal{H}_S if and only if $\|A(x - y)\|_S^2 \leq \|x - y\|_S^2 - \|A(x - y) - (x - y)\|_S^2$ for all $x, y \in \mathbb{R}^n$. Equivalently, we get $(x - y)^\top A^\top SA(x - y) \leq (x - y)^\top S(x - y) - (x - y)^\top (A - I)^\top S(A - I)(x - y)$ for all $x, y \in \mathbb{R}^n$, that is $A^\top SA \preccurlyeq S - (A - I)^\top S(A - I) = S - A^\top SA + A^\top S + SA - S \Leftrightarrow 2A^\top SA \preccurlyeq A^\top S + SA$. The mapping f is **SMON** in \mathcal{H}_S if and only if there exists $\epsilon > 0$ such that $(x - y)^\top A^\top S(x - y) \geq \epsilon\|x - y\|_S^2 = \epsilon(x - y)^\top S(x - y)$ for all $x, y \in \mathbb{R}^n$, that is equivalent to $\frac{1}{2}(A^\top S + SA) \succcurlyeq \epsilon S$. Since $S \succ 0$, the latter matrix inequality is equivalent to $(A^\top S + SA) \succ 0$. An analogous proof with $\epsilon = 0$ shows that the mapping f is **MON** in \mathcal{H}_S if and only if $A^\top S + SA \succcurlyeq 0$. \square

3.2 Fixed points

In Chapter 4 we illustrate how the coordination problem for large populations of myopic agents can be solved by finding the fixed point of a suitable mapping. In the following, we denote a generic fixed point by the symbol \bar{z} for consistency with the notation used in Chapter 4.

Definition 3.3 (Fixed point). *A point $\bar{z} \in \mathbb{R}^n$ is a fixed point for the mapping $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ if $\bar{z} = f(\bar{z})$.* \square

Fixed point problems are ubiquitous in mathematics and engineering and consequently conditions guaranteeing their existence, uniqueness and algorithms to compute them have been extensively investigated. In the interest of space we report here only the results that are needed in Chapter 4. We refer to [Ber07, Chi09, BC10, BT97, Sma74], for a more comprehensive survey of fixed point theory.

3.2.1 Existence and uniqueness

The main existence result that we use is a generalization of Brouwer fixed point theorem as reported in the next lemma.

Lemma 3.2.1 (Schauder [Sma74]). *Consider a continuous mapping $f : X \rightarrow Y$, where $X \subseteq \mathbb{R}^n$ is convex and $Y \subseteq X$ is compact. Then f admits a fixed point $\bar{z} = f(\bar{z})$. \square*

This result can be refined if the mapping $f(\cdot)$ possesses one of the regularity properties presented in Definition 3.2. Specifically, if one assume that the mapping is a CON then it is possible to prove existence and uniqueness without any assumption on the domain and codomain.

Lemma 3.2.2 (Banach theorem [Ber07, Theorem 2.1]). *Consider a mapping $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$. If f is a CON then it admits a unique fixed point $\bar{z} = f(\bar{z})$. \square*

Unlike CON mappings, NEX mappings may have more than one fixed point (e.g. the identity mapping $f(x) = x$ is NEX but not a CON and has infinitely many fixed points). The same is true also for the subclass of FNE mappings (e.g. the projection operator onto a convex and compact set is FNE [BC10, Proposition 4.8] but not a CON and has infinitely many fixed points). On the other hand, SPC mappings maintain the uniqueness property.

Lemma 3.2.3 (Chidume 1987). *Consider a mapping $f : X \rightarrow X$, where $X \subset \mathbb{R}^n$ is compact and convex. If f is SPC then it admits at most a unique fixed point $\bar{z} = f(\bar{z})$. If f is continuous and SPC then it admits a unique fixed point $\bar{z} = f(\bar{z})$.*

Proof. This is a classical result, reported for completeness. Suppose by contradiction that there are two fixed points $\bar{z}_1 = f(\bar{z}_1), \bar{z}_2 = f(\bar{z}_2)$. Then by definition of SPC, there exists $\epsilon > 0$ such that

$$\begin{aligned} \|\bar{z}_1 - \bar{z}_2\|_S^2 &\leq \frac{1}{\epsilon} (\bar{z}_1 - f(\bar{z}_1) - \bar{z}_2 + f(\bar{z}_2))^\top S (\bar{z}_1 - \bar{z}_2) \\ &= \frac{1}{\epsilon} (\bar{z}_1 - \bar{z}_1 - \bar{z}_2 + \bar{z}_2)^\top S (\bar{z}_1 - \bar{z}_2) = 0. \end{aligned}$$

Hence $\bar{z}_1 = \bar{z}_2$. Existence of a fixed point under the continuity assumption follows from Lemma 3.2.1. \square

3.2.2 Fixed point iterations

Finding the fixed point of a mapping f is in general a very difficult task. However, if the mapping f has one of the regularity properties given in Definition 3.2, one can use known fixed point algorithms that aim at iteratively constructing a sequence $\{z^{(k)}\}_{k=1}^\infty$ such that $z^{(k)} \rightarrow \bar{z} = f(\bar{z})$.

Picard-Banach iteration

The simplest fixed point iteration is the Picard–Banach iteration

$$z_{(k+1)} = f(z_{(k)}) =: \Phi^{\text{P-B}}(z_{(k)}, f(z_{(k)})), \quad (3.1)$$

which is obtained by simply applying the operator $f(\cdot)$ iteratively to an initial point $z_{(0)}$, so that

$$z_{(1)} = f(z_{(0)}), \quad z_{(2)} = f(f(z_{(0)})), \quad \dots, \quad z_{(k)} = f^k(z_{(0)}),$$

where f^k denotes the k^{th} functional power of the mapping $f(\cdot)$, which is recursively defined as $f^k := f \circ f^{k-1}$, $f^1 = f$. Even though quite simple, the Picard–Banach iteration is at the core of many of the algorithms suggested in game theory to coordinate the agents strategies to a Nash equilibrium [MCH13, HCM07]. Not surprisingly, the sufficient conditions needed to guarantee its convergence are however quite restrictive. We summarize them in the next two lemmas.

Lemma 3.2.4 ([Ber07, Theorem 2.1]). *If the mapping $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a CON then the Picard–Banach iteration (3.1) converges from any initial point $z_{(0)} \in \mathbb{R}^n$ to the unique fixed point $\bar{z} = f(\bar{z})$ with geometric rate, that is $\|z_{(k)} - \bar{z}\| \leq \delta^k \|z_{(0)} - \bar{z}\|$. \square*

Lemma 3.2.5 ([GR84, Theorem 15.1]). *If the mapping $f : X \rightarrow X$, with $X \subseteq \mathbb{R}^n$ closed and convex, is FNE then the Picard–Banach iteration (3.1) converges from any initial point $z_{(0)} \in X$ to a fixed point of f , if such a point exists. \square*

One of the main objective of Chapter 4 is to suggest the use of alternative, slightly more complex, fixed point iterations, whose convergence can be guaranteed under weaker conditions. We report them in the following.

Krasnoselskij iteration

According to the previous lemma the Picard-Banach iteration converges if f is a CON or if it is FNE. Those results cannot be generalized to the case of NEX mappings; for example, $z \mapsto f(z) := -z$ is NEX, but not a CON, and the Picard–Banach iteration $z_{(k+1)} = f(z_{(k)}) = -z_{(k)}$ oscillates indefinitely between $z_{(0)}$ and $-z_{(0)}$. In this case one can use the Krasnoselskij iteration

$$z_{(k+1)} = (1 - \lambda)z_{(k)} + \lambda f(z_{(k)}) =: \Phi^{\text{K}}(z_{(k)}, f(z_{(k)})) \quad (3.2)$$

that selects as next vector $z_{(k+1)}$ the convex combination of the previous vector $z_{(k)}$ and $f(z_{(k)})$, according to a fixed parameter $\lambda \in (0, 1)$, thus introducing a memory.

Lemma 3.2.6 ([Ber07, Theorem 3.2]). *If a mapping $f : X \rightarrow X$ is NEX, with $X \subset \mathbb{R}^n$ compact and convex, then the Krasnoselskij iteration (3.2) where $\lambda \in (0, 1)$, converges, for any initial condition $z_{(0)} \in X$, to a fixed point of f . \square*

Note that **CON** and **FNE** mappings are special cases of **NEX** mappings, hence the previous lemma guarantees convergence also for these cases. Finally, the Krasnoselskij iteration is also guaranteed to converge in the case of **SPC** mappings, if these are Lipschitz and the parameter λ is sufficiently small.

Lemma 3.2.7 ([Ber07, Theorem 3.6]). *If a mapping $f : X \rightarrow X$ with $X \subset \mathbb{R}^n$ compact and convex is Lipschitz in \mathcal{H}_S with constant L and **SPC** in the same space \mathcal{H}_S with constant ϵ , then the Krasnoselskij iteration (3.2) where $\lambda \in \Lambda := (0, a)$ and*

$$a = \frac{2\epsilon}{L^2 - 1 + 2\epsilon}$$

converges, for any initial condition $z_{(0)} \in X$, to the unique fixed point $\bar{z} = f(\bar{z})$. Moreover, $\|z_{(k)} - \bar{z}\| \leq q^k \|z_{(0)} - \bar{z}\|$, with

$$q = \sqrt{(1 - \lambda)^2 + (\lambda L)^2 + 2\lambda(1 - \lambda)(1 - \epsilon)}.$$

Proof. This proof follows the lines of [Ber07, Theorem 3.6] and is here reported, for convenience, using the notation of this thesis. Consider the auxiliary mapping $g(x) = (1 - \lambda)x + \lambda f(x)$. We are going to show that, under the given assumptions, g is a **CON** in X with rate $\delta = q \in [0, 1)$. The conclusion then follows from Lemma 3.2.4.

$$\begin{aligned} \|g(x) - g(y)\|_S^2 &= (1 - \lambda)^2 \|x - y\|_S^2 + \lambda^2 \|f(x) - f(y)\|_S^2 + 2\lambda(1 - \lambda)(f(x) - f(y))^T S(x - y) \\ &\leq (1 - \lambda)^2 \|x - y\|_S^2 + \lambda^2 L^2 \|x - y\|_S^2 + 2\lambda(1 - \lambda)(1 - \epsilon) \|x - y\|_S^2 \\ &\leq [(1 - \lambda)^2 + \lambda^2 L^2 + 2\lambda(1 - \lambda)(1 - \epsilon)] \|x - y\|_S^2 = q^2 \|x - y\|_S^2. \end{aligned}$$

Note that the quantity q is well defined since $(1 - \lambda)^2 + \lambda^2 L^2 + 2\lambda(1 - \lambda)(1 - \epsilon)$ is the sum of three non-negative terms (one can assume $\epsilon < 1$ without loss of generality). To guarantee that $g(x)$ is a **CON** we need to impose $q < 1$. However, $q^2 = \lambda^2(L^2 - 1 + 2\epsilon) - 2\lambda\epsilon + 1 < 1 \Leftrightarrow \lambda(L^2 - 1 + 2\epsilon) - 2\epsilon < 0 \Leftrightarrow \lambda < a$. \square

Note that in order to select a value of λ that guarantees convergence one needs to know the Lipschitz constant L of f . If this is not known one can resort to a iteration varying parameter, as discussed next.

Mann iteration

The most general fixed point iteration that we consider is the Mann iteration

$$z_{(k+1)} = (1 - \alpha_k)z_{(k)} + \alpha_k f(z_{(k)}) =: \Phi_k^M(z_{(k)}, f(z_{(k)})), \quad (3.3)$$

which is a Krasnoselskij iteration where the fixed parameter λ is substituted by a sequence of parameters $(\alpha_k)_{k=0}^\infty$ such that $\alpha_k \in (0, 1) \forall k \geq 0$, $\lim_{k \rightarrow \infty} \alpha_k = 0$ and $\sum_{k=0}^\infty \alpha_k = \infty$.

Lemma 3.2.8 ([Ber07, Theorem 4.11]). *If a mapping $f : X \rightarrow X$ is Lipschitz and SPC with $X \subset \mathbb{R}^n$ compact and convex, then the Mann iteration in (3.3) converges, for any initial condition $z_{(0)} \in X$, to the unique fixed point of f . \square*

Lemma 3.2.9 ([Ber07, Fact 4.9, p.112]). *If a mapping $f : X \rightarrow X$ is NEX with $X \subset \mathbb{R}^n$ compact and convex, then the Mann iteration in (3.3) converges, for any initial condition $z_{(0)} \in X$, to a fixed point of f . \square*

To summarize, if we consider a Lipschitz mapping $f : X \rightarrow X$ with X compact and convex the following pairs of regularity conditions and fixed point iterations guarantee convergence to a fixed point $\bar{z} = f(\bar{z})$.

		Picard–Banach	Krasnoselskij	Mann
1.	CON	✓	✓	✓
2.	NEX		✓	✓
3.	FNE	✓	✓	✓
4.	SPC		✓*	✓

* with $\lambda \in \Lambda$ as defined in Lemma 3.2.7.

Table 3.1: Conditions on regularity properties and iterations that ensure convergence to a fixed point.

3.3 Variational inequalities

In Chapter 5 we illustrate how Nash equilibria can be equivalently characterized in terms of solutions to a suitable variational inequality. We then exploit this knowledge to suggest decentralized and distributed algorithms to solve the coordination problem in the case of boundedly rational agents. To this end, we use some known results on existence, uniqueness and convergence to solutions of variational inequalities, that we briefly recap in this section. We refer the reader to [FP03, BT97] for a comprehensive review. We note that in this section, as well as in Chapter 5, we assume to work in the finite dimensional Hilbert space \mathcal{H}_I , instead of a generic space \mathcal{H}_S .

Definition 3.4 (Variational inequality (VI)). *Consider a set $\mathcal{X} \subseteq \mathbb{R}^n$ and an operator $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$. The variational inequality problem $VI(\mathcal{X}, f)$ is the problem of finding a vector $\bar{x} \in \mathcal{X}$ that solves*

$$f(\bar{x})^\top (x - \bar{x}) \geq 0 \quad \forall x \in \mathcal{X}.$$

The set of vectors that solve the $VI(\mathcal{X}, f)$ is denoted by $SOL(\mathcal{X}, f)$. \square

Connection with convex optimization

Variational inequalities are a generalization of convex optimization problems. In fact, let

$$\begin{aligned} \bar{x} \in \arg \min_{x \in \mathbb{R}^n} J(x) \\ \text{s.t. } x \in \mathcal{X}. \end{aligned} \quad (3.4)$$

By using the minimum principle [FP03, Section 1.3.1], one can show the following statement.

Proposition 3.3.1. (*Minimum principle*) Consider the optimization problem in (3.4) and assume that $J : \mathbb{R}^n \rightarrow \mathbb{R}$ is \mathcal{C}^1 on an open superset of the closed and convex set $\mathcal{X} \subseteq \mathbb{R}^n$. Then any local minimizer \bar{x} of (3.4) must satisfy

$$\nabla_x J(\bar{x})^\top (x - \bar{x}) \geq 0 \quad \forall x \in \mathcal{X}. \quad (3.5)$$

If, additionally, the function $J(\cdot)$ is convex then \bar{x} is a global minimum of (3.4) if and only if it satisfies (3.5). \square

Intuitively, this result says that a point $\bar{x} \in \mathcal{X}$ is a minimum of J if by taking any feasible step (i.e., a step that leads inside the set \mathcal{X}) the objective function can only increase (because the gradient of J has a positive component in that direction). See Figure 3.2.

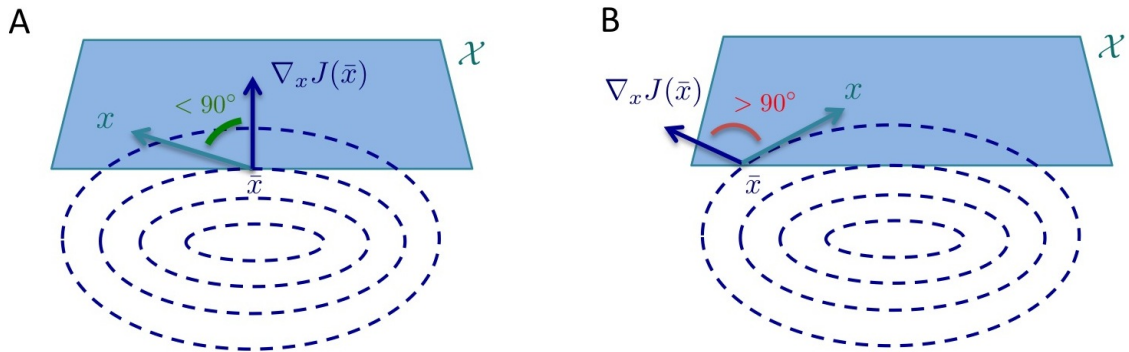


Figure 3.2: Illustration of the minimum principle, the set \mathcal{X} is in blue and the dotted lines are the contour plot of $J(x)$. A) The point \bar{x} is a minimizer of J in \mathcal{X} because $\nabla_x J(\bar{x})^\top (x - \bar{x}) \geq 0$ for all $x \in \mathcal{X}$ or equivalently the angle between $x - \bar{x}$ and $\nabla_x J(\bar{x})$ is less than 90° . B) The point \bar{x} is not a minimizer of J in \mathcal{X} because there exists $x \in \mathcal{X}$ such that the angle between $x - \bar{x}$ and $\nabla_x J(\bar{x})$ is more than 90° .

The following conclusion is then immediate.

Corollary 3.3.2 (Equivalence of VI and convex optimization). *Consider the optimization problem in (3.4), assume that $J : \mathbb{R}^n \rightarrow \mathbb{R}$ is \mathcal{C}^1 and convex and that \mathcal{X} is closed and convex, then*

$$\bar{x} \in \arg \min_{x \in \mathcal{X}} J(x) \quad \Leftrightarrow \quad \bar{x} \in \text{SOL}(\mathcal{X}, \nabla_x J).$$

□

In other terms, the previous corollary says that whenever the operator f of the VI(\mathcal{X}, f) is the gradient of a convex function J , then the set $\text{SOL}(\mathcal{X}, f)$ can be found by applying standard optimization tools to the problem $\arg \min_{x \in \mathcal{X}} J(x)$. A necessary and sufficient condition for the existence of a function J such that $f(x) = \nabla_x J(x)$ is $\nabla_x f(x) = \nabla_x f(x)^\top$ for all x , as in that case $\nabla_x f(x)$ would be the Hessian of $J(x)$ [FP03, Theorem 1.3.1]. A necessary and sufficient condition for such a J to be convex is that the operator $f(x) = \nabla_x J(x)$ is **MON**. If the operator f is additionally **SMON**, then J is strongly convex. This implies that the minimizer of (3.4), and consequently the solution of VI(\mathcal{X}, f), is unique. In Section 3.3.1 we show that **SMON** implies uniqueness of the solution of a VI, even when f is not a gradient operator.

Connection with fixed point theory

Another interesting characterization of the solution of a variational inequality is in terms of fixed points, see Figure 3.3.

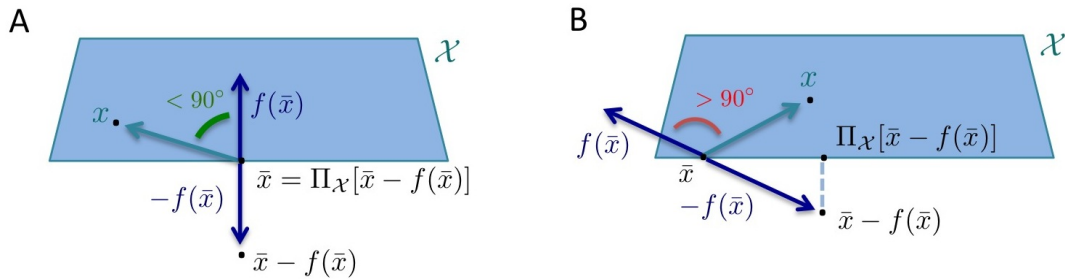


Figure 3.3: Illustration of VI and fixed point equivalence. A) The point $\bar{x} \in \text{SOL}(\mathcal{X}, f)$ because $f(\bar{x})^\top(x - \bar{x}) \geq 0$ for all $x \in \mathcal{X}$ or, equivalently, the angle between $x - \bar{x}$ and $f(\bar{x})$ is less than 90° for all $x \in \mathcal{X}$. Equivalently $\bar{x} = \Pi_{\mathcal{X}}[\bar{x} - f(\bar{x})]$. B) The point $\bar{x} \notin \text{SOL}(\mathcal{X}, f)$ because there exists $x \in \mathcal{X}$ such that the angle between $x - \bar{x}$ and $f(\bar{x})$ is more than 90° . Equivalently, $\bar{x} \neq \Pi_{\mathcal{X}}[\bar{x} - f(\bar{x})]$.

Many algorithms for solving VIs are indeed fixed point algorithms, similar to those described in Section 3.2.2.

Proposition 3.3.3. (Fixed point equivalence) *Suppose that $\mathcal{X} \subseteq \mathbb{R}^n$ is closed and convex. Then, for any $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$*

$$\bar{x} \in \text{SOL}(\mathcal{X}, f) \quad \Leftrightarrow \quad \bar{x} = \Pi_{\mathcal{X}}(\bar{x} - f(\bar{x})).$$

Proof. This is a classical result. See for example [FP03, Proposition 1.5.8] or [BT97, Chapter 3, Proposition 5.1]. \square

Karush-Kuhn-Tucker (KKT) conditions

Consider again the optimization problem (3.4) and assume that the set \mathcal{X} can be parametrized as follows

$$\mathcal{X} = \{x \in \mathbb{R}^n \mid h(x) = 0, g(x) \leq 0\}, \quad (3.6)$$

for suitable \mathcal{C}^1 functions $h : \mathbb{R}^n \rightarrow \mathbb{R}^l$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$. To derive the following statements we need to assume that the set \mathcal{X} is regular enough. In the literature, a large number of conditions that ensure regularity of a constraint set have been derived under the name of **constraint qualification (CQ)** conditions. In the following, we use one of the most general, which is known as Abadie CQ [FP03, Section 1.3.1]. Since verifying this property can be tedious, we report in the next lemma some sufficient conditions that ensure its validity.

Lemma 3.3.4. (Constraint qualification [FP03, Section 3.2]) *Let \mathcal{X} be given by (3.6), where the functions $h : \mathbb{R}^n \rightarrow \mathbb{R}^l$, $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are \mathcal{C}^1 , and $\bar{x} \in \mathcal{X}$. Let $I(\bar{x}) := \{i \in \mathbb{Z}[1, m] \mid g_i(\bar{x}) = 0\}$ be the set of active inequality constraints at \bar{x} and define the following CQs.*

1. [ACQ: Abadie CQ] *The tangent cone of \mathcal{X} at \bar{x} , which is defined as*

$$\mathcal{T}(\bar{x}; \mathcal{X}) := \left\{ v \in \mathbb{R}^n \mid \lim_{\nu \rightarrow \infty} \frac{x_\nu - \bar{x}}{\tau_\nu} = v, \exists \{x_\nu \in \mathcal{X}\}_{\nu=1}^\infty \rightarrow \bar{x}, \exists \{\tau_\nu \in \mathbb{R}_{>0}\}_{\nu=1}^\infty \rightarrow 0 \right\}$$

is equal to the linearization cone of \mathcal{X} at \bar{x} , which is defined as

$$\mathcal{L}(\bar{x}; \mathcal{X}) := \{v \in \mathbb{R}^n \mid v^\top \nabla_x h_j(\bar{x}) = 0, \forall j \in \mathbb{Z}[1, l], v^\top \nabla_x g_i(\bar{x}) \leq 0, \forall i \in I(\bar{x})\}.$$

Mathematically, $\mathcal{T}(\bar{x}; \mathcal{X}) = \mathcal{L}(\bar{x}; \mathcal{X})$.

2. [MFCQ: Mangasarian and Fromovitz CQ] *The gradients $\{\nabla_x h_j(\bar{x}), j \in \mathbb{Z}[1, l]\}$ are linearly independent and there exists $d \in \mathbb{R}^n$ such that $d^\top \nabla_x g_i(\bar{x}) < 0$ for all $i \in I(\bar{x})$, $d^\top \nabla_x h_j(\bar{x}) = 0$ for all $j \in \mathbb{Z}[1, l]$.*
3. [LICQ: Linear independence CQ] *The gradients $\{\nabla_x h_j(\bar{x}), j \in \mathbb{Z}[1, l]\} \cup \{\nabla_x g_i(\bar{x}), i \in I(\bar{x})\}$, are linearly independent.*

4. [SLCQ: Slater CQ] The functions $h_j(x)$ for all $j \in \mathbb{Z}[1, l]$ are affine and $g_i(x)$ for all $i \in \mathbb{Z}[1, m]$ are convex. The gradients $\{\nabla_x h_j(x), j \in \mathbb{Z}[1, l]\}$ are linearly independent and there exists a point $x^S \in \mathcal{X}$ such that $h(x^S) = 0$ and $g(x^S) < 0$.

Then

$$\begin{aligned} \text{LICQ}(\bar{x}) &\Rightarrow \text{MFCQ}(\bar{x}) \Rightarrow \text{ACQ}(\bar{x}) \\ \text{SLCQ} &\Rightarrow \text{MFCQ}(\bar{x}) \quad \text{for all } \bar{x} \in \mathcal{X} \end{aligned}$$

□

Remark 3.1. Note that while $\text{ACQ}(\bar{x})$, $\text{MFCQ}(\bar{x})$, $\text{LICQ}(\bar{x})$ are local properties that need to be checked at the (usually a priori unknown) optimal point \bar{x} , SLCQ is a global property that can be easily checked. This is the property that we use in Chapter 5. □

Proposition 3.3.5. (*Karush-Kuhn-Tucker conditions (KKT) for optimization problems*) Let \mathcal{X} be given by (3.6). Assume that the functions $h : \mathbb{R}^n \rightarrow \mathbb{R}^l$, $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $J : \mathbb{R}^n \rightarrow \mathbb{R}$ are \mathcal{C}^1 and that Abadie's CQ holds at a point $\bar{x} \in \mathcal{X}$. Then \bar{x} solves (3.5) if and only if there exist $\bar{\mu} \in \mathbb{R}^l$, $\bar{\lambda} \in \mathbb{R}^m$ such that

$$\begin{aligned} \nabla_x J(\bar{x}) + \sum_{j=1}^l \bar{\mu}_j \nabla_x h_j(\bar{x}) + \sum_{i=1}^m \bar{\lambda}_i \nabla_x g_i(\bar{x}) &= 0 \\ h(\bar{x}) &= 0, \quad g(\bar{x}) \leq 0, \quad \bar{\lambda} \geq 0 \\ \bar{\lambda} &\perp g(\bar{x}) \end{aligned} \tag{3.7}$$

Let $\Lambda(\bar{x})$ be the set of vector pairs $(\bar{\mu}, \bar{\lambda})$ satisfying (3.7). If $\text{MFCQ}(\bar{x})$ holds then $\Lambda(\bar{x})$ is compact and if $\text{LICQ}(\bar{x})$ holds then $\Lambda(\bar{x})$ is a singleton.

Proof. See [FP03, Section 1.3.1] and [Wac13]. □

As we noted before, (3.5) is a special case of VI. It turns out that the previous proposition can be extended to any VI, under the assumption that the set \mathcal{X} is convex.

Proposition 3.3.6. (*KKT for VIs* [FP03, Proposition 1.3.4 and 3.2.1]) Let \mathcal{X} be given by (3.6), where the functions $h : \mathbb{R}^n \rightarrow \mathbb{R}^l$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ are \mathcal{C}^1 . Let f be a mapping from \mathcal{X} into \mathbb{R}^n . The following two statements hold.

1. Let $\bar{x} \in \text{SOL}(\mathcal{X}, f)$. If $\text{ACQ}(\bar{x})$ holds, then there exist vectors $\bar{\mu} \in \mathbb{R}^l$, $\bar{\lambda} \in \mathbb{R}^m$ such that

$$\begin{aligned} f(\bar{x}) + \sum_{j=1}^l \bar{\mu}_j \nabla_x h_j(\bar{x}) + \sum_{i=1}^m \bar{\lambda}_i \nabla_x g_i(\bar{x}) &= 0 \\ h(\bar{x}) &= 0, \quad g(\bar{x}) \leq 0, \quad \bar{\lambda} \geq 0 \\ \bar{\lambda} &\perp g(\bar{x}) \end{aligned} \tag{3.8}$$

2. Conversely, if each function h_j is affine, each function g_i is convex and if $(\bar{x}, \bar{\mu}, \bar{\lambda})$ satisfies (3.8), then $\bar{x} \in \text{SOL}(\mathcal{X}, f)$.
3. Let f be continuous, $\bar{x} \in \text{SOL}(\mathcal{X}, f)$ and $\Lambda(\bar{x})$ be the set of vector pairs $(\bar{\mu}, \bar{\lambda})$ satisfying (3.8). If $\text{MFCQ}(\bar{x})$ holds then $\Lambda(\bar{x})$ is non-empty and bounded and if $\text{LICQ}(\bar{x})$ holds then $\Lambda(\bar{x})$ is a singleton. \square

3.3.1 Existence and uniqueness

We report in the following a main result on existence and uniqueness of the solution to $\text{VI}(\mathcal{X}, f)$, [FP03, Corollary 2.2.5, Theorem 2.3.3].

Proposition 3.3.7 (Existence and uniqueness). *Consider the $\text{VI}(\mathcal{X}, f)$. The following statements hold.*

1. If \mathcal{X} is compact and convex and f is continuous then $\text{SOL}(\mathcal{X}, f)$ is non-empty and compact.
2. If \mathcal{X} is closed and convex and f is continuous and **SMON** then the $\text{VI}(\mathcal{X}, f)$ has a unique solution. \square

3.3.2 Iterative algorithms

Finding a solution to a **VI** is equivalent to finding a fixed point of a n -dimensional mapping, which is in general a very difficult task. However, if the mapping f has one of the regularity properties given in Definition 3.1, then one can use known algorithms that aim at iteratively constructing a sequence $\{x_{(k)}\}_{k=1}^{\infty}$ such that $x_{(k)} \rightarrow \bar{x} \in \text{SOL}(\mathcal{X}, f)$. In the following we review the ones that are used in the thesis.

SMON operator

The simplest case is when the operator f is **SMON** and Lipschitz and the set \mathcal{X} is closed and convex. According to Proposition 3.3.7, in this case, the $\text{VI}(\mathcal{X}, f)$ has a unique solution that can be found using the following projection algorithm, with a sufficiently small step size $\alpha > 0$.

Projection algorithm

$$x_{(k+1)} = \Pi_{\mathcal{X}}(x_{(k)} - \alpha f(x_{(k)})) \quad (3.9)$$

Proposition 3.3.8. (Projection algorithm: **SMON** [FP03, Theorem 12.1.2]) *Assume that $\mathcal{X} \subseteq \mathbb{R}^n$ is closed and convex and that f is **SMON** with constant ϵ and Lipschitz*

with constant L . For any $\alpha < \frac{2\epsilon}{L^2}$, Algorithm (3.9) converges for any initial condition $x_{(0)} \in \mathcal{X}$ to the unique solution \bar{x} of $VI(\mathcal{X}, f)$. Moreover, let $q := \sqrt{1 + (\alpha L)^2 - 2\alpha\epsilon}$, then $\|x_{(k)} - \bar{x}\| \leq q^k \|x_{(0)} - \bar{x}\|$. \square

Convergence of the projection algorithm can be proven also under a slightly different condition (i.e., the co-coercivity property) and for time-varying step sizes (i.e., using α_k instead of α in (3.9)).

Definition 3.5. A mapping $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is called *co-coercive (COC)* in \mathcal{H}_S if there exists $c > 0$ such that $(x - y)^\top S(f(x) - f(y)) \geq c \|f(x) - f(y)\|_S^2, \forall x, y \in \mathbb{R}^n$. \square

Note that if a mapping f is COC with constant $c \geq 1$ then it is FNE (by Lemma 3.1.1). Moreover, COC implies MON but not necessarily SMON (e.g., take $f(x) = x$).

Proposition 3.3.9. (Projection algorithm: COC [FP03, Theorem 12.1.8]) Assume that $\mathcal{X} \subseteq \mathbb{R}^n$ is closed and convex and that $f : \mathcal{X} \rightarrow \mathbb{R}^n$ is COC on \mathcal{X} with constant c . Suppose that $SOL(\mathcal{X}, f)$ is non-empty. Consider a sequence $\{\alpha_k \in \mathbb{R}\}_{k=1}^\infty$, such that,

$$0 < \inf \alpha_k \leq \sup \alpha_k < 2c$$

then the algorithm

$$x_{(k+1)} = \Pi_{\mathcal{X}}(x_{(k)} - \alpha_k f(x_{(k)}))$$

produces a sequence $x_{(k)}$ converging to a solution of the $VI(\mathcal{X}, f)$. \square

MON operator

If the operator is MON but not SMON, then the projection algorithm is not guaranteed to converge [FP03, Example 12.1.3]. One can however, use the slightly more complex extragradient algorithm.

Extragradient algorithm

$$\begin{aligned} \tilde{x}_{(k)} &= \Pi_{\mathcal{X}}(x_{(k)} - \alpha f(x_{(k)})) \\ x_{(k+1)} &= \Pi_{\mathcal{X}}(x_{(k)} - \alpha f(\tilde{x}_{(k)})) \end{aligned} \tag{3.10}$$

Proposition 3.3.10. (Extragradient algorithm [FP03, Theorem 12.1.11]) Assume that $\mathcal{X} \subseteq \mathbb{R}^n$ is closed and convex and that f is MON and Lipschitz with constant L . Suppose that $SOL(\mathcal{X}, f)$ is non-empty. If $\alpha < \frac{1}{L}$ then the sequence $x_{(k)}$ generated by Algorithm (3.10) converges to a solution of the $VI(\mathcal{X}, f)$. \square

The main disadvantage of the extragradient algorithm is that it requires two gradient computations and two projections at every step, thus slowing the overall performance. A large number of alternative algorithms have been suggested in the literature to solve, either exactly or approximately, a VI with MON operators. We refer the reader to [FP03, Chapter 12]. In the next subsection we propose a new approach that can potentially speed up convergence under additional assumptions on the problem structure (in line with the requirements of Chapter 5).

MON operator: a special case

In Chapter 5 we will consider games where, additionally to the local constraint sets $\mathcal{X}^i \subset \mathbb{R}^n$, there is also a constraint $\mathcal{C} := \{x = [x^1; \dots; x^N] \mid Ax \leq b\} \subset \mathbb{R}^{Nn}$ coupling together the strategies of all the players, where $A \in \mathbb{R}^{m \times Nn}$, $b \in \mathbb{R}^m$. We show in Chapter 5 that the coordination problem, in this case, can be solved by finding the solution of a variational inequality of the form $\text{VI}(\mathcal{Y}, t)$ where $\mathcal{Y} := \mathcal{X} \times \mathbb{R}_{\geq 0}^m$, \mathcal{X} is a convex and compact set and the operator $t : \mathbb{R}^{Nn+m} \rightarrow \mathbb{R}^{Nn+m}$ is defined as follows

$$t(y) := t\left(\begin{bmatrix} x \\ \lambda \end{bmatrix}\right) = \begin{bmatrix} f(x) + A^\top \lambda \\ -(Ax - b) \end{bmatrix}, \quad (3.11)$$

for a suitable operator $f : \mathbb{R}^{Nn} \rightarrow \mathbb{R}^{Nn}$. The additional variable $\lambda \in \mathbb{R}_{\geq 0}^m$ can be thought of as a dual variable associated to the coupling constraints.

In this section we propose a new algorithm to find a solution $\bar{y} := \begin{bmatrix} \bar{x} \\ \bar{\lambda} \end{bmatrix}$ of the $\text{VI}(\mathcal{Y}, t)$, under the assumption that the operator f is SMON and affine. We briefly note that, under these assumptions, the operator t in (3.11) is MON and Lipschitz. Consequently, a solution $\bar{y} \in \text{SOL}(\mathcal{Y}, t)$, if it exists, can be found by applying the extragradient algorithm presented in (3.10). Our objective is to propose an alternative approach that, contrary to the extragradient, requires only one projection operation for each step. To this end, we start by introducing a class of algorithms, parametrized by a matrix $D \succ 0$ not necessarily symmetric, that has been proposed in the literature to solve a generic $\text{VI}(\mathcal{Y}, t)$, as illustrated in (3.12).

Asymmetric projection algorithm for a generic $D \succ 0$

$$y_{(k+1)} = \text{SOL}(\mathcal{Y}, t_D^k), \quad t_D^k(y) := t(y_{(k)}) + D(y - y_{(k)}) \quad (3.12)$$

Intuitively, an APA aims at solving the original $\text{VI}(\mathcal{Y}, t)$ by solving at every step a different affine VI that depends on the previous vector $y_{(k)}$ and on the matrix D . Therefore APAs are a special class of linearized algorithms for VIs [BT97, Section 3.5.4]. Note that, since $D \succ 0$, the affine operator t_D^k is SMON and continuous, therefore if \mathcal{Y} is convex and closed the solution of $\text{VI}(\mathcal{Y}, t_D^k)$ exists and is unique (Proposition 3.3.7),

thus guaranteeing that the update step in (3.12) is well-defined. Each choice of matrix $D \succ 0$ leads to a different APA. The simplest one is obtained by selecting $D = \frac{1}{\alpha}I$. In this case in fact, if \mathcal{Y} is closed and convex, one can see that the corresponding APA coincides with the projection algorithm given in (3.9) since, by the minimum principle,

$$\bar{y} = \text{SOL}(\mathcal{Y}, t_D^k) \Leftrightarrow \bar{y} = \arg \min_{y \in \mathcal{Y}} \left(\frac{1}{2\alpha} y^\top y + [t(y^{(k)}) - \frac{1}{\alpha} y^{(k)}]^\top y \right) \Leftrightarrow \bar{y} = \Pi_{\mathcal{Y}}[y^{(k)} - \alpha t(y^{(k)})].$$

Unfortunately, the choice $D = \frac{1}{\alpha}I$ does not guarantee convergence in the case that we are considering here because the operator t as defined in (3.11) is neither SMON nor COC (to see this set $y_1 = [\lambda_1]$ and $y_2 = [\lambda_2]$ with $x_1 = x_2$). For APAs with a generic matrix D a sufficient condition for convergence has been derived in [FP03, Proposition 12.5.2] and is reported in the next lemma.

Lemma 3.3.11. *Consider the VI(\mathcal{Y}, t) with \mathcal{Y} convex and closed and t continuous. Suppose that $\text{SOL}(\mathcal{Y}, t)$ is non-empty. Fix a matrix $D \succ 0$, set $D_s = (D + D^\top)/2$ and denote by $D_s^{-1/2}$ the principal square root of the symmetric positive definite matrix D_s^{-1} . If the operator*

$$g(y) = D_s^{-1/2} t(D_s^{-1/2} y) - D_s^{-1/2} (D - D_s) D_s^{-1/2} y$$

is COC with modulus greater than 1/2 then Algorithm (3.12) converges to a solution of VI(\mathcal{Y}, T).

The convergence rate of the APAs has been studied in [FP03, Chapter 12]. In the following, we consider the specific APA obtained by selecting

$$D := \begin{bmatrix} \frac{1}{\tau} I_{Nn} & 0 \\ -2A & \frac{1}{\tau} I_m \end{bmatrix} \Rightarrow D_s = \begin{bmatrix} \frac{1}{\tau} I_{Nn} & -A^\top \\ -A & \frac{1}{\tau} I_m \end{bmatrix} \quad (3.13)$$

which, as proven in the next theorem, results in the following algorithm which alternates updates in the primal (x) and dual (λ) variables and depends on a step size parameter $\tau > 0$.

Asymmetric projection algorithm

$$\begin{aligned} x_{(k+1)} &= \Pi_{\mathcal{X}}[x_{(k)} - \tau (f(x_{(k)}) + A^\top \lambda_{(k)})] \\ \lambda_{(k+1)} &= \Pi_{\mathbb{R}_{\geq 0}^m}[\lambda_{(k)} - \tau (b - 2Ax_{(k+1)} + Ax_{(k)})] \end{aligned} \quad (3.14)$$

We note that the proposed scheme differs from the projection algorithm, which is not guaranteed to converge for MON operators, since the λ -update depends not only on $Ax_{(k)}$ but also on $Ax_{(k+1)}$. Our main result is to prove that if t has the structure given in (3.11) and $f(x) := Fx + \tilde{f}$ is affine and SMON, then the sufficient condition of Lemma 3.3.11 is met for D as in (3.13), even if $F \neq F^\top$.

We briefly note that convergence for the symmetric case $F = F^\top$ has been already proven in the literature [FP03, Proposition 12.5.3 (b)]. Even if seemingly mild, the missing assumption $F \neq F^\top$ implies that there exists no cost function $J(x)$ such that $f(x) = \nabla_x J(x)$. Consequently, the considered VI(\mathcal{Y}, t) cannot be solved using convex optimization tools. On the contrary, we show in the next theorem that the condition $F + F^\top \succ 0$ is sufficient to guarantee convergence of Algorithm (3.14) for this new class of affine VIs.

Theorem 3.3.12. *Consider the VI(\mathcal{Y}, t), where $t(y)$ is defined as in (3.11) and suppose that $SOL(\mathcal{Y}, t)$ is non-empty. Suppose that there exists a function $g : \mathbb{R}^{N_n} \rightarrow \mathbb{R}^{m_g}$, $g \in \mathcal{C}^1$ such that $\mathcal{X} = \{x \in \mathbb{R}^{N_n} \mid g(x) \leq 0\}$, that \mathcal{X} satisfies the Slater CQ condition and that the operator f is affine (i.e., $f(x) = Fx + \tilde{f}$, $F \in \mathbb{R}^{N_n \times N_n}$, $\tilde{f} \in \mathbb{R}^{N_n}$) and SMON. Let $y_{(k)} := \begin{bmatrix} x_{(k)} \\ \lambda_{(k)} \end{bmatrix}$ be the state of Algorithm (3.14) at iteration k and set*

$$\kappa_F := \frac{\sigma_{\max}^2(F)}{\sigma_{\min}(F_s)}, \quad 0 < \tau < \left(\frac{\kappa_F + \sqrt{\kappa_F^2 + 4\|A^\top A\|}}{2} \right)^{-1},$$

where $F_s := (F + F^\top)/2$. Then the sequence $y_{(k)}$ converges, for any initial condition, to $\bar{y} \in SOL(\mathcal{Y}, t)$.

Proof. We divide the proof into two parts: (i) we prove that Algorithm (3.14) is a particular case of APAs, applied to the VI(\mathcal{Y}, t); (ii) we prove that Algorithm (3.14) satisfies the sufficient convergence condition for APAs. Throughout the proof we use that for any matrix B , $\sigma_{\max}(B) = \|B\|$ and if $B = B^\top \succeq 0$ then $\sigma_{\max}(B) = \lambda_{\max}(B)$ and $\sigma_{\max}(B)I \succeq B \succeq \sigma_{\min}(B)I$. Moreover, we define $B_s := (B + B^\top)/2$.

(i) We are going to show that Algorithm (3.14) coincides with the APA defined by the matrix D given in (3.13). Note that, from the Schur complement, $D \succ 0 \Leftrightarrow D_s \succ 0 \Leftrightarrow \frac{1}{\tau}I - \tau A^\top A \succ 0$. The assumption on τ implies $2/\tau > \kappa_F + \sqrt{\kappa_F^2 + 4\|A^\top A\|} > \sqrt{\kappa_F^2 + 4\|A^\top A\|}$ therefore $4/\tau^2 > 4\|A^\top A\|$ and

$$\tau^2\|A^\top A\| < 1. \tag{3.15}$$

Therefore, since $A^\top A \succeq 0$, $\frac{1}{\tau}I - \tau A^\top A \succeq \frac{1}{\tau}(1 - \tau^2\|A^\top A\|)I \succ 0$ and thus $D \succ 0$. Let us characterize the set $\mathbb{R}_{\geq 0}^m = \{\lambda \in \mathbb{R}^m \mid h(\lambda) := -\lambda \leq 0\}$. By the Slater CQ condition and by Proposition 3.3.6 the unique solution of VI(\mathcal{Y}, t_D^k) coincides with the solution $[x, \lambda, \mu_1, \mu_2]$ of the KKT system

$$\begin{aligned} t_D^k(y) + \nabla_y g(x)\mu_1 + \nabla_y h(\lambda)\mu_2 &= 0 \\ \mu_1 &\geq 0, \quad g(x) \leq 0, \quad \mu_1 \perp g(x) \\ \mu_2 &\geq 0, \quad h(\lambda) \leq 0, \quad \mu_2 \perp h(\lambda). \end{aligned} \tag{3.16}$$

The equality in (3.16) can be rewritten as

$$\begin{bmatrix} \frac{1}{\tau}(x - x_{(k)}) + Fx_{(k)} + \tilde{f} + A^\top \lambda_{(k)} + \nabla_x g(x) \mu_1 \\ \frac{1}{\tau}(\lambda - \lambda_{(k)}) + b - 2Ax + Ax_{(k)} + \nabla_\lambda h(\lambda) \mu_2 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}.$$

Therefore (3.16) holds iff both K_1 and K_2 , defined as

$$K_1 := \begin{cases} \frac{1}{\tau}(x - x_{(k)}) + Fx_{(k)} + \tilde{f} + A^\top \lambda_{(k)} + \nabla_x g(x) \mu_1 = 0 \\ \mu_1 \geq 0, g(x) \leq 0, \mu_1 \perp g(x) \end{cases}$$

$$K_2 := \begin{cases} \frac{1}{\tau}(\lambda - \lambda_{(k)}) + b - 2Ax + Ax_{(k)} + \nabla_\lambda h(\lambda) \mu_2 = 0 \\ \mu_2 \geq 0, h(\lambda) \leq 0, \mu_2 \perp h(\lambda) \end{cases}$$

hold. These are the KKT systems of two optimization problems, respectively (see Proposition 3.3.5). More in detail, K_1 does not depend on $[\lambda, \mu_2]$ and $[x, \mu_1]$ solves K_1 if and only if

$$x := \arg \min_{x \in \mathcal{X}} \left\{ \frac{1}{2} \|x\|^2 - x_{(k)}^\top x + \tau(Fx_{(k)} + \tilde{f} + A^\top \lambda_{(k)})^\top x \right\}$$

$$= \Pi_{x \in \mathcal{X}} [x_{(k)} - \tau(Fx_{(k)} + \tilde{f} + A^\top \lambda_{(k)})].$$

For every value of x (and hence for the one determined by K_1), $[\lambda, \mu_2]$ solves K_2 if and only if λ solves the optimization problem

$$\lambda := \arg \min_{\lambda \in \mathbb{R}_{\geq 0}^m} \left\{ \frac{1}{2} \|\lambda\|^2 - \lambda_{(k)}^\top \lambda + \tau(b - 2Ax + Ax_{(k)})^\top \lambda \right\}$$

$$= \Pi_{\mathbb{R}_{\geq 0}^m} [\lambda_{(k)} - \tau(b - 2Ax + Ax_{(k)})].$$

Summing up, $y_{(k+1)} := \begin{bmatrix} x_{(k+1)} \\ \lambda_{(k+1)} \end{bmatrix}$, which is by definition the unique solution of (3.16), can be explicitly characterized as

$$x_{(k+1)} = \Pi_{x \in \mathcal{X}} [x_{(k)} - \tau(Fx_{(k)} + \tilde{f} + A^\top \lambda_{(k)})]$$

$$\lambda_{(k+1)} = \Pi_{\mathbb{R}_{\geq 0}^m} [\lambda_{(k)} + \tau(b - 2Ax_{(k+1)} + Ax_{(k)})]$$

which coincide with the update steps of Algorithm (3.14).

(ii) As illustrated in the previous point, Algorithm (3.14) is the specific APA associated with the choice of D given in (3.13). According to Lemma 3.3.11, any APA converges if the mapping $g(y) = D_s^{-1/2}(T - D + D_s)D_s^{-1/2}y + D_s^{-1/2}\tilde{t} =: Gy + \tilde{g}$ is COC with modulus greater than 1/2. We are now going to show that this sufficient condition is met for the specific operator $t(y)$ and matrix D that we are considering. For simplicity let us rename $L := D_s^{-1/2}$. Note that $G = L \begin{bmatrix} F & 0 \\ 0 & 0 \end{bmatrix} L$. We are going to show that, under the given condition on τ , $g(y)$ is COC with modulus 1 or equivalently $v^\top G_s v - v^\top G^\top G v \geq 0, \forall v$. Note that

$$v^\top G_s v - v^\top G^\top G v = v^\top L \left(\begin{bmatrix} F_s & 0 \\ 0 & 0 \end{bmatrix} - \begin{bmatrix} F^\top & 0 \\ 0 & 0 \end{bmatrix} L^\top \begin{bmatrix} F & 0 \\ 0 & 0 \end{bmatrix} \right) L v$$

$$= w^\top \left(\begin{bmatrix} F_s - F^\top [L^\top]_{11} F & 0 \\ 0 & 0 \end{bmatrix} \right) w,$$

where $w := Lv$ and $[L^2]_{11}$ denotes the block in position $(1, 1)$ of the matrix L^2 . Since L is invertible, the statement is proven if $F_s - F^\top [L^2]_{11} F \succeq 0$. To this end, let us compute $[L^2]_{11} \equiv [D_s^{-1}]_{11}$. From the the formula for the inverse of a block matrix applied to D_s we get

$$[L^2]_{11} = \tau(I - \tau^2 A^\top A)^{-1}. \quad (3.17)$$

Since $\tau^2 A^\top A$ is symmetric positive semidefinite, $\lambda_{max}(\tau^2 A^\top A) = \tau^2 \|A^\top A\| < 1$ (by (3.15)) and the matrix is convergent. Consequently, the Neumann series $\sum_{k=0}^{\infty} (\tau^2 A^\top A)^k$ converges to $(I - \tau^2 A^\top A)^{-1}$. Substituting in (3.17) yields $[L^2]_{11} = \tau \sum_{k=0}^{\infty} (\tau^2 A^\top A)^k \succeq 0$ and

$$\|[L^2]_{11}\| \leq \tau \sum_{k=0}^{\infty} (\tau^2 \|A^\top A\|)^k = \frac{\tau}{1 - \tau^2 \|A^\top A\|},$$

where we used the fact that the geometric series converges since $\tau^2 \|A^\top A\| < 1$ (by (3.15)). Finally, $F \succ 0 \Rightarrow F_s = F_s^\top \succ 0$ and $[L^2]_{11} = [L^2]_{11}^\top \succeq 0 \Rightarrow F^\top [L^2]_{11} F = (F^\top [L^2]_{11} F)^\top \succeq 0$. Therefore

$$\begin{aligned} F_s - F^\top [L^2]_{11} F &\succeq (\sigma_{min}(F_s) - \sigma_{max}(F^\top [L^2]_{11} F))I \\ &\succeq (\sigma_{min}(F_s) - \|F\|^2 \|[L^2]_{11}\|)I \\ &= (\sigma_{min}(F_s) - \sigma_{max}(F)^2 \|[L^2]_{11}\|)I \\ &\succeq \sigma_{min}(F_s) \left(1 - \kappa_F \frac{\tau}{1 - \tau^2 \|A^\top A\|} \right) I \succ 0. \end{aligned}$$

In the last line we used $1 - \kappa_F \frac{\tau}{1 - \tau^2 \|A^\top A\|} > 0$, this is in fact equivalent (by (3.15)) to $1 - \tau^2 \|A^\top A\| - \kappa_F \tau > 0 \Leftrightarrow \frac{1}{\tau^2} - \|A^\top A\| - \frac{1}{\tau} \kappa_F > 0 \Leftrightarrow \frac{1}{\tau} > \frac{\kappa_F + \sqrt{\kappa_F^2 + 4\|A^\top A\|}}{2}$, which is true by assumption. □

Myopic agents

In this chapter we derive algorithms to coordinate the strategies of a population of *myopic* agents to a Nash equilibrium of different types of aggregative games

$$\mathcal{G} = \left\{ \begin{array}{l} x_{\text{br}}^i(x^{-i}) := \arg \min_{x^i \in \mathbb{R}^n} J^i(x^i, \sigma^i(x)) \\ \text{s.t. } x^i \in \mathcal{X}^i \end{array} \quad \forall i = 1, \dots, N. \right. \quad (4.1)$$

Specifically, in Section 4.2 we propose a *decentralized algorithm* that can be used to steer a population of agents to a Nash equilibrium of an **AAG** by means of a reference signal z , that is broadcast by a central operator. In Section 4.3 we adapt the proposed algorithm to the case of **NAGs**. To this end, we assume that the agents, instead of reacting to a global reference z , react to a local estimate z^i of the neighbors aggregate state which they update by *network communications*. Finally, in Section 4.4 we illustrate how the setting developed for **NAGs** can be used to coordinate the agents in **AAGs** *distributedly*, that is, by using local communications instead of relying on the presence of a central operator (thus complementing the results of Section 4.2). All the proofs are given in the Appendix (Section 4.5).

All the algorithms suggested in this chapter are derived under the fundamental assumption that the agents are *myopic*. This notion is formalized in Section 4.1; however, it intuitively corresponds to the assumption that, at every iteration of the coordination algorithm, each agent selects as next strategy the one that minimizes its cost function, given the current reference z or its current estimate z^i . This assumption is motivated by those applications where the agents are fully noncooperative and/or the algorithmic iterations correspond to repetitions of the game over time (e.g. in opinion dynamics). The assumption of myopic agents significantly complicates the coordination problem, since it limits the class of allowed update rules. In Chapter 5 we show how different algorithms, based on gradient steps, can be used when this assumption is relaxed to bounded rationality.

The results of this chapter have been published in [PCGL14, PGCL15, GPCL16, PGGL15a, PGGL15b]

4.1 Assumptions on the agents and optimal responses

Motivated by large-scale applications and privacy issues, we assume throughout this chapter that the agents do not have access to the individual strategies of the other agents, but can instead compute their *optimal response* to a fixed signal $z^i \in \mathbb{R}^n$, that is

$$x^{i*}(z^i) := \arg \min_{x^i \in \mathbb{R}^n} J^i(x^i, z^i) \quad (4.2)$$

s.t. $x^i \in \mathcal{X}^i$

Note that this mapping differs from the *best response* mapping, given in (4.1), since the second argument in the cost function $J^i(\cdot, \cdot)$ is not the aggregate state $\sigma^i(x)$ (that might be unknown and in general could depend on x^i), but is a fixed exogenous vector z^i (that does not depend on the optimization variable x^i). Specifically, in Section 4.2 we assume that $z^i = \bar{z}$ for all $i \in \mathbb{Z}[1, N]$ where \bar{z} is a reference broadcast by a central operator to the whole population, while in Section 4.3 and 4.4 we assume that z^i is different for each agent and represents a local estimate of the neighbors aggregate state.

Since the difference between optimal response and best response, as well as the relation between $J^i(x^i, \sigma^i(x))$ and $J^i(x^i, z^i)$ is of fundamental importance for the subsequent results, we present here a clarifying example.

Example 4.1. Consider a game with $N = 3$ players, with strategies x^1, x^2 and x^3 , and whose cost functions are

$$\begin{aligned} J^1(x^1, \sigma^1(x)) &:= [\sigma^1(x)]^\top x^1 := [(x^1 + x^2)/2]^\top x^1 =: \tilde{J}^1(x^1, x^2, x^3) =: \tilde{J}^1(x^1, x^{-1}), \\ J^2(x^2, \sigma^2(x)) &:= [\sigma^2(x)]^\top x^2 := [x^1]^\top x^2 =: \tilde{J}^2(x^1, x^2, x^3) =: \tilde{J}^2(x^2, x^{-2}), \\ J^3(x^3, \sigma^3(x)) &:= [\sigma^3(x)]^\top x^3 := [(x^1 + x^2 + x^3)/3]^\top x^3 =: \tilde{J}^3(x^1, x^2, x^3) =: \tilde{J}^3(x^3, x^{-3}), \end{aligned}$$

where we used $\sigma^1(x) := \frac{x^1 + x^2}{2}$, $\sigma^2(x) := x^1$ and $\sigma^3(x) := \frac{x^1 + x^2 + x^3}{3}$ as aggregator functions. We firstly note that the cost functions depend on the other players strategies only via the aggregate quantities σ^i , therefore this is an aggregative game. More in detail, it is a network aggregative game relative to the matrix

$$P = \begin{bmatrix} 1/2 & 1/2 & 0 \\ 1 & 0 & 0 \\ 1/3 & 1/3 & 1/3 \end{bmatrix}.$$

For simplicity we assume that there are no constraints, so that $x^1, x^2, x^3 \in \mathbb{R}^n$. The best

response for the three players, given the strategies of the others are

$$\begin{aligned}
x_{\text{br}}^1(x^{-1}) &:= \arg \min_{x^1} J^1(x^1, \sigma^1(x)) = \arg \min_{x^1} \tilde{J}^1(x^1, x^{-1}) = \arg \min_{x^1} [(x^1 + x^2)/2]^\top x^1 \\
x_{\text{br}}^2(x^{-2}) &:= \arg \min_{x^2} J^2(x^2, \sigma^2(x)) = \arg \min_{x^2} \tilde{J}^2(x^2, x^{-2}) = \arg \min_{x^2} [x^1]^\top x^2 \\
x_{\text{br}}^3(x^{-3}) &:= \arg \min_{x^3} J^3(x^3, \sigma^3(x)) = \arg \min_{x^3} \tilde{J}^3(x^3, x^{-3}) = \arg \min_{x^3} [(x^1 + x^2 + x^3)/3]^\top x^3.
\end{aligned} \tag{4.3}$$

Note that, since the best responses depend on the strategies of the other players the three optimization problems in (4.3) are coupled. A joint solution is, by definition, a Nash equilibrium. On the other hand, consider three reference vectors z^1, z^2, z^3 , that are fixed and independent on x^1, x^2, x^3 . The optimal responses are

$$\begin{aligned}
x^{1*}(z^1) &:= \arg \min_{x^1} J^1(x^1, z^1) = \arg \min_{x^1} [z^1]^\top x^1 \\
x^{2*}(z^2) &:= \arg \min_{x^2} J^2(x^2, z^2) = \arg \min_{x^2} [z^2]^\top x^2 \\
x^{3*}(z^3) &:= \arg \min_{x^3} J^3(x^3, z^3) = \arg \min_{x^3} [z^3]^\top x^3.
\end{aligned} \tag{4.4}$$

The most important thing to notice is that the three optimization problems in (4.4) are decoupled. In other words, each agent i can compute its optimal response by knowing only z^i . Moreover, our definition of optimal response is based on the definition of aggregative games (i.e., it is not possible to express the optimal responses in terms of the cost function \tilde{J} because therein the aggregative structure is not made explicit). The main objective of this chapter is to propose algorithms that can be used to distributedly construct three vectors $\bar{z}^1, \bar{z}^2, \bar{z}^3 \in \mathbb{R}^n$ such that the joint solution of (4.4), in response to these vectors, is a Nash equilibrium of the original game. \square

In the rest of the chapter we make the following regularity assumptions on the cost functions, the constraint sets and the optimal responses, according to Definition 3.2.

Assumption 4.1.1 (Convexity of the cost function and constraints). (a) Each agent $i \in \mathbb{Z}[1, N]$ is subject to personalized convex and compact constraints $\mathcal{X}^i \subset \mathbb{R}^n$. There exists a convex and compact set $\bar{\mathcal{X}}$ such that $\mathcal{X}^i \subseteq \bar{\mathcal{X}}$ for all $i \in \mathbb{Z}[1, N]$. (b) The cost functions $\{J^i(x^i, z^i) : \mathcal{X}^i \times \bar{\mathcal{X}} \rightarrow \mathbb{R}\}_{i=1}^N$ are uniformly Lipschitz with constant \bar{L}_J .

Assumption 4.1.2 (Regularity of the optimal responses 1). (a) The mappings $\{x^{i*}\}_{i=1}^N$ in (4.2) are single valued and (b) uniformly Lipschitz with constant \bar{L}_x .

Assumption 4.1.3 (Regularity of the optimal responses 2). There exists a matrix $S = S^\top \in \mathbb{R}^{n \times n}$, $S \succ 0$, such that at least one of the following statements holds:

- (a) x^{i*} is a contraction in \mathcal{H}_S for all $i \in \mathbb{Z}[1, N]$;
- (b) x^{i*} is non-expansive in \mathcal{H}_S for all $i \in \mathbb{Z}[1, N]$;

(c) $x^{i\star}$ is firmly non-expansive in \mathcal{H}_S for all $i \in \mathbb{Z}[1, N]$;

(d) $x^{i\star}$ is anti-monotone in \mathcal{H}_S for all $i \in \mathbb{Z}[1, N]$.

In the following analysis we provide asymptotic guarantees in terms of the population size N . To this end, we need to define how the game in (4.1) scales when the population grows and we need to guarantee that, even though the population increases, the basic properties of each agent remain well defined. Mathematically, we do so by assuming that the set $\bar{\mathcal{X}}$ and the Lipschitz constants \bar{L}_J , \bar{L}_x , as defined in Assumptions 4.1.1 and 4.1.2, are uniformly bounded over N .

Assumption 4.1.4 (Asymptotic properties). *Consider a sequence of games $\mathcal{G}(N)$, of increasing population size N . (a) Denote by $\bar{\mathcal{X}}(N)$ the set $\bar{\mathcal{X}}$ as defined in Assumption 4.1.1 to stress its dependence on the population size N . There exists a convex and compact set $\mathcal{X} \subset \mathbb{R}^n$ such that $\bar{\mathcal{X}}(N) \subseteq \mathcal{X}$ for all population sizes N . (b) Denote by $\bar{L}_J(N)$ and $\bar{L}_x(N)$ the Lipschitz constants \bar{L}_J and \bar{L}_x defined in Assumptions 4.1.1 and 4.1.2 for a population of size N . There exist $L_J, L_x > 0$ such that $\bar{L}_J(N) \leq L_J$ and $\bar{L}_x(N) \leq L_x$ for all population sizes N .*

4.1.1 The quadratic case

A very important class of aggregative games is the one where the agents have a quadratic cost function

$$J^i(x^i, z^i) := q_i x^{i\top} Q x^i + 2(Cz^i + c^i)^\top x^i, \quad (4.5)$$

where $x^i, z^i \in \mathbb{R}^n$, $Q = Q^\top \succ 0$, $q_i > 0$, $C \in \mathbb{R}^{n \times n}$ and $c^i \in \mathbb{R}^n$. This type of cost function has been used for example in [HCM07, HCM12, BP13, GTL13, GPCL16]. In this case, the optimal response can be explicitly characterized as follows.

Lemma 4.1.1 (Optimal response for quadratic cost). *The unconstrained optimizer of the problem (4.2) with cost (4.5) is*

$$\hat{x}^{i\star}(z^i) := \arg \min_{x^i \in \mathbb{R}^n} J^i(x^i, z^i) = -(q_i Q)^{-1} (Cz^i + c^i); \quad (4.6)$$

the (constrained) optimizer in (4.2) reads as

$$x^{i\star}(z^i) = \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, z^i) = \Pi_{\mathcal{X}^i}^Q(\hat{x}^{i\star}(z^i)), \quad (4.7)$$

where $\Pi_{\mathcal{X}}^S(y) := \arg \min_{x \in \mathcal{X}} \|y - x\|_S$ denotes the projection on the set \mathcal{X} in the Hilbert space \mathcal{H}_S . \square

The next lemma provides sufficient conditions for Assumptions 4.1.1, 4.1.2, 4.1.3 and 4.1.4 to hold when the cost functions are as in (4.5) and the agents have convex and compact local constraint sets \mathcal{X}^i . For simplicity, let

$$M_i := \begin{bmatrix} q_i Q & -C \\ -C^\top & q_i Q \end{bmatrix}, \quad i \in \mathbb{Z}[1, N]. \quad (4.8)$$

Lemma 4.1.2 (Sufficient conditions for the convex quadratic case). *Under Assumption 4.1.1(a) the cost function in (4.5) satisfies Assumption 4.1.1(b) and the mappings $x^{i*}(z^i)$ with cost defined as in (4.5) satisfy Assumption 4.1.2. Moreover, the following statements hold.*

1.	$M_i \succ 0 \forall i$	\Rightarrow	Assumption 4.1.3.(a) holds with	$S = Q$
2.	$M_i \succcurlyeq 0 \forall i$	\Rightarrow	Assumption 4.1.3.(b) holds with	$S = Q$
3.	$-q_i Q \preccurlyeq C = C^\top \prec 0 \forall i$	\Rightarrow	Assumption 4.1.3.(c) holds with	$S = -C$
4.	$C = C^\top \succ 0$	\Rightarrow	Assumption 4.1.3.(d) holds with	$S = C$

Under Assumption 4.1.1(a) and Assumption 4.1.4.(a) if there exist $\underline{q}, \bar{q} > 0$ such that $\underline{q} \leq q_i \leq \bar{q}$ for all population sizes N and for all agents $i \in \mathbb{Z}[1, N]$, then Assumption 4.1.4.(b) holds. \square

Remark 4.1. We note that if $C = C^\top$ (e.g., as in the cases 3. and 4. of the previous lemma) then the AAG with cost function as in (4.5) is a potential game, according to Definition 2.2, with potential function

$$S(x) := x^\top \left(\begin{bmatrix} q_1 & & \\ & \ddots & \\ & & q_N \end{bmatrix} \otimes Q + \frac{1}{N} (\mathbb{1}_N \mathbb{1}_N^\top + I_N) \otimes C \right) x + 2[c^1; \dots; c^N]^\top x.$$

Consequently, if Assumption 4.1.1(a) holds, $q_i Q + \frac{2}{N} C \succ 0$ for all $i \in \mathbb{Z}[1, N]$ and $S(x)$ is convex, then the sequential BR dynamics in (2.5) are guaranteed to converge by Proposition 2.3.1. The same does not hold in general for the simultaneous BR dynamics in (2.6) (a quite restrictive sufficient condition is derived in Section 5.1.2). In the following section we exploit the aggregative structure of the game to suggest an alternative scheme, which uses again simultaneous updates, but is guaranteed to converge under different sufficient conditions.

4.2 Average aggregative games: schemes with central operator

We start by considering the case of average aggregative games, that is, games where the aggregator function is the same for all the players and coincides with the population average, $\sigma^i(x) = \bar{\sigma}(x) := \frac{1}{N} \sum_{j=1}^N x^j$ for all $i \in \mathbb{Z}[1, N]$.¹ For each agent $i \in \mathbb{Z}[1, N]$ the best response to the other agents strategies x^{-i} is therefore given by

$$x_{\text{br}}^i(x^{-i}) := \begin{cases} \arg \min_{x^i \in \mathbb{R}^n} J^i(x^i, \bar{\sigma}(x)) \\ \text{s.t. } x^i \in \mathcal{X}^i \end{cases} = \begin{cases} \arg \min_{x^i \in \mathbb{R}^n} J^i\left(x^i, \frac{1}{N} x^i + \frac{1}{N} \sum_{j \neq i} x^j\right) \\ \text{s.t. } x^i \in \mathcal{X}^i \end{cases} \quad (4.9)$$

¹The theory presented in this section can be generalized to the weighted average $\bar{\sigma}(x) := \frac{1}{N} \sum_{j=1}^N a^j x^j$, $\sum_{j=1}^N a^j = N$, under the assumption that the influence of each agent on the weighted average becomes negligible when the population size increases, [GPCL16].

The concept of Nash equilibrium can be specialized for AAGs as follows.

Definition 4.1 (Average aggregative (AA) Nash equilibrium). *Given N cost functions $\{J^i : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}\}_{i=1}^N$, a set of strategies $\{\bar{x}^i \in \mathcal{X}^i \subseteq \mathbb{R}^n\}_{i=1}^N$ is an AA ε -Nash equilibrium, with $\varepsilon > 0$, if for all $i \in \mathbb{Z}[1, N]$*

$$J^i(\bar{x}^i, \bar{\sigma}(\bar{x})) = J^i\left(\bar{x}^i, \frac{1}{N} \sum_{j=1}^N \bar{x}^j\right) \leq \min_{x^i \in \mathcal{X}^i} J^i\left(x^i, \frac{1}{N} \left(x^i + \sum_{j \neq i}^N \bar{x}^j\right)\right) + \varepsilon. \quad (4.10)$$

It is an AA Nash equilibrium if (4.10) holds with $\varepsilon = 0$. \square

4.2.1 Structural assumptions and communication requirements

According to the discussion in Section 4.1, we assume that the agents do not know the strategies of the other players and therefore cannot compute their best responses as given in (4.9). On the other hand, we assume that they react optimally to a common reference \bar{z} that is broadcast by a central operator (e.g., a price signal in the case of demand-response). In other words the agents update their strategies by using the optimal response mapping in (4.2) with $z^i = \bar{z}$ for all $i \in \mathbb{Z}[1, N]$. The main objective of the section is to derive schemes that can be used by the central operator to iteratively design the signal \bar{z} to be broadcast so that the set of optimal responses $\{x^{i*}(\bar{z})\}_{i=1}^N$ is an AA Nash equilibrium.

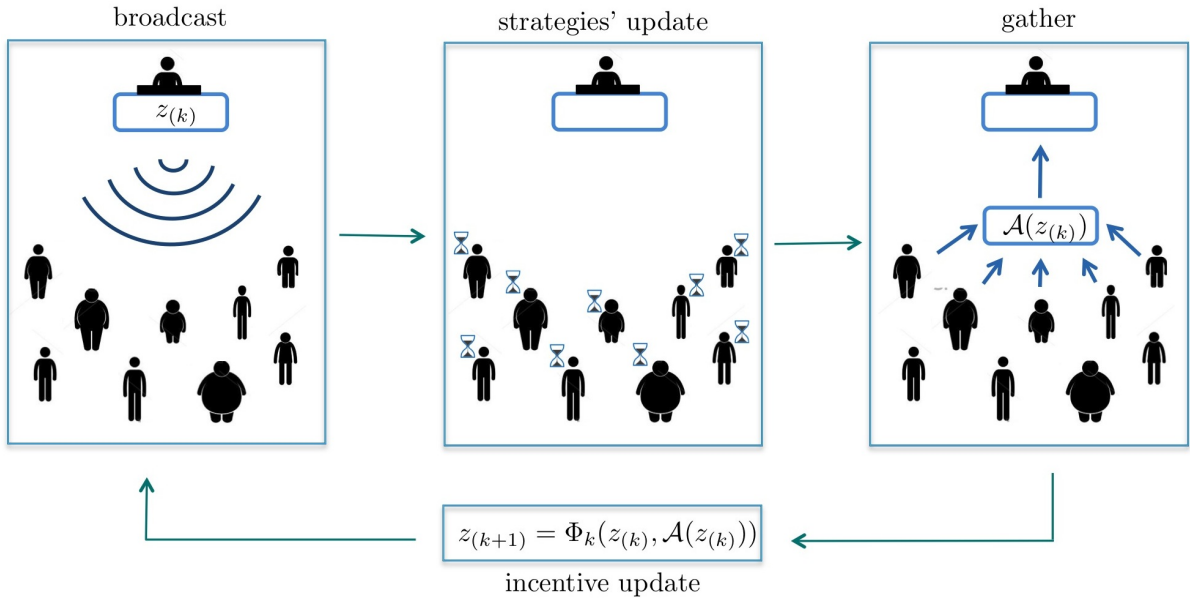


Figure 4.1: Coordination algorithm with central operator

Figure 4.1 illustrates the communication set up. Specifically, we assume that the agents and the central operator can iteratively communicate in a gather-and-broadcast

fashion, so that at every iteration k the central operator broadcasts a tentative reference signal $z_{(k)}$ to the population, the agents update their strategies accordingly and the aggregate of the optimal responses $\mathcal{A}(z_{(k)}) := \frac{1}{N} \sum_{i=1}^N x^{i\star}(z_{(k)})$ is gathered by the central operator. This information can then be used by the central operator to design the next reference $z_{(k+1)}$. Note that according to this scheme, at every iteration, the agents select as next strategy the one that minimizes their cost given the current reference $z_{(k)}$, thus complying with the myopic assumption.

Algorithm 2: AAG: Picard–Banach (decentralized)

Initialization: Set $k \leftarrow 0$. The central operator sets and broadcasts $z_{(0)} \in \mathbb{R}^n$.

Iterate until convergence:

Local: strategy update

$$x_{(k+1)}^i \leftarrow x^{i\star}(z_{(k)}) := \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, z_{(k)})$$

Central: average and reference update and broadcast

$$\begin{aligned} \mathcal{A}(z_{(k)}) &\leftarrow \frac{1}{N} \sum_{i=1}^N x_{(k+1)}^i \\ z_{(k+1)} &\leftarrow \mathcal{A}(z_{(k)}) \end{aligned}$$

The simplest possible scheme consistent with our structural assumptions is illustrated in Algorithm 2: at every iteration the central operator broadcasts as new reference $z_{(k+1)}$ the average $\mathcal{A}(z_{(k)}) := \frac{1}{N} \sum_{i=1}^N x^{i\star}(z_{(k)}) =: \bar{\sigma}_{(k+1)}$ that he has previously received from the population. We note that Algorithm 2 is similar in spirit to the simultaneous best response dynamics except for the fact that the agents update their strategies by taking the optimal response to the current reference $z_{(k)} = \bar{\sigma}_{(k)}$, considered as a fixed external signal, instead of the best response to the current average $\bar{\sigma}_{(k)}$. Consequently, in Algorithm 2 the agents do not take into account their contribution in the average when optimizing. This difference becomes negligible for large populations (see also Theorem 4.2.1). Conditions guarantying the convergence of Algorithm 2 to an AA Nash equilibrium are derived in Section 4.2.3 where a more general class of coordination algorithms, including Algorithm 2 as a special case, is proposed. To this end, we start by introducing the concept of aggregation mapping.

4.2.2 The aggregation mapping

We formalize the average population behavior obtained when all the agents react optimally to a reference z by defining the aggregation mapping $\mathcal{A} : \mathbb{R}^n \rightarrow \left(\frac{1}{N} \sum_{i=1}^N \mathcal{X}^i \right) \subset \bar{\mathcal{X}} \subset \mathbb{R}^n$ as

$$\mathcal{A}(z) := \frac{1}{N} \sum_{i=1}^N x^{i\star}(z). \tag{4.11}$$

Since the objective of our coordination problem is to find an **AA** Nash equilibrium for *large population* size, in the next theorem we exploit the Nash certainty equivalence principle or mean field approximation idea [HCM07, Section IV.A] to prove that when N is large the optimal responses to any fixed point of the aggregation mapping are an almost Nash equilibrium. The cornerstone to this result is the fact that the contribution of an individual strategy x^i to the average population behavior $\bar{\sigma}(x)$ becomes negligible when $N \rightarrow \infty$. Therefore, if $\bar{z} = \mathcal{A}(\bar{z}) = \frac{1}{N} \sum_{j=1}^N x^{j*}(\bar{z})$, then the optimal response $x^{i*}(\bar{z})$ is a good approximation of the best response $x_{\text{br}}^i(x^{-i*}(\bar{z}))$ of agent i to the strategies $\{x^{j*}(\bar{z})\}_{j \neq i}$ of the other players.

Theorem 4.2.1 (**AAG** Nash equilibrium). *Suppose Assumptions 4.1.1, 4.1.2(a) and 4.1.4 hold. For all $\varepsilon > 0$, there exists $\bar{N}_\varepsilon \in \mathbb{N}$ such that, for all $N \geq \bar{N}_\varepsilon$, if \bar{z} is a fixed point of \mathcal{A} in (4.11), that is, $\bar{z} = \frac{1}{N} \sum_{i=1}^N x^{i*}(\bar{z})$, then the set $\{x^{i*}(\bar{z})\}_{i=1}^N$, with x^{i*} as in (4.2), is an **AA** ε -Nash equilibrium. \square*

Remark 4.2. *It follows from the proof of Theorem 4.2.1, given in the Appendix, that the set $\{x^{i*}(\bar{z})\}_{i=1}^N$ is an **AA** ε_N -Nash equilibrium with $\varepsilon_N = \mathcal{O}(1/N)$. \square*

4.2.3 A class of decentralized coordination schemes

According to Theorem 4.2.1 the optimal responses to any fixed point of the aggregation mapping \mathcal{A} in (4.11) are an almost **AA** Nash equilibrium, for large population sizes. Consequently, if the central operator designs and broadcasts references $z_{(k)}$ that asymptotically converge to a fixed point, that is $z_{(k)} \rightarrow \bar{z} = \mathcal{A}(\bar{z})$, then the agents strategies converge to the desired configuration. Let us consider again Algorithm 2. By using the concept of aggregation mapping, a single iteration of this algorithm can be rewritten as

$$z_{(k+1)} = \mathcal{A}(z_{(k)}),$$

which is the Picard-Banach iteration relative to the aggregation mapping \mathcal{A} . From the theory revised in Section 3.2.2 it is then clear that a sufficient condition for Algorithm 2 to converge to a fixed point is that the mapping \mathcal{A} is either a **CON** or **FNE**. Motivated by this observation, we investigate the regularity properties of the aggregation mapping.

Proposition 4.2.2 (Regularity of the aggregation mapping). *Suppose that assumptions 4.1.1(a) and 4.1.2 hold. For all $i \in \mathbb{Z}[1, N]$, let x^{i*} be defined as in (4.2). The mapping \mathcal{A} in (4.11) is Lipschitz continuous, has a fixed point, and is*

-
1. **CON** in \mathcal{H}_S under assumption 4.1.3.(a)
 2. **NEX** in \mathcal{H}_S under assumption 4.1.3.(b)
 3. **FNE** in \mathcal{H}_S under assumption 4.1.3.(c)
 4. **SPC** in \mathcal{H}_S under assumption 4.1.3.(d)
-

\square

An immediate consequence of the previous proposition is that under Assumption 4.1.3.(a) or 4.1.3.(c), Algorithm 2 converges. This proposition however shows that, even if these assumptions are violated, different regularity conditions of \mathcal{A} can be guaranteed under different sufficient conditions. Leveraging on this fact and on the fixed point iterations described in Section 3.2.2 we propose an extension of Algorithm 2.

Algorithm 3: AAG: fixed point iterations (decentralized)

Initialization: Set $k \leftarrow 0$. The central operator sets the mappings Φ_k and broadcasts $z_{(0)} \in \mathbb{R}^n$.

Iteration:

Local: strategy update

$$x_{(k+1)}^i \leftarrow x^{i*}(z_{(k)}) := \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, z_{(k)})$$

Central: average and reference update and broadcast

$$\begin{aligned} \mathcal{A}(z_{(k)}) &\leftarrow \frac{1}{N} \sum_{i=1}^N x_{(k+1)}^i \\ z_{(k+1)} &\leftarrow \Phi_k(z_{(k)}, \mathcal{A}(z_{(k)})) \end{aligned}$$

The main difference between Algorithm 2 and the proposed extension, detailed in Algorithm 3, is that, in the latter, we allow the central operator to design the new signal $z_{(k+1)}$ by means of a (possibly iteration dependent) mapping $\Phi_k(\cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ that depends not only on the received aggregate $\mathcal{A}(z_{(k)})$ but also on the previous signal $z_{(k)}$, thus introducing a memory in the coordination scheme. A single iteration of Algorithm 3 results in

$$z_{(k+1)} = \Phi_k(z_{(k)}, \mathcal{A}(z_{(k)})).$$

We show in the next theorem that by choosing the mapping Φ_k appropriately, based on the theory of fixed point iterations, it is possible to guarantee convergence under less restrictive assumptions than those of Algorithm 2. Specifically, besides the already mentioned Picard–Banach iteration

$$\Phi^{\text{P-B}}(z_{(k)}, \mathcal{A}(z_{(k)})) := \mathcal{A}(z_{(k)}), \quad (4.12)$$

for which Algorithm 3 coincides with Algorithm 2, we suggest the use of the Krasnoselskij iteration

$$\Phi^K(z_{(k)}, \mathcal{A}(z_{(k)})) := (1 - \lambda)z_{(k)} + \lambda\mathcal{A}(z_{(k)}) \quad (4.13)$$

with $\lambda \in (0, 1)$, and the step-dependent Mann iteration

$$\Phi_k^{\text{M}}(z_{(k)}, \mathcal{A}(z_{(k)})) := (1 - \alpha_k)z_{(k)} + \alpha_k\mathcal{A}(z_{(k)}), \quad (4.14)$$

where the sequence $(\alpha_k)_{k=1}^{\infty}$ is such that $\alpha_k \in (0, 1) \forall k \geq 0$, $\lim_{k \rightarrow \infty} \alpha_k = 0$ and $\sum_{k=1}^{\infty} \alpha_k = \infty$ (e.g., $\alpha_k = 1/k$), as described in Section 3.2.2.

Theorem 4.2.3 (Decentralized convergence to a fixed point). *Under Assumptions 4.1.1(a) and 4.1.2, the following iterations and conditions guarantee global convergence of Algorithm 3 to a fixed point of \mathcal{A} in (4.11), where x^{i*} is as in (4.2) for all $i \in \mathbb{Z}[1, N]$:*

	Picard–Banach (4.12)	Krasnoselskij (4.13)	Mann (4.14)
1. Assumption 4.1.3.(a)	✓	✓	✓
2. Assumption 4.1.3.(b)		✓	✓
3. Assumption 4.1.3.(c)	✓	✓	✓
4. Assumption 4.1.3.(d)		✓*	✓

* with $\lambda < \frac{2}{1+L^2}$, where L is the Lipschitz constant of \mathcal{A} in \mathcal{H}_S .

The fixed point is unique in cases 1. and 4. □

Remark 4.3 (Rate of Convergence). *We emphasize that each mapping presented in Theorem 4.2.3 has its specific range of applicability. This allows us to select one or more fixed point iterations based on the specific regularity property at hand. Which mapping provides the best convergence performance is in general problem dependent [Ber07, Chapter 9]. We note however that under Assumption 4.1.3.(a), using Φ^{P-B} , and Assumption 4.1.3.(d), using Φ^K , Algorithm 3 exhibits geometric convergence (see Lemma 3.2.4 and 3.2.7). □*

An important feature of Theorem 4.2.3 is that decentralized convergence is guaranteed under conditions that are independent of the individual constraints $\{\mathcal{X}^i\}_{i=1}^N$. Therefore, our results and methods apply naturally to populations of heterogeneous agents.

Connection with other decentralized convergence results in aggregative games

Decentralized convergence to Nash equilibria in terms of fixed point iterations has been previously studied in the literature of aggregative game theory; mostly to show convergence of the sequential BR dynamics in (2.5) [Kuk04, Cournot path] [Jen10, Theorem 2], under the assumption that the BR mappings of the players are non-increasing [Jen10, Assumption 1], besides being continuous and compact valued. In large-scale games, however, simultaneous responses as in Algorithms 2 and 3 are computationally more convenient with respect to sequential ones. Within the literature of aggregative games, the Mann iteration in (4.24) has been proposed in [DHZ06, Remark 2] for the simultaneous BRs of the agents. See [Hei06] for an application to distributed power allocation and scheduling in congested distributed networks. The aggregative game setup in these papers considers the strategy of the players to be a 1-dimensional variable taking values in a compact interval of the real numbers. Convergence is then guaranteed if the BR mappings of the players are continuous, compact and non-increasing [DHZ06, conditions (i)–(iii), p. 81]. It actually follows from the proof of Proposition 4.2.2 that

Assumption 4.1.3.(d) implies that the aggregation mapping is anti-monotone, which is the n -dimensional generalization of the non-increasing property used in [DHZ06], for the case of large populations.

4.2.4 An application of AAGs: PEV coordination

As an application of AAGs, we investigate the problem of coordinating the charging of a large population of **plug-in electric vehicles (PEVs)**, introduced in [MCH13] and extended to the constrained case in [GTL13, PCGL14]. For each agent/PEV $i \in \mathbb{Z}[1, N]$, we consider the discrete-time linear dynamics $s_{t+1}^i = s_t^i + b^i u_t^i$, where $t \in \mathbb{N}$, $s_t^i \in [0, 1]$ is the state of charge, u_t^i is the charging control input and $b^i > 0$ represents the charging efficiency. The objective of each PEV i is to select a charging profile $u^i := [u_0^i, \dots, u_{T-1}^i]^\top$ that acquires a charge amount $\bar{\gamma}^i \in [0, 1]$, within a finite charging horizon $[0, T]$, $T \in \mathbb{N}$, and that minimizes the price that the agent has to pay for the electricity. Note that $s_T^i = s_0^i + b^i \sum_{t=0}^{T-1} u_t^i$, therefore the requirement $s_T^i - s_0^i = \bar{\gamma}^i$ is equivalent to the charging constraint $\sum_{t=0}^{T-1} u_t^i = \mathbb{1}_T^\top u^i = \frac{\bar{\gamma}^i}{b^i} =: \gamma^i$. Moreover, we consider a dynamic pricing scheme, where the price of the electricity depends on the overall demand, namely the inflexible demand $D \in \mathbb{R}_{\geq 0}^T$ (i.e., the demand coming from the non-PEVs) plus the aggregate PEV demand $N\bar{\sigma} \in \mathbb{R}_{\geq 0}^T$. In line with the (almost-affine) price function in [MCH13, PCGL14], we consider the affine price function $p(\bar{\sigma}_t) := 2a(\bar{\sigma}_t + d_t)$, where a is a positive constant and $d_t := \frac{D_t}{N}$ denotes the average inflexible demand² in the time interval t . A typical profile of average inflexible demand is illustrated in Figure 4.2 A). To sum up, each agent aims at minimizing the total electricity cost according to the following AAG

$$\mathcal{G} := \begin{cases} \min_{u^i \in \mathbb{R}^T} & p(\bar{\sigma}(u))^T u^i \\ \text{s.t.} & 0 \leq u^i \leq U^i, \mathbb{1}_T^\top u^i = \gamma_i, \end{cases} \quad (4.15)$$

where $p(\bar{\sigma}) := [p(\bar{\sigma}_0), \dots, p(\bar{\sigma}_{T-1})]^\top$ is the vector of prices over the interval $[0, T]$ and $U^i \in \mathbb{R}_{\geq 0}^T$ is a vector of desired upper bounds on the charging inputs.³

According to the previous theory, we assume that the agents do not have access to the charging profile of the other PEVs (so that they cannot compute the actual price $p(\bar{\sigma})$) but react to a reference $z = [z_0, \dots, z_{T-1}] \in \mathbb{R}^T$, broadcast by the central operator⁴, and set their electricity profile $u_{\text{LP}}^{i*}(z) \in \mathbb{R}^T$ for the whole interval $[0, T]$ by solving the

²We refer to [MCH13] for a discussion on why the average quantities $\bar{\sigma}_t(u_t) = \frac{1}{N} \sum_{j=1}^N u_t^j$ and $d_t = \frac{D_t}{N}$ are used instead of the total quantities $\sum_{j=1}^N u_t^j$ and D_t . We briefly note here that such assumption is justifiable if the grid infrastructure (e.g., the number of production plants) scales with the PEV population size N .

³More general convex constraints could be considered (e.g., to model multiple charging intervals, charging rates, vehicle-to-grid operations). We keep the same setting of [MCH13, PCGL14] for simplicity.

⁴The central operator can, without loss of generality, broadcast the price $p(z)$ instead of z .

following optimization problem

$$u_{\text{LP}}^{i*}(z) := \begin{cases} \arg \min_{u^i \in \mathbb{R}^T} & p(z)^\top u^i \\ \text{s.t.} & 0 \leq u^i \leq U^i, \mathbb{1}^\top u^i = \gamma_i, \end{cases} := \begin{cases} \arg \min_{u^i \in \mathbb{R}^T} & 2a(z+d)^\top u^i \\ \text{s.t.} & 0 \leq u^i \leq U^i, \mathbb{1}^\top u^i = \gamma_i. \end{cases} \quad (4.16)$$

The optimal response in (4.16) is the optimizer of a linear program, hence it can be set valued and thus violate Assumption 4.1.2. Following [MCH13, PCGL14], we therefore consider a quadratic relaxation of (4.16) and we define as optimal charging control u^{i*} for each PEV $i \in \mathbb{Z}[1, N]$, given the reference z , the solution of

$$u^{i*}(z) := \begin{cases} \arg \min_{u^i \in \mathbb{R}^T} & \delta \|u^i - z\|^2 + 2a(z+d)^\top u^i \\ \text{s.t.} & 0 \leq u^i \leq U^i, \mathbb{1}_T^\top u^i = \gamma_i, \end{cases} \quad (4.17)$$

where $\delta > 0$ should be chosen small to approximate the original linear cost. We refer to [PCGL14, Section V] for numerical evidence of the beneficial effect of choosing a small $\delta > 0$ for the perturbed cost in (4.17). We here report a theoretical result on the relation among the Nash equilibria of the game in (4.15) and the relaxed game whose optimal responses are as in (4.17)

$$\mathcal{G}_\delta := \begin{cases} \min_{u^i \in \mathbb{R}^T} & \delta \|u^i - \bar{\sigma}(u)\|^2 + 2a(\bar{\sigma}(u) + d)^\top u^i \\ \text{s.t.} & 0 \leq u^i \leq U^i, \mathbb{1}_T^\top u^i = \gamma_i. \end{cases} \quad (4.18)$$

Lemma 4.2.4 (Convergence of Nash equilibria for small δ). *For any $\delta < a$, the game \mathcal{G} in (4.15) and its quadratic relaxation \mathcal{G}_δ in (4.18), have a unique Nash equilibrium which we denote by \bar{u}_{LP} and \bar{u}_δ , respectively. Moreover, $\bar{u}_\delta \rightarrow \bar{u}_{\text{LP}}$, as $\delta \rightarrow 0$, and for any $\varepsilon > 0$ there exists $\bar{\delta} > 0$ such that, for any $\delta < \bar{\delta}$, \bar{u}_δ is an ε -Nash equilibrium for the original game (4.16).*

We assume that the central operator (e.g., the grid operator) aims at designing a reference $\bar{z} \in \mathbb{R}^T$ such that the set of optimal responses $\{u^{i*}(\bar{z})\}$ is a Nash equilibrium of the game \mathcal{G}_δ in (4.18). In view of Theorem 4.2.1, a solution to such coordination problem can be found by selecting as reference a fixed point of the mapping

$$\mathcal{A}(z) := \frac{1}{N} \sum_{i=1}^N u^{i*}(z)$$

which represents the average among the optimal charging control inputs $\{u^{i*}(z)\}_{i=1}^N$, for a given price $p(z)$. According to Theorem 4.2.3, such a fixed point can be found by applying Algorithm 3. Note that in the PEV case this algorithm can be implemented as follows. At each iteration k :

1. the grid operator broadcasts a tentative price $p(z_{(k)})$;
2. the PEVs solve (4.17) and compute what would be their optimal charging profile if the price was indeed $p(z_{(k)})$;

3. the aggregate charging profile is send back to the grid operator which utilizes it to update and broadcast a new tentative price $p(z_{(k+1)})$.

It is important to stress that these algorithmic iterations take place before the charging interval $[0, T]$ begins. Consequently, the charging profiles $u^{i*}(z_{(k)})$ are never implemented⁵. Once the algorithm reaches convergence, the grid operator sets the actual price to $p(\bar{z})$ and the agents implement $u^{i*}(\bar{z})$. Since the cost function in (4.18) can be rewritten up to constant terms as the quadratic cost function in (4.5) with $Q = I$, $C = (a - \delta)I$, $c^i = ad$, $q_i = \delta$ for all $i \in \mathbb{Z}[1, N]$, we can establish conditions on δ under which a specific fixed point iteration, that in this context represents a price update law, converges to an AA ε_N -Nash solution for the constrained charging control problem.

Corollary 4.2.5 (Decentralized PEV coordination). *The following iterations and conditions guarantee global convergence to an ε_N -Nash equilibrium of (4.18) with $\varepsilon_N = \mathcal{O}(\frac{1}{N})$.*

	Picard–Banach (4.12)	Krasnoselskij (4.13)	Mann (4.14)
1. $\delta > a/2$	✓	✓	✓
2. $\delta \geq a/2$		✓	✓
3. $a > \delta > 0$		✓*	✓

* with $\lambda < \frac{2}{1+L^2}$, where L is the Lipschitz constant of \mathcal{A} in \mathcal{H}_S . □

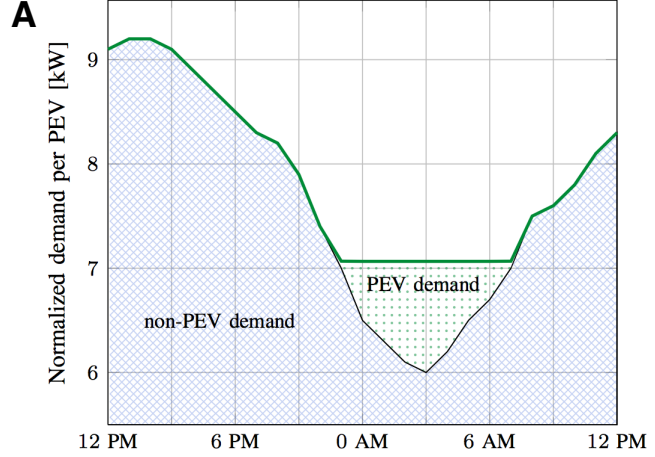
The use of the Picard–Banach iteration was firstly suggested in [MCH13], where the authors prove convergence of Algorithm 2 for $\delta > a/2$. For small values of δ , it was however shown both in [MCH13] and [PCGL14] that the Picard–Banach iteration causes permanent price oscillations. The previous corollary proves that, on the other hand, the Mann iteration converges for any value of $\delta > 0$, as empirically observed in [PCGL14], thus allowing one to approximate the original problem (4.15) to any desired precision.

Using the same numerical values as in [MCH13], Figure 4.2 shows that by using the Mann iteration, the central operator can successfully coordinate the PEV population to an AA ε_N -Nash equilibrium, even for $\delta < a/2$. Actually it appears from Figure 4.2 A) that the aggregate of the charging strategies at the obtained Nash equilibrium is valley-filling, that is, there exists $v > 0$ such that for each time t , $\frac{1}{N} \sum_{j=1}^N u_t^{j*}(\bar{z}) + d_t = \max\{d_t, v\}$. Such aggregate behavior is desirable in practice since it has no demand peaks [MCH13, Section I] and can be shown to be welfare-optimal for the original game \mathcal{G} , in the case when there are no charging upper bounds [MCH13, Lemma 3.1]. Figure 4.2 B) shows some charging trajectories of the AA ε_N -Nash equilibrium $\{u^{i*}(\bar{z})\}_{i=1}^N$ obtained with the Mann iteration. Figure 4.2 A) and B) illustrate the strategies at the Nash equilibrium, that is, the strategies obtained at convergence of Algorithm 3. Figure 4.2 C), on the other hand, is related to the iterations of the algorithm. Specifically, it shows that by using the Picard–Banach iteration to update the price Algorithm 3 would oscillate

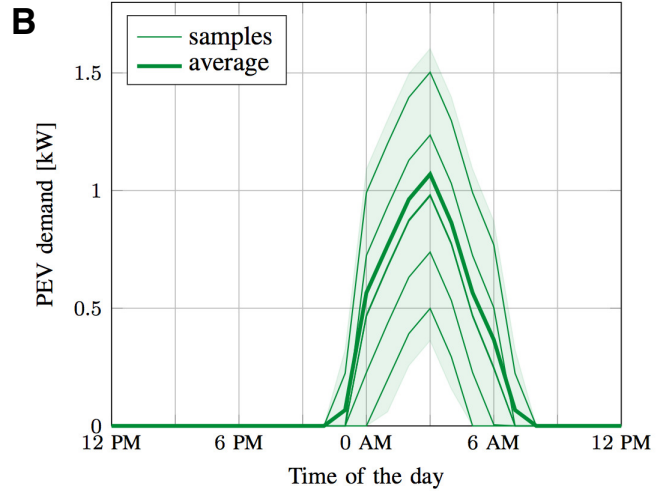
⁵We note here that since the charging profiles $u^{i*}(z_{(k)})$ at iteration k are not implemented, the agents have actually no incentive in using the optimal response as strategies update rule, but could instead cheat and send back to the central

Figure 4.2: Charging setting without upper bounds ($\delta = 10^{-4}$).

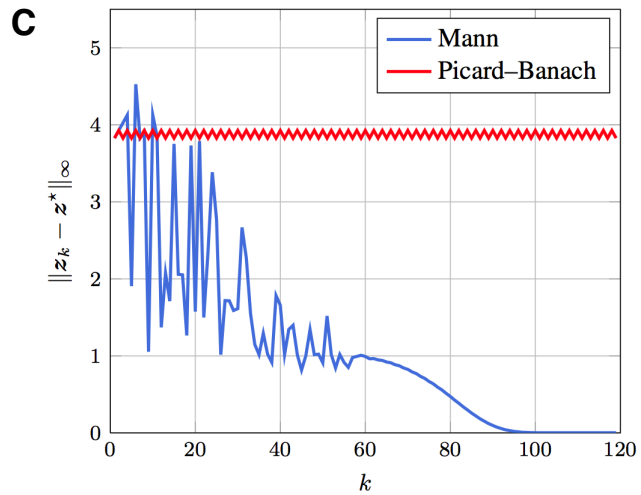
A) The green line represents the sum of the average non-PEV demand d added to the aggregate charging profile $\bar{\sigma}(u^*(\bar{z}))$ corresponding to the AA ε_N -Nash equilibrium $\{u^{i*}(\bar{z})\}_{i=1}^N$, obtained with Algorithm 3 using the Mann iteration. The obtained profile appears to be valley-filling.



B) Individual strategies at the ε_N -Nash equilibrium $\{u^{i*}(\bar{z})\}_{i=1}^N$. The area between the maximum and minimum charging strategies is shown in transparent green. The thin lines are some charging strategies sampled from the AA ε_N -Nash equilibrium; the thick line is the average of all the charging strategies.



C) Distance between the reference z_k , obtained by applying Algorithm 3 with the Picard–Banach iteration (red) or the Mann iteration (blue), as price update mappings, and the reference \bar{z} obtained at convergence of the algorithm with the Mann iteration. Note that the Picard–Banach iteration oscillates indefinitely.



indefinitely, while by using the Mann it reaches convergence. We refer to [PCGL14] for further discussions and numerical simulations.

4.3 Network aggregative games: distributed schemes

In this section we consider a generalization of the NAG setup introduced in Section 2.1.2 that allows for multiple rounds of communications. We briefly note here that communicating ν times over a graph with associated adjacency matrix P is mathematically equivalent to communicating once over a graph with associated adjacency matrix P^ν . In fact, if the agents have initial state $x := [x^1; \dots; x^N]$, after one round of communications the vector of neighbors aggregate states is $\sigma_1 := [\sigma_1^1; \dots; \sigma_1^N] = Px$, after two rounds it is $\sigma_2 := P\sigma_1 = P^2x$ and in general the vector of neighbors aggregate states after ν rounds of communications is $\sigma_\nu := P^\nu x$. Practically, however, communicating ν times over the network P is more convenient than communicating once over P^ν since in the latter case the agents would need to exchange information with a larger subset of the population (actually all the agents up to ν hops away in the graph P). In the following we denote by $[P]_{ij} = P_{ij}$ the element in position (i, j) of the matrix P and by $P_{ij}^\nu := [P^\nu]_{ij}$ the element in position (i, j) of the matrix P^ν .

To extend the NAG setup to multiple rounds of communications, we assume that each agent i aims at minimizing a cost function $J^i(x^i, \sigma_\nu^i)$ that depends on its own strategy x^i and on the aggregate quantity $\sigma_\nu^i(x) := \sum_{j=1}^N P_{ij}^\nu x^j$ which is obtained by communicating and averaging the strategies of its neighbors $\mathcal{N}^i := \{j \neq i \mid P_{ij} > 0\}$ a number ν of times, a priori fixed. In the following, we reserve special attention to the cases: $\nu = 1$ (typical of standard NAGs) and (ii) $\nu \rightarrow \infty$ (used in the next section to approximate AAGs). Formally, each agent $i \in \mathbb{Z}[1, N]$ aims at computing its *best response* to the neighbors' aggregate state σ_ν^i

$$\begin{aligned} x_{\text{br}}^i(x^{-i}) &:= \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, \sigma_\nu^i(x)) \\ &= \arg \min_{x^i \in \mathcal{X}^i} J^i\left(x^i, P_{ii}^\nu x^i + \sum_{j \neq i}^N P_{ij}^\nu x^j\right). \end{aligned} \quad (4.19)$$

The definition of Nash equilibrium can be specialized to NAGs, with ν communications, as follows.

Definition 4.2 (Network aggregative (NA) Nash equilibrium with ν communications). *Given N cost functions $\{J^i : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}\}_{i=1}^N$, a weighted adjacency matrix $P \in \mathbb{R}^{N \times N}$, a fixed number of communications $\nu \in \mathbb{N}$ and $\varepsilon > 0$, a set of strategies $\{\bar{x}^i \in \mathcal{X}^i\}_{i=1}^N$ is a NA ε -Nash equilibrium for (4.19) if, for all $i \in \mathbb{Z}[1, N]$, it holds*

$$J^i(\bar{x}^i, \sigma_\nu^i(\bar{x})) = J^i\left(\bar{x}^i, \sum_{j=1}^N P_{ij}^\nu \bar{x}^j\right) \leq \min_{x^i \in \mathcal{X}^i} J^i\left(x^i, P_{ii}^\nu x^i + \sum_{j \neq i}^N P_{ij}^\nu \bar{x}^j\right) + \varepsilon.$$

If the above inequality holds for $\varepsilon = 0$ then $\{\bar{x}^i\}_{i=1}^N$ is an NA Nash equilibrium.

4.3.1 Structural assumptions and communication requirements

We assume that the communication network satisfies the following assumption.

Assumption 4.3.1 (Graph property for NAGs). *For the a priori fixed number $\nu \in \mathbb{N}$ of communications and the given population size N , the weighted adjacency matrix $P \in \mathbb{R}^{N \times N}$ is row stochastic and satisfies $P_{ii}^\nu = 0$ for all $i \in \mathbb{Z}[1, N]$.*

Assumption 4.3.1 is equivalent to the absence of cycles of length ν in the graph associated with P . In the typical case of single-communication NAGs, that is if $\nu = 1$, Assumption 4.3.1 is equivalent to the absence of self-loops. Note that, according to this assumption, $x_{\text{br}}^i(x^{-i}) = x^{i*}(\sigma_\nu^i(x))$, since $\sigma_\nu^i(x)$ does not depend on x^i .

4.3.2 The extended aggregation mapping (ν fixed)

Consider a game with an a priori fixed number ν of communications. Let $\mathbf{z} := [z^1; \dots; z^N] \in \mathbb{R}^{Nn}$ be a vector of (possibly different) references for each agent and define the mapping $\mathbf{x}^* : \mathbb{R}^{Nn} \rightarrow \mathcal{X}_{1 \times N}$ as

$$\mathbf{x}^*(\mathbf{z}) := [x^{1*}(z^1); \dots; x^{N*}(z^N)] \in \mathbb{R}^{Nn}, \quad (4.20)$$

whose components are the optimal responses $x^{i*}(z^i)$ of each agent i to the signal z^i , as defined in (4.2). The mapping \mathbf{x}^* in (4.20) can be used to define an *extended aggregation mapping* \mathcal{A}_ν that, given a vector \mathbf{z} , returns the updated neighbors aggregate states, after one optimization and ν communication steps. Formally, $\mathcal{A}_\nu : \mathbb{R}^{Nn} \rightarrow (P^\nu \otimes I_n)\mathcal{X}_{1 \times N} \subset \mathbb{R}^{Nn}$ is defined as

$$\mathcal{A}_\nu(\mathbf{z}) := \begin{bmatrix} \mathcal{A}_\nu^1(\mathbf{z}) \\ \vdots \\ \mathcal{A}_\nu^N(\mathbf{z}) \end{bmatrix} := \begin{bmatrix} \sum_{j=1}^N P_{1j}^\nu x^{j*}(z^j) \\ \vdots \\ \sum_{j=1}^N P_{Nj}^\nu x^{j*}(z^j) \end{bmatrix} = (P^\nu \otimes I_n)\mathbf{x}^*(\mathbf{z}) =: \mathcal{P}_\nu \mathbf{x}^*(\mathbf{z}). \quad (4.21)$$

Note that, with respect to the aggregation mapping $\mathcal{A}(z)$ defined for AAGs in Section 4.2, the mapping $\mathcal{A}_\nu(\mathbf{z})$ takes values in \mathbb{R}^{Nn} instead of \mathbb{R}^n . In fact, in NAGs we have N local references $\{z^i\}_{i=1}^N$ instead of one reference z which is the same for all players. The next theorem shows that, nonetheless, the fixed points of the extended aggregation mapping \mathcal{A}_ν can be used to find a Nash equilibrium of the NAG game with ν communications, for any population size N .

Theorem 4.3.1 (NAG Nash equilibria). *Under Assumptions 4.1.1(a) and 4.1.2 the mapping \mathcal{A}_ν in (4.21) admits at least one fixed point $\bar{\mathbf{z}} = \mathcal{A}_\nu(\bar{\mathbf{z}})$. If also Assumption 4.3.1 holds and $\bar{\mathbf{z}}$ is a fixed point of \mathcal{A}_ν , then the set of strategies $\{x^{i*}(\bar{z}^i)\}_{i=1}^N$, with x^{i*} as in (4.2) for all $i \in \mathbb{Z}[1, N]$, is an NA Nash equilibrium for (4.19). \square*

Note that Theorem 4.3.1 implicitly ensures the existence of at least one Nash equilibrium for all NAGs satisfying Assumptions 4.1.1, 4.3.1 and 4.1.2 (consistently with Proposition 2.1.1). If one relaxes Assumption 4.3.1 to $P_{ii}' = \mathcal{O}(\frac{1}{N})$ for all $i \in \mathbb{Z}[1, N]$, then it is possible to show that the optimal responses to a fixed point are an ε_N -Nash equilibrium, with $\varepsilon = \mathcal{O}(\frac{1}{N})$, following the same procedure as in Theorem 4.2.1.

4.3.3 A class of distributed coordination schemes

Similarly to what was done in Section 4.2 for AAGs, we suggest here a class of update rules, summarized in Algorithm 4, that aims at reaching an NA Nash equilibrium by finding a fixed point of the extended aggregation mapping (thus exploiting Theorem 4.3.1). The main difference between Algorithm 3 and Algorithm 4 is that the former requires the presence of a central operator, that broadcasts the common reference $z^{(k)}$ to the whole population, while in the latter coordination is achieved in a totally distributed fashion, see Figure 4.3. Specifically, we assume that at the beginning of step k each agent i has a local reference vector $z_{(k)}^i$, which could for example be its local estimate of the neighbors aggregate state, and

1. communicates ν_1 times to compute a local a priori aggregate $\mathcal{E}_{\nu_1,0}^i(\mathbf{z}^{(k)}) := \sum_{j=1}^N P_{ij}^{\nu_1} z_{(k)}^j$, where $\mathbf{z}^{(k)} = [z_{(k)}^1; \dots; z_{(k)}^N]$;
2. computes the optimal response $x_{(k+1)}^i$ to $\mathcal{E}_{\nu_1,0}^i(\mathbf{z}^{(k)})$;
3. communicates ν_2 times to compute a local a posteriori aggregate $\mathcal{A}_{\nu_1,\nu_2}^i(\mathbf{z}^{(k)}) := \sum_{j=1}^N P_{ij}^{\nu_2} x_{(k+1)}^j$;
4. updates the local reference z^i by filtering the a posteriori aggregate $\mathcal{A}_{\nu_1,\nu_2}^i(\mathbf{z}^{(k)})$ with the previous local reference $z_{(k)}^i$ through a mapping $\Phi_k : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$.

Note that the two subscripts ν_1, ν_2 in \mathcal{A}^i refer to the number of communications executed before (ν_1) and after (ν_2) the optimization step 2. Performing communications both before and after optimizing allows us to guarantee convergence under more general conditions than those typically assumed in the literature (see Theorem 4.3.2 in the next subsection).

Algorithm 4 describes a class of learning dynamics, parameterized by ν_1, ν_2 and by the mapping Φ_k . The simplest dynamics is obtained by setting $(\nu_1, \nu_2) = (0, \nu)$ and by using as mapping the Picard–Banach iteration, yielding the update

$$\mathbf{z}_{(k+1)} = \Phi^{\text{P-B}}(\mathbf{z}_{(k)}, \mathcal{A}_{\nu_1,\nu_2}(\mathbf{z}_{(k)})) := \mathcal{A}_{\nu_1,\nu_2}(\mathbf{z}_{(k)}). \quad (4.22)$$

Under these settings and under Assumption 4.3.1, Algorithm 4 coincides with the *simultaneous best response dynamics* in (2.6) for the NAG in (4.19). As mentioned in Section 2.3, these are the learning dynamics that are obtained when, at every iteration,

Algorithm 4: NAG: fixed point iterations (distributed)

Initialization. Set $k \leftarrow 0$, choose an initial reference $z_{(1)}^i \in \mathbb{R}^n$ for every agent i , the mappings $\{\Phi_k\}_{k=1}^\infty$ and $\nu_1, \nu_2 \in \mathbb{Z}_{\geq 0}$, such that $\nu_1 + \nu_2 = \nu$.

Iterate until convergence. Each agent i :

1) Communicates ν_1 times

$$\begin{aligned} \mathcal{E}_{0,0}^i &\leftarrow z_{(k)}^i \\ \text{for } s = 0 \text{ to } s = \nu_1 - 1 \text{ do} \\ &\mathcal{E}_{s+1,0}^i \leftarrow \sum_{j=1}^N P_{ij} \mathcal{E}_{s,0}^j \\ \text{end} \end{aligned}$$

2) Computes its optimal strategy with respect to $\mathcal{E}_{\nu_1,0}^i$

$$x_{(k+1)}^i \leftarrow x^{i*}(\mathcal{E}_{\nu_1,0}^i) := \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, \mathcal{E}_{\nu_1,0}^i)$$

3) Communicates ν_2 times

$$\begin{aligned} \mathcal{A}_{\nu_1,0}^i &\leftarrow x_{(k+1)}^i \\ \text{for } s = 0 \text{ to } s = \nu_2 - 1 \text{ do} \\ &\mathcal{A}_{\nu_1,s+1}^i \leftarrow \sum_{j=1}^N P_{ij} \mathcal{A}_{\nu_1,s}^j \\ \text{end} \end{aligned}$$

4) Updates the reference

$$\begin{aligned} z_{(k+1)}^i &\leftarrow \Phi_k(z_{(k)}^i, \mathcal{A}_{\nu_1,\nu_2}^i) \\ k &\leftarrow k + 1 \end{aligned}$$

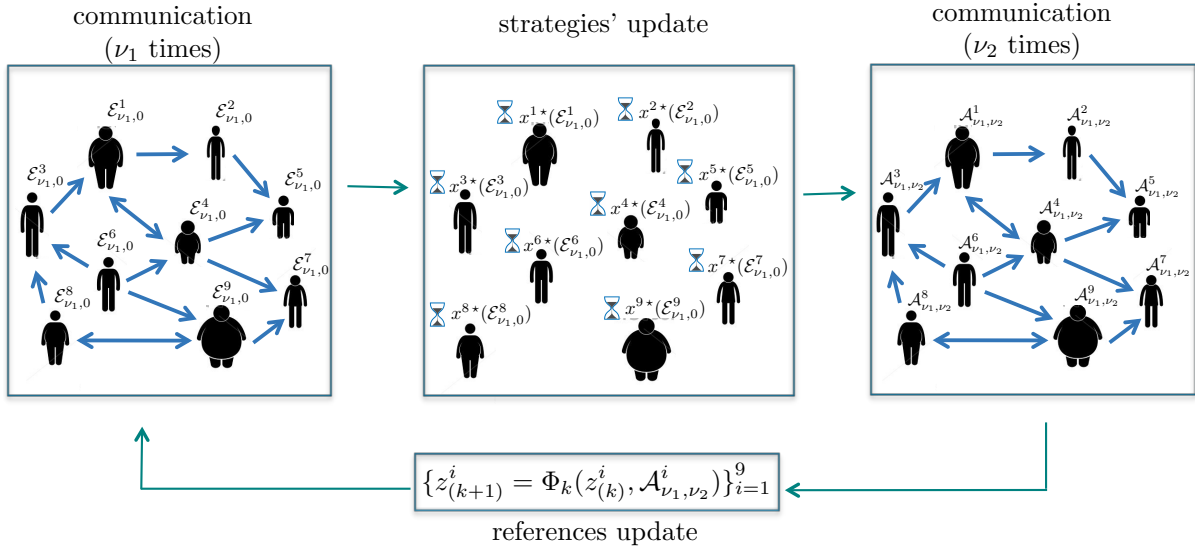


Figure 4.3: Coordination algorithm with local communications only

all the players synchronously compute their BRs to the current strategies of the neighbors. Conditions under which the BR dynamics are guaranteed to converge are derived in Theorem 4.3.2 and Lemma 5.1.3. To ensure convergence under weaker assumptions, we consider again the use of more general fixed point iterations as the Krasnoselskij iteration

$$\Phi^K(\mathbf{z}^{(k)}, \mathcal{A}_{\nu_1, \nu_2}(\mathbf{z}^{(k)})) := (1 - \lambda)\mathbf{z}^{(k)} + \lambda\mathcal{A}_{\nu_1, \nu_2}(\mathbf{z}^{(k)}) \quad (4.23)$$

with $\lambda \in (0, 1)$, and the step-dependent Mann iteration

$$\Phi_k^M(\mathbf{z}^{(k)}, \mathcal{A}_{\nu_1, \nu_2}(\mathbf{z}^{(k)})) := (1 - \alpha_k)\mathbf{z}^{(k)} + \alpha_k\mathcal{A}_{\nu_1, \nu_2}(\mathbf{z}^{(k)}), \quad (4.24)$$

where the sequence $(\alpha_k)_{k=1}^\infty$ is such that $\alpha_k \in (0, 1) \forall k \geq 0$, $\lim_{k \rightarrow \infty} \alpha_k = 0$ and $\sum_{k=1}^\infty \alpha_k = \infty$ (e.g., $\alpha_k = 1/k$), see Section 3.2.2. Note that each agent can compute the i^{th} component of these mappings knowing only the i^{th} component of its arguments, therefore the update of $\mathbf{z}^{(k)}$ can be performed distributedly (as in point 4 of Algorithm 4).

Global convergence of Algorithm 4

The following result provides conditions on the cost functions and on the network structure under which the sequence of vectors used to compute the optimal responses,

$$([\mathcal{E}_{\nu_1, 0}^1(\mathbf{z}^{(k)}); \dots; \mathcal{E}_{\nu_1, 0}^N(\mathbf{z}^{(k)})]_{k=1}^\infty), \quad (4.25)$$

in Algorithm 4 converges, as k tends to infinity, to a fixed point of the extended aggregation mapping \mathcal{A}_ν in (4.21). As a consequence, $\{x_{(k)}^i\}_{i=1}^N$ converges to the desired Nash equilibrium configuration, according to Theorem 4.3.1.

Theorem 4.3.2 (Distributed convergence to a fixed point). *Under Assumptions 4.1.1(a) and 4.1.2, the following iterations and conditions guarantee that the sequence in (4.25) converges, for any initial condition $\mathbf{z}_0 \in \mathbb{R}^{Nn}$, to a fixed point of \mathcal{A}_ν in (4.21)*

	regularity	network	(ν_1, ν_2)	Φ^{P-B} in (4.22)	Φ^K in (4.23)	Φ_k^M in (4.24)
1.	A.4.1.3.(a)	$\ P\ \leq 1$	$(0, \nu)$	✓	✓	✓
2.	A.4.1.3.(b)	$\ P\ \leq 1$	$(0, \nu)$		✓	✓
3.	A.4.1.3.(c)	$P = P^\top$	$(\nu/2, \nu/2)$	✓	✓	✓
4.	A.4.1.3.(d)	$P = P^\top$	$(\nu/2, \nu/2)$		✓*	✓

* = with $\lambda < \frac{2}{1+L^2}$ where L is the Lipschitz constant of $\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}$ in \mathcal{H}_S .

The mapping \mathcal{A}_ν has a unique fixed point in case 1. □

We note that the condition $\|P\| \leq 1$ is satisfied by any double stochastic matrix (thanks to Hölder's inequality). In particular, it is therefore satisfied if the network is symmetric (i.e., $P = P^\top$).

Remark 4.4 (Rate of Convergence). *Under A.4.1.3.(a), using Φ^{P-B} , and A.4.1.3.(d), using Φ^K , Algorithm 4 exhibits geometric convergence (see Lemma 3.2.4 and 3.2.7). □*

4.3.4 An application of NAGs: multidimensional opinion dynamics

We consider here the problem of modeling how ideas, innovations or behaviors spread in a social network [Lor07, MS10, BHT10, MB12, PPTF15]. Opinion dynamics have been formulated as games in the literature in [GS14] and [EB15]. Here we generalize these models to the case of constrained opinion sets. Specifically, we assume that each agent $i \in \mathbb{Z}[1, N]$ has a vector $x^i \in [0, 1]^n$ of opinions regarding $n \in \mathbb{N}$ topics. Each component $x_s^i \in [0, 1]$ represents the opinion of agent i about topic $s \in \mathbb{Z}[1, n]$, where 0 represents an extremely negative and 1 an extremely positive opinion. We denote by $x_{(0)}^i \in [0, 1]^n$ the initial opinion of agent i and by $\theta^i > 0$ its stubbornness regarding its initial opinion. To describe the opinion dynamics, we consider a synchronous repetitive game where at every iteration k each agent i communicates once ($\nu = 1$) with its neighbors $\mathcal{N}^i := \{j \in \mathbb{Z}[1, N] \mid P_{ij} > 0\}$ and updates its opinion according to the optimization problem

$$\begin{aligned} x_{\text{br}}^i(x^{-i}) &:= \arg \min_{x^i \in \mathcal{X}^i} \sum_{j \neq i}^N (P_{ij} \|x^i - x^j\|^2) + \theta_i \|x^i - x_{(0)}^i\|^2 \\ \text{s.t.} \quad &x^i \in \mathcal{X}^i. \end{aligned} \tag{4.26}$$

The cost function in (4.27) comprises two terms: the first one models the influence of the neighbors to the new opinion of agent i , the second one models the “stubbornness” of agent i about its initial opinion. Additional constraints on the agents’ opinions across the n topics, as for example the fact that the opinions regarding two topics should not differ more than a given threshold, or hard constraints on single topics can be encoded via the constraint set $\mathcal{X}^i \subseteq [0, 1]^n$. The agents are assumed to be heterogeneous in the sense that the stubbornness parameter $\theta_i \geq 0$, the constraint set \mathcal{X}^i and the weights $\{P_{ij}\}_{j \neq i}^N$ may be different for every agent. We refer to agents for which $\theta_i = 0$ and $\mathcal{X}^i = [0, 1]^n$ as *followers* and to all the remaining ones as *stubborn*. We note that (4.27) can be equivalently reformulated as

$$\begin{aligned} x_{\text{br}}^i(x^{-i}) &:= \arg \min_{x^i \in \mathcal{X}^i} \|x^i\|^2 - 2 \left(\sum_{j \neq i}^N P_{ij} x^j \right)^\top x^i + \theta_i \|x^i - x_{(0)}^i\|^2 \\ \text{s.t.} \quad &x^i \in \mathcal{X}^i. \end{aligned}$$

thus it is a network aggregative game with $\nu = 1$. Furthermore, since $P_{ii} = 0$, Assumption 4.3.1 is met and $x_{\text{br}}^i(x^{-i}) = x^{i*}(\sigma_1^i(x))$.

In the absence of constraints, the solution to (4.27) for each topic decouples, hence one can consider $n = 1$ without loss of generality, and the BR can be computed explicitly, leading to the simultaneous BR dynamics

$$x_{(k+1)}^i = x_{\text{br}}^i(x_{(k)}^{-i}) = \frac{1}{1+\theta_i} \sum_{j \neq i} P_{ij} x_{(k)}^j + \frac{\theta_i}{1+\theta_i} x_{(0)}^i,$$

which is a particular case of the standard Friedkin and Johnsen model [FJ99], with parameters $\Lambda := \text{diag}\left(\frac{1}{1+\theta_1}, \dots, \frac{1}{1+\theta_N}\right)$ and $W := P$. For the case in which there is at least one stubborn agent it is possible to show, with the same argument used in [GS14], that the simultaneous BR dynamics converge to a Nash equilibrium of the game in (4.27). If on the other hand, all the agents are followers, one recovers the standard DeGroot model [DeG74], whose convergence properties have been exhaustively investigated using consensus theory. In this case, the Nash equilibria coincide with the right eigenvectors of the matrix P corresponding to the eigenvalue $\lambda = 1$. Theorems 4.3.1 and 4.3.2 allow us to extend the analysis to the *multi-dimensional* case with *stubborn agents and generic convex constraints*.

Corollary 4.3.3 (Distributed opinion dynamics). *Suppose that Assumption 4.1.1(a) holds. The following iterations and conditions guarantee convergence of Algorithm 4, from any initial configuration, to a Nash equilibrium for (4.27).*

	<i>cost</i> ($\forall i$)	<i>network</i>	(ν_1, ν_2)	Φ^{P-B} in (4.22)	Φ^K in (4.23)
1.	$\theta_i > 0$	$\ P\ \leq 1$	$(0, 1)$	✓	✓
2.		$\ P\ \leq 1$	$(0, 1)$		✓

□

In words, if $\theta_i > 0$ for all $i \in \mathbb{Z}[1, N]$ and $\|P\| \leq 1$, the simultaneous BR dynamics (i.e., Algorithm 4 for $(\nu_1, \nu_2) = (0, 1)$ by using Φ^{P-B}) converge to the unique Nash equilibrium. If some *followers* are present in the population, then a Nash equilibrium can be reached using an update rule with memory (i.e., Algorithm 4 for $(\nu_1, \nu_2) = (0, 1)$ by using Φ^K). We also note that in the first case (i.e. when $\theta_i > 0$ for all $i \in \mathbb{Z}[1, N]$), the rate of convergence can be explicitly characterised. In fact, as noted in Remark 4.4, the simultaneous BR dynamics converge with a geometric rate that depends on the contraction constant of the extended aggregation mapping. It is easy to see from the proof of Corollary 4.3.3 that such constant is $\delta := \frac{1}{1+\min\{\theta_i\}_{i=1}^N} < 1$. In other words $\|x_{(k)} - \bar{x}\|_2 \leq \delta^k \|x_{(0)} - \bar{x}\|_2$, where \bar{x} is the unique Nash equilibrium. Let us define the time of convergence as $T(\epsilon) := \min\{\tau \mid \frac{\|x_{(k)} - \bar{x}\|_\infty}{\|x_{(0)} - \bar{x}\|_\infty} \leq \epsilon, \forall k \geq \tau, \forall x_{(0)} \in \mathcal{X}\}$. Then we get

$$\frac{\|x_{(k)} - \bar{x}\|_\infty}{\|x_{(0)} - \bar{x}\|_\infty} \leq Nn \frac{\|x_{(k)} - \bar{x}\|_2}{\|x_{(0)} - \bar{x}\|_2} \leq Nn\delta^k \quad (4.27)$$

hence $\frac{\|x_{(k)} - \bar{x}\|_\infty}{\|x_{(0)} - \bar{x}\|_\infty} \leq \epsilon$ if $Nn\delta^k \leq \epsilon$ or equivalently if $k \geq \frac{\log(\epsilon) - \log(N) - \log(n)}{\log(\delta)}$. Consequently, $T(\epsilon) \leq \frac{\log(\epsilon) - \log(N) - \log(n)}{\log(\delta)}$. If δ is uniformly upper bounded on N (i.e. θ_i are uniformly lower bounded), then the convergence time increases at most with rate $\mathcal{O}(\log(N))$. We note that the upper bound derived with this approach does not depend on the network topology (as long as $\|P\| \leq 1$). The investigation of tighter upper bounds that exploit specific network properties (e.g. the essential spectral radius that appears in the

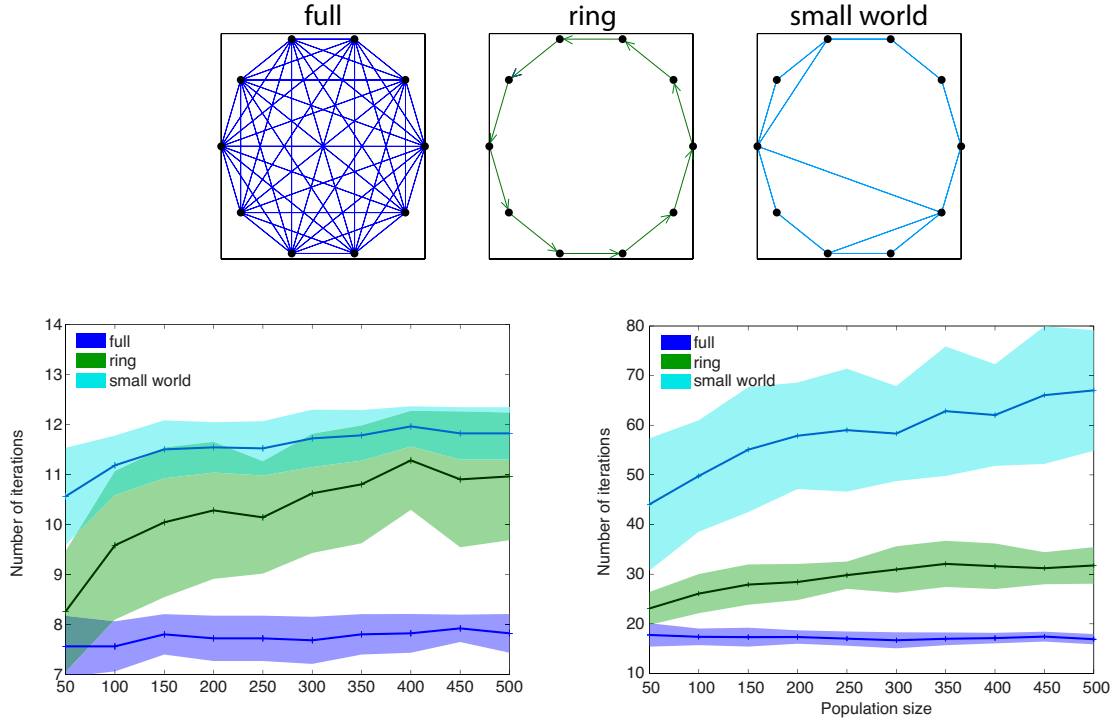


Figure 4.4: Average number of iterations (solid line) and 90% confidence intervals as functions of the population size N for the three different network topologies (top). The plot on the left refers to a population where only stubborn agents are present and update their opinion using the BR scheme ($\Phi^{\text{P-B}}$), while the plot on the right refers to a population composed by half stubborn and half follower agents, using the scheme with memory (Φ^{K}). In each case, 50 different networks and populations were simulated with initial opinions chosen according to a uniform distribution in $[0, 1]^2$. The stopping criterion is $\|\mathbf{z}^{(k)} - \mathbf{z}^{(k-1)}\|_{\infty} \leq 10^{-5}$.

convergence rate of opinion dynamics without stubborn agents [OT09]) is left as future work.

To numerically investigate the performance of the two schemes we consider a case study where each agent i has $n = 2$ opinions $x^i = [x_1^i, x_2^i]^{\top}$, regarding two different topics, taking values in $\mathcal{X}^i := \{[x_1, x_2]^{\top} \mid \|x_1 - x_2\|^2 \leq 0.3\}$ and is either a follower or stubborn with $\theta_i = 1$. Figure 4.4 reports the number of iterations required to reach convergence as a function of the population size N for two different compositions of the population and three different network topologies. These simulations show that the convergence speed depends only mildly on the population size in the case when all the agents are stubborn, as expected, but also in the mixed population case, suggesting that our approach is scalable.

4.4 Average aggregative games: distributed schemes

AAGs can be seen as NAGs over a complete network, so that, $\sigma_\nu^i(x) = \bar{\sigma}(x) = \frac{1}{N} \sum_{j=1}^N x^j \in \mathbb{R}^n$ for all i . Consequently, Algorithm 4 can readily be applied, with $P = \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top$ and $\nu = 1$, to distributedly steer the population to a Nash equilibrium of the AAG.⁶ The drawback is that, to this end, each agent needs to communicate with all the other agents in the population, which is not feasible in large-population games. If there is a central operator, that can compute and broadcast the average, then one can apply the coordination algorithm already suggested in Section 4.2. Unfortunately, solutions relying on a central operator may not always be implementable.

The main result of this section is to show how Algorithm 4 can additionally be used to coordinate the agents strategies to an almost Nash equilibrium in large-scale AAGs distributedly, that is by means of local communications over a sparse network $P \neq \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top$, by performing a large enough number ν of communications in between two strategies updates.

4.4.1 Structural assumptions and communication requirements

To apply Algorithm 4 to AAGs we need to ensure that $\lim_{\nu \rightarrow \infty} \sigma_\nu^i(x) = \bar{\sigma}(x) = \frac{1}{N} \sum_{j=1}^N x^j$. In other words, the network P should be such that, by iteratively communicating, the agents asymptotically reach consensus on the population average. To this aim we introduce the following assumption.

Assumption 4.4.1 (Asymptotic average consensus). *For all population sizes N , the weighted adjacency matrix P , which is here denoted by $P(N)$ to stress its dependence on the population size N , satisfies $\lim_{\nu \rightarrow \infty} P^\nu(N) = \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top$. Equivalently, $P(N)$ is primitive and doubly stochastic [OSAFM07].*

4.4.2 The extended aggregation mapping ($\nu \rightarrow \infty$)

We start by analyzing the relation between the extended aggregation mapping \mathcal{A}_ν in (4.21), obtained via ν communications over the local network P , and the extended aggregation mapping $\mathcal{A} : \mathbb{R}^{Nn} \rightarrow (\frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top \otimes I_n) \mathcal{X}_{1 \times N}$,

$$\mathcal{A}(z) := (\frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top \otimes I_n) \mathbf{x}^*(z) =: \mathcal{I} \mathbf{x}^*(z), \quad (4.28)$$

that arises from a complete network in AAGs.

⁶Note that rigorously speaking AAGs do not satisfy Assumption 4.3.1, however one can consider the equivalent cost functions $\hat{J}^i(x^i, \sigma^{-i}(x)) = J^i(x^i, \sigma(x))$, where $\sigma^{-i}(x) = \frac{1}{N} \sum_{j \neq i} x^j$ and then study the regularity properties of the game with costs \hat{J}^i and matrix $P = \frac{1}{N} (\mathbb{1}_N \mathbb{1}_N^\top - I_N)$.

Lemma 4.4.1 (Convergence of aggregation mappings). *Under Assumptions 4.1.1 and 4.1.2, the mappings \mathcal{A} in (4.28) and \mathcal{A}_ν in (4.21), for all $\nu \in \mathbb{N}$, are continuous and have at least one fixed point. If additionally Assumption 4.4.1 holds, then*

$$\lim_{\nu \rightarrow \infty} \sup_{z \in \mathbb{R}^{Nn}} \|\mathcal{A}_\nu(z) - \mathcal{A}(z)\| = 0. \quad \square$$

This lemma allows us to prove that the optimal responses to the fixed point of \mathcal{A}_ν in (4.21), which can be computed distributedly via Algorithm 4, are an almost Nash equilibrium for the AAG.

Theorem 4.4.2 (AAG Nash equilibria). *Suppose that Assumptions 4.1.1, 4.1.2, 4.1.4 and 4.4.1 hold. For all $\varepsilon > 0$ there exists \bar{N} such that: for all $N > \bar{N}$, there exists $\bar{\nu} > 0$ such that, for all $\nu \geq \bar{\nu}$, if \bar{z} is a fixed point of \mathcal{A}_ν in (4.21), then the set $\{x^{i*}(\bar{z}^i)\}_{i=1}^N$, with x^{i*} as in (4.2), is an ε -Nash equilibrium for the AAG in (4.9).* \square

We note that, contrary to the case of NAGs, where the fixed points of \mathcal{A}_ν lead, under Assumption 4.3.1, to a Nash equilibrium for finite ν and finite N , in the case of AAGs a Nash equilibrium is recovered only asymptotically. This is due to the fact that: i) the agents are required to “almost reach consensus” on the population average (i.e., a number ν large enough of consensus steps should be performed) and ii) the population average depends on the strategy of agent i with contribution proportional to $1/N$ (hence, to be negligible, N should be large enough). However, for any desired $\varepsilon > 0$, the proof of Theorem 4.4.2 allows one to derive lower bounds on N and ν ensuring that $\{x^{i*}(\bar{z}^i)\}_{i=1}^N$ is an ε -Nash equilibrium. We note that the minimum number of required communications $\bar{\nu}$ can be computed in a distributed fashion [PGGL15b, Appendix C]. For symmetric networks this dependence can be further specified in terms of the spectral properties of P .

Assumption 4.4.2 (Spectral properties). *For all population sizes N the weighted adjacency matrix $P = P(N)$ is symmetric. Moreover, there exists $\mu \in [0, 1)$ such that $\mu_N := \max_{\lambda \in \Lambda(P(N)) \setminus \{1\}} \{|\lambda|\} \leq \mu$ for all N .*

Corollary 4.4.3 (ε -convergence rate). *Suppose that Assumptions 4.1.1, 4.1.2, 4.1.4, 4.4.1 and 4.4.2 hold. If \bar{z} is a fixed point of \mathcal{A}_ν in (4.21), then the set of strategies $\{x^{i*}(\bar{z}^i)\}_{i=1}^N$, with x^{i*} as in (4.2) for all $i \in \mathbb{Z}[1, N]$, is an ε -Nash equilibrium for (4.9), with $\varepsilon = \mathcal{O}(\frac{1}{N} + \sqrt{N}\mu^\nu)$.* \square

In other words, Assumption 4.4.2 allows us to derive an upper bound on ε that is composed by two terms: one that decreases linearly in N , as for the solution with central coordinator (Section 4.2), and one that, for any fixed N , decreases exponentially fast with the number of communication steps ν . Assumption 4.4.2 is satisfied, e.g., by the degree-normalized adjacency matrices of any family of d -regular undirected ϵ -expander graphs [HLW06, Definition 2.2 and Example 2.2]. In fact $\mu_N \leq 1 - \frac{\epsilon^2}{2d^2} =: \mu$ [HLW06, Theorem 2.4], [Sin92, Theorem 2].

4.4.3 An application of AAGs: hierarchical demand-response

Consider a population of N loads whose electricity consumption $u^i = [u_1^i, \dots, u_T^i] \in \mathbb{R}^T$ over the horizon $\mathcal{T} = \mathbb{Z}[0, T - 1]$ is scheduled according to the demand-response scheme

$$\begin{aligned} u^{i*}(\bar{\sigma}) &:= \arg \min_{u^i \in \mathbb{R}^T} \sum_{t \in \mathcal{T}} \left(\rho_i \|u_t^i - \hat{u}_t^i\|^2 + p(\bar{\sigma}_t) u_t^i \right) \\ \text{s.t.} \quad & s_{t+1}^i = a^i s_t^i + \gamma^i u_t^i \quad \forall t \in \mathcal{T} \\ & [s^i(u^i); u^i] \in (\mathcal{S}^i \times \mathcal{B}^i) \cap \mathcal{R}^i, \end{aligned} \quad (4.29)$$

where $s_t^i = s_t^i(u^i)$ is the state of the load at time t (e.g., its temperature in case of heating ventilation air conditioning systems [MHS14] and thermostatically controlled loads [GGPL15] or its state of charge in case of plug-in electric vehicles [MCH13, PCGL14]), $s_0^i \in \mathbb{R}$ is the given initial state, $a^i, \gamma^i \in \mathbb{R} \setminus \{0\}$ are parameters modeling the dynamics and the efficiency of load i , $\bar{\sigma}_t = \frac{1}{N} \sum_{i=1}^N u_t^i \in \mathbb{R}$ is the population average energy demand at time t , $\bar{\sigma} := [\bar{\sigma}_0; \dots; \bar{\sigma}_{T-1}] \in \mathbb{R}^T$ and $\hat{u}^i \in \mathbb{R}^T$ is the preferred schedule of agent i . The energy consumption u^i and state vector $s^i(u^i)$ are constrained by the personalized sets $\mathcal{B}^i \subset \mathbb{R}^T$ and $\mathcal{S}^i \subset \mathbb{R}^T$, respectively, and by the joint constraint set $\mathcal{R}^i \subset \mathbb{R}^{2T}$. The first term in the cost function of (4.29) models the curtailment cost that each agent encounters for deviating from its nominal energy schedule \hat{u}_t^i , according to the Taguchi loss function [TCW04], $\rho_i > 0$ being a constant weighting parameter. The second term models the demand-response mechanism: the energy price varies according to a price function $p(\bar{\sigma}_t)$ that is assumed to be an affine increasing function of the total energy demand as in [MHS14, Equation (15)] and Section 4.2.4,

$$p(\bar{\sigma}_t) := \lambda \bar{\sigma}_t + p_0, \quad \lambda > 0. \quad (4.30)$$

Approaches to find a Nash equilibrium iteratively via a central coordinator have been proposed in [MCH13, PCGL14, GPCL16, MHS14]. Here we note that (4.29) can be rewritten as the AAG in (4.9) with $x^i = u^i$ and cost function as in (4.5) with $q_i = \rho_i$, $Q = I_T, C = \frac{\lambda}{2} I_T, c = \frac{p_0}{2} - \theta \hat{u}^i$. Note that, since $s^i(u^i)$ is an affine function of the input u^i , the constraint $[s^i(u^i); u^i] \in (\mathcal{S}^i \times \mathcal{B}^i) \cap \mathcal{R}^i$ can be rewritten as a unique constraint on the input, $u^i \in \mathcal{U}^i$, that is convex and compact if $\mathcal{B}^i, \mathcal{S}^i, \mathcal{R}^i$ are convex and compact.

Corollary 4.4.4 (Distributed demand-response). *Suppose that Assumptions 4.1.1, 4.1.4 and 4.4.1 hold, and let $p(\bar{\sigma}_t)$ be as in (4.30). The following iterations and conditions guarantee convergence of the strategies in Algorithm 4, from any initial point, to an $\varepsilon_{(N,\nu)}$ -Nash equilibrium for (4.29) with $\varepsilon_{(N,\nu)} = \mathcal{O}(\|P^\nu - \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top\|_\infty + \frac{1}{N})$.*

	cost ($\forall i$)	network	(ν_1, ν_2)	Φ^{P-B} in (4.22)	Φ^K in (4.23)	Φ_k^M in (4.24)
1.	$\rho_i > \lambda/2$	$\ P\ \leq 1$	$(0, \nu)$	✓	✓	✓
2.	$\rho_i \geq \lambda/2$	$\ P\ \leq 1$	$(0, \nu)$		✓	✓
3.		$P = P^\top$	$(\nu/2, \nu/2)$		✓*	✓

* = with $\lambda < \frac{2}{1+L^2}$ where L is the Lipschitz constant of $\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}$ in \mathcal{H}_S . □

The model given in (4.29) can be used for example to describe demand-response methods for heating ventilation air conditioning (HVAC) systems in smart buildings, as suggested in [MHS14], by selecting $\rho_i = \theta\gamma_i^2$, where $\theta > 0$ is the cost coefficient of the Taguchi loss function and $\gamma_i > 0$ specifies the thermal characteristic of the HVAC system. In [MHS14, Theorems 1, 2] it is shown that, for $N > 3$, if $\gamma^i = \gamma > 0$ for all i , $\mathcal{U}^i = [u_{\min}^i, u_{\max}^i] \subset \mathbb{R}^T$ with $u_{\min}^i, u_{\max}^i \in \mathbb{R}^T$, and $\lambda \leq \frac{2\theta\gamma^2}{N-3}$, then the Nash equilibrium is unique and can be computed using a control algorithm involving a central coordinator. Corollary 4.4.4 proves that, Algorithm 4, on the other hand, can be used to find an ε -Nash equilibrium in a distributed fashion, under arbitrary convex constraints $u^i \in \mathcal{U}^i$ and for all values of λ .

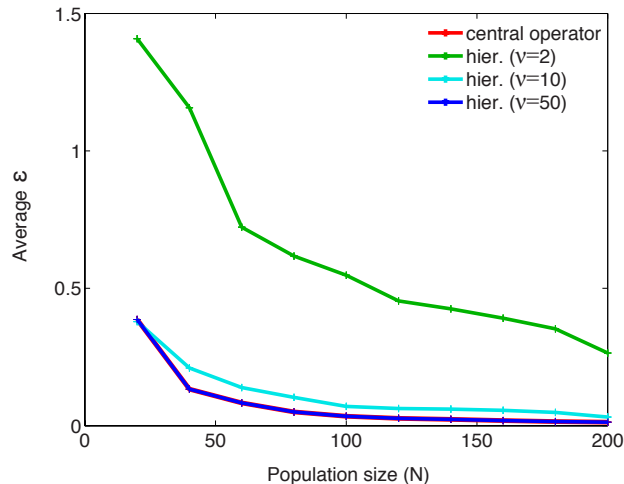


Figure 4.5: Comparison between the approach with central operator and the distributed hierarchical approach with $M = 5$ and $\nu = 2, 10, 50$. The matrix P_M is the adjacency matrix corresponding to a symmetric ring, so that $[P_M]_{ij} = \frac{1}{2}$ if $j = i + 1$ or $j = i - 1$ and $[P_M]_{ij} = 0$ otherwise, with the convention $N + 1 = -1 = 1$. The plot shows the average cost improvement $\varepsilon_N := \max_{i \in \{1, \dots, N\}} \bar{J}^i - J^{i*}$, where $\bar{J}^i := J\left(\bar{u}^i, \frac{1}{N} \sum_{j=1}^N \bar{u}^j\right)$ and $J^{i*} := \min_{u^i \in \mathcal{X}^i} J\left(u^i, \frac{1}{N} \left(u^i + \sum_{j \neq i}^N \bar{u}^j\right)\right)$, that an agent can achieve by unilateral deviations from the set of strategies $\{\bar{u}^j\}_{j=1}^N$ obtained at convergence of the algorithm. We set $\lambda = 2$ and $\theta\gamma_i^2 = 0.1$ for all i . We consider an horizon of $T = 24$ hrs and we assume a baseline energy consumption σ_0 as illustrated in [MCH13, Figure 1]. We set the baseline energy price to $p_0 := \lambda\sigma_0$. The average is computed over 10 different populations with prescheduled energy consumption \hat{u}^i uniformly sampled in $[0, 1]^T$ and constraint set $\mathcal{U}^i := \{u \in \mathbb{R}_{\geq 0}^T \mid \sum_{t=1}^T u_t^i = \sum_{t=1}^T \hat{u}_t^i, u_t^i = 0 \text{ if } t \notin [T_{\text{start}}^i, T_{\text{end}}^i]\}$, with T_{start}^i uniformly sampled in $\{1, T\}$, T_{end}^i uniformly sampled in $\{T_{\text{start}}^i + 1, T\}$. To guarantee convergence, we use Φ_k^M and $\nu_1 = \nu_2 = \frac{\nu}{2}$. The stopping criterion is $\|\mathcal{A}_{\nu}(\mathcal{A}_{\nu_1, 0}(z^{(k)})) - \mathcal{A}_{\nu_1, 0}(z^{(k)})\|_{\infty} \leq 10^{-3}$.

As a particular case of network we consider a hierarchical communication structure that models the fact that groups of buildings are managed by the same company. For simplicity, let us assume that there are M companies and each one manages B buildings, for a total of $N = MB$ buildings. At every communication step the managers compute the aggregate power demand of their buildings, then communicate among each other using a network $P_M \in \mathbb{R}^{M \times M}$ and finally broadcast the price signal to their buildings. It is easy to prove that the corresponding network $P = P_M \otimes \frac{1}{B} \mathbb{1}_B \mathbb{1}_B^\top$ satisfies Assumption 4.4.1, $\|P\|_2 \leq 1$ and $P = P^\top$ if and only if P_M does. In Figure 4.5 we report the values of ε obtained at convergence using Algorithm 4 with the mapping Φ^M and different number of communications, and we compare it with the performance of the centrally coordinated control scheme described in Section 4.2, Algorithm 3. We empirically verify that ε decreases when ν increases, as guaranteed by Theorem 4.4.2.

We conclude this section by noting that the model given in (4.29) can be also used, by setting ρ_i sufficiently small, to compute the optimal charging strategy for large populations of plug-in electrical vehicles (Section 4.2.4). Under the same conditions of Corollary 4.4.4 (point 3.), Algorithm 4 allows one to recover an ε -Nash equilibrium using a symmetric network P instead of the central coordinator.

4.5 Appendix

4.5.1 Proofs of the results stated in Section 4.1

Proof of Lemma 4.1.1 (Optimal response for quadratic cost)

The closed-form expression of the (unique) unconstrained optimizer $\hat{x}^{i\star}(z)$ directly follows from the equation $0 = \frac{\partial}{\partial x^i} J^i(x^i, z^i) = 2q_i Q x^i + 2(Cz^i + c^i)$. Then the following equalities hold:

$$\begin{aligned} \Pi_{\mathcal{X}^i}^Q(\hat{x}^{i\star}(z)) &= \arg \min_{x^i \in \mathcal{X}^i} \|x^i - \hat{x}^{i\star}(z)\|_Q^2 \\ &= \arg \min_{x^i \in \mathcal{X}^i} [x^i + (q_i Q)^{-1} (Cz^i + c^i)]^\top Q \cdot [x^i + (q_i Q)^{-1} (Cz^i + c^i)] \\ &= \arg \min_{x^i \in \mathcal{X}^i} x^{i\top} Q x^i + 2q_i^{-1} x^{i\top} (Cz^i + c^i) = \arg \min_{x^i \in \mathcal{X}^i} x^{i\top} q_i Q x^i + 2(Cz^i + c^i)^\top x^i \\ &= x^{i\star}(z^i). \end{aligned}$$

Proof of Lemma 4.1.2 (Sufficient conditions for the convex quadratic case)

The cost function $J^i(x^i, z^i)$ is \mathcal{C}^1 and is defined in the compact set $\mathcal{X}^i \times \bar{\mathcal{X}}$, therefore it is Lipschitz. One can prove that $L_{J^i} \leq 2(q_i \|Q\| \bar{D} + \|C\| \bar{D} + \|c\|) \leq 2(\max_i \{q_i\} \|Q\| \bar{D} + \|C\| \bar{D} + \|c\|) =: \bar{L}_J$, where $\bar{D} := \max_{x \in \bar{\mathcal{X}}} \|x\|$. Hence Assumption 4.1.1(b) holds. The mapping $\hat{x}^{i\star}$ in (4.6) is a single valued affine mapping. Consequently, it is Lipschitz

with constant $L_{x^i} = \frac{1}{q_i} \|Q^{-1}C\|_Q \leq \frac{1}{\min_i \{q_i\}} \|Q^{-1}C\|_Q =: \bar{L}_x$ in \mathcal{H}_Q . The projection operator $\Pi_{\mathcal{X}^i}^Q$ has Lipschitz constant 1 in \mathcal{H}_Q [BC10, Proposition 4.8], therefore the compositions $x^{i*}(\cdot) = \Pi_{\mathcal{X}^i}^Q(\hat{x}^{i*}(\cdot))$ in (4.7) are uniformly Lipschitz with constant \bar{L}_x and Assumption 4.1.2 holds.

Statement 1. and 2. It follows from Lemma 3.1.4 that the unconstrained optimizer \hat{x}^{i*} in (4.6) is a **CON** in \mathcal{H}_Q if and only if $-(q_i Q)^{-1}C)^\top Q(-(q_i Q)^{-1}C) - Q \prec 0$ or equivalently $C^\top (q_i Q)^{-1}C - q_i Q \prec 0$. By Shur complement this is equivalent to $M_i > 0$. The proof that \hat{x}^{i*} in (4.6) is **NEX** in \mathcal{H}_Q if and only if $M_i \succeq 0$ is analogous. Since $\Pi_{\mathcal{X}^i}^Q$ is **FNE** [BC10, Proposition 4.8] and hence **NEX** in \mathcal{H}_Q it follows that the composition $x^{i*}(\cdot) = \Pi_{\mathcal{X}^i}^Q(\hat{x}^{i*}(\cdot))$ is a **CON** in \mathcal{H}_Q if \hat{x}^{i*} is a **CON** in \mathcal{H}_Q , **NEX** in \mathcal{H}_Q if \hat{x}^{i*} is **NEX** in \mathcal{H}_Q .

Statement 3. and 4. The projection operator $\Pi_{\mathcal{X}^i}^Q$ is **FNE** in \mathcal{H}_Q and $x^{i*}(\cdot) = \Pi_{\mathcal{X}^i}^Q(\hat{x}^{i*}(\cdot))$ therefore, by Lemma 3.1.1, for all v, w we have

$$\begin{aligned} 0 \leq \|x^{i*}(v) - x^{i*}(w)\|_Q^2 &\leq (x^{i*}(v) - x^{i*}(w))^\top Q(\hat{x}^{i*}(v) - \hat{x}^{i*}(w)) \\ &= -(x^{i*}(v) - x^{i*}(w))^\top Q(q_i Q)^{-1}C(v - w) \\ &= -\frac{1}{q_i} (x^{i*}(v) - x^{i*}(w))^\top C(v - w). \end{aligned}$$

And consequently

$$0 \leq \|x^{i*}(v) - x^{i*}(w)\|_{q_i Q}^2 \leq -(x^{i*}(v) - x^{i*}(w))^\top C(v - w) \quad (4.31)$$

If $C = C^\top$ and $0 \prec -C \preceq q_i Q$, then $\|x^{i*}(v) - x^{i*}(w)\|_{-C}^2 \leq \|x^{i*}(v) - x^{i*}(w)\|_{q_i Q}^2$ for all $v, w \in \mathbb{R}^n$. Therefore, it follows from (4.31) that

$$\|x^{i*}(v) - x^{i*}(w)\|_{-C}^2 \leq (x^{i*}(v) - x^{i*}(w))^\top (-C)(v - w)$$

for all $v, w \in \mathbb{R}^n$, which is equivalent to x^{i*} being **FNE** in \mathcal{H}_{-C} by Lemma 3.1.1.

On the other hand, from (4.31) we get

$$0 \leq (x^{i*}(w) - x^{i*}(v))^\top C(v - w)$$

for all v, w , which for $C = C^\top \succ 0$ is equivalent to $-x^{i*}(\cdot)$ being **MON** in \mathcal{H}_C by Definition 3.2.

Finally, note that if $\underline{q} \leq q_i \leq \bar{q}$ for all population size N and for all agents $i \in \mathbb{Z}[1, N]$, then $\bar{L}_J(N) = 2(\max_i \{q_i\} \|Q\| \bar{D} + \|C\| \bar{D} + \|c\|) \leq 2(\bar{q} \|Q\| D + \|C\| D + \|c\|) =: L_J$ and $\bar{L}_x(N) = \frac{1}{\min_i \{q_i\}} \|Q^{-1}C\|_Q \leq \frac{1}{\underline{q}} \|Q^{-1}C\|_Q =: L_x$, where $D = \max_{x \in \mathcal{X}} \|x\|$. Therefore, Assumption 4.1.4.(b) holds. ■

4.5.2 Proofs of the results stated in Section 4.2

Proof of Theorem 4.2.1 (AAG Nash equilibrium)

By Assumptions 4.1.1 and 4.1.4 there exist a Lipschitz constant $L_J > 0$, such that for all population size N and all $i \in \mathbb{Z}[1, N]$, $|J^i(v, z_1) - J^i(w, z_2)| \leq L_J (\|v - w\| + \|z_1 - z_2\|)$ for all $v, w \in \mathcal{X}^i$, $z_1, z_2 \in \bar{\mathcal{X}}$ (see Lemma 3.1.3). Let us also define $D := \max_{x \in \mathcal{X}} \|x\|$, where \mathcal{X} is as defined in Assumption 4.1.4.

We now consider an arbitrary fixed point $\bar{z} = \frac{1}{N} \sum_{j=1}^N x^{j\star}(\bar{z}) =: \frac{1}{N} \sum_{j=1}^N \bar{x}^j$ of the aggregation mapping \mathcal{A} in (4.11). We show that an arbitrary agent i can improve its cost only by an amount $\varepsilon = \varepsilon_N = \mathcal{O}(1/N)$ if we fix the strategies $\{\bar{x}^j := x^{j\star}(\bar{z})\}_{j \neq i}^N$ of all other agents. By definition of optimal response, since $\bar{x}^i := x^{i\star}(\bar{z})$, we have

$$J^i(\bar{x}^i, \bar{z}) \leq J^i(x^i, \bar{z}) \quad \text{for all } x^i \in \mathcal{X}^i.$$

By substituting the equivalence $\bar{z} = \frac{1}{N} \sum_{j=1}^N \bar{x}^j$ we get that, for all $x^i \in \mathcal{X}^i$,

$$\begin{aligned} J^i(\bar{x}^i, \frac{1}{N} \sum_{j=1}^N \bar{x}^j) &\leq J^i(x^i, \frac{1}{N} \sum_{j=1}^N \bar{x}^j) \\ &= J^i(x^i, \frac{1}{N}(x^i + \sum_{j \neq i}^N \bar{x}^j)) + \left[J^i(x^i, \frac{1}{N} \sum_{j=1}^N \bar{x}^j) - J^i(x^i, \frac{1}{N}(x^i + \sum_{j \neq i}^N \bar{x}^j)) \right] \\ &\leq J^i(x^i, \frac{1}{N}(x^i + \sum_{j \neq i}^N \bar{x}^j)) + L_J \left\| \frac{1}{N} \sum_{j=1}^N \bar{x}^j - \frac{1}{N}(x^i + \sum_{j \neq i}^N \bar{x}^j) \right\| \\ &= J^i(x^i, \frac{1}{N}(x^i + \sum_{j \neq i}^N \bar{x}^j)) + \frac{L_J}{N} \|\bar{x}^i - x^i\| \leq J^i(x^i, \frac{1}{N}(x^i + \sum_{j \neq i}^N \bar{x}^j)) + \frac{2L_J D}{N}. \end{aligned}$$

Since this holds for all $i \in \mathbb{Z}[1, N]$, the set $\{\bar{x}^i\}_{i=1}^N$ is an ε_N -Nash equilibrium with $\varepsilon_N := \frac{2L_J D}{N}$. Finally, for all $\varepsilon > 0$ there exists $\bar{N}_\varepsilon := \frac{2L_J D}{\varepsilon}$ such that the cost of any agent i at a fixed point \bar{z} is ε -close to its optimal cost, for all population sizes $N \geq \bar{N}_\varepsilon$. ■

Proof of Proposition 4.2.2 (Regularity of the aggregation mapping)

The mapping \mathcal{A} in (4.11) is a convex combination of the mappings $\{x^{i\star}\}_{i=1}^N$, that are uniformly Lipschitz by Assumption 4.1.2, therefore \mathcal{A} is Lipschitz continuous as well. Moreover, \mathcal{A} is compact valued on the convex set $\frac{1}{N} \sum_{i=1}^N \mathcal{X}^i$, thus it has at least one fixed point by Lemma 3.2.1. The first three statements follow from the fact that the convex combination of CON/NEX/FNE mappings is CON/NEX/FNE itself. For the fourth statement, note that the convex combination of MON mappings is MON. Therefore $-\mathcal{A}(z) = \frac{1}{N} \sum_{i=1}^N -x^{i\star}(z)$ is MON. By Lemma 3.1.2, we then get that $\text{Id} - \mathcal{A}$ is SMON and consequently \mathcal{A} is SPC by Definition 3.2 with $\varepsilon = 1$. ■

Proof of Theorem 4.2.3 (Decentralized convergence to a fixed point)

Remember that an iteration of the algorithm is described by $z_{(k+1)} = \Phi_k(z_{(k)}, \mathcal{A}(z_{(k)}))$. From Proposition 4.2.2, if Assumption 4.1.3.(a) holds, then \mathcal{A} in (4.11) is a CON and if

Assumption 4.1.3.(c) holds, then \mathcal{A} is **FNE**. In both cases by Lemma 3.2.4 and 3.2.5, the Picard–Banach iteration converges to a fixed point of \mathcal{A} , which is unique if \mathcal{A} is a **CON**. For the other two fixed point iterations, we need to consider \mathcal{A} in (4.11) as a mapping from a compact convex set to itself. This can be assumed without loss of generality (that is, up to discarding the initial condition $z_{(0)}$) since \mathcal{A} takes values in $\frac{1}{N} \sum_{i=1}^N \mathcal{X}^i$, which is a linear transformation of the compact convex sets $\{\mathcal{X}^i\}_{i=1}^N$, as hence compact and convex as well [Roc70, Section 3, Theorem 3.1].

If Assumption 4.1.3.(a) or 4.1.3.(b) or 4.1.3.(c) holds then \mathcal{A} is **NEX** from Proposition 4.2.2 (see also Figure 3.1). If Assumption 4.1.3.(d) holds then \mathcal{A} is **SPC** and Lipschitz. In both cases, the Krasnoselskij iteration converges to a fixed point of \mathcal{A} by Lemma 3.2.6 and 3.2.7 and the Mann iteration converges by Lemma 3.2.9 and 3.2.8. ■

Proof of Lemma 4.2.4 (Convergence of Nash equilibria for small δ)

To prove this lemma we use an equivalent characterization of Nash equilibria in terms of variational inequalities as discussed in the following Section 5.1 and Proposition 5.1.1. Specifically, \bar{u}_δ is a Nash equilibrium if and only if it solves $\text{VI}(\mathcal{U}_{1 \times N}, f_\delta)$, where $\mathcal{U}_{1 \times N} = \mathcal{U}^1 \times \dots \times \mathcal{U}^N$, $\mathcal{U}^i := \{u^i \in \mathbb{R}^T \mid 0 \leq u^i \leq U^i, \mathbb{1}_T^\top u^i = \gamma^i\}$, $f_\delta(u) = F_\delta u + \tilde{f}$,

$$F_\delta = \frac{2}{N} [(a + N\delta - \delta)I_N + (a - \delta + \delta/N)\mathbb{1}_N \mathbb{1}_N^\top] \otimes I_n, \quad \tilde{f} = 2a \begin{bmatrix} d \\ \vdots \\ d \end{bmatrix},$$

and \bar{u}_{LP} is a Nash equilibrium if and only if it solves the $\text{VI}(\mathcal{U}_{1 \times N}, f)$, where $f(u) = Fu + \tilde{f}$ with $F := F_{[\delta=0]}$. For $\delta < a$, one can show that both F and F_δ are positive definite matrices. Therefore, by Lemma 3.1.4, the operators $f(u)$ and $f_\delta(u)$ are **SMON** (i.e., $F_\delta, F \succ 0$) and the VIs above have a unique solution (Proposition 3.3.7), proving the uniqueness of the Nash equilibria. The convergence result follows from [Mos69, Theorem A (b)] leveraging on the fact that the two VIs are defined over the same convex and compact set $\mathcal{U}_{1 \times N}$, the mappings $f_\delta(u)$ are monotone, continuous and uniformly bounded in $\mathcal{U}_{1 \times N}$ for $\delta \in [0, a)$ and the operator $f_\delta(u) \rightarrow f(u)$ uniformly on $\mathcal{U}_{1 \times N}$ as $\delta \rightarrow 0$. The fact that \bar{u}_δ is a ε -Nash equilibrium for the original game (4.15) for $\delta < \bar{\delta}$ follows from

$$J_{\text{LP}}(\bar{u}_\delta^i, \bar{u}_\delta^{-i}) \leq J_\delta(\bar{u}_\delta^i, \bar{u}_\delta^{-i}) \leq J_\delta(u^i, \bar{u}_\delta^{-i}) \leq J_{\text{LP}}(u^i, \bar{u}_\delta^{-i}) + \delta D, \quad \forall u^i \in \mathcal{U}^i$$

where $D := \max\{\|u^i - z\|^2 \mid u^i \in \mathcal{U}^i, z \in \frac{1}{N} \sum_{j=1}^N \mathcal{U}^j, i \in \mathbb{Z}[1, N]\}$ by imposing $\bar{\delta} < \frac{\varepsilon}{D}$. ■

Proof of Corollary 4.2.5 (Decentralized PEV coordination)

Assumptions 4.1.1, 4.1.2, 4.1.4 are met by the problem definition and Lemma 4.1.2. We consider the matrix inequality (4.8) in Lemma 4.1.2 with $Q = I$, $q_i = \delta$ and $C = (a - \delta)I$, $\delta, a > 0$ and note that $M_i = M$ for all $i \in \mathbb{Z}[1, N]$. The condition

$$M = \begin{bmatrix} \delta I & (\delta - a)I \\ (\delta - a)I & \delta I \end{bmatrix} \succ 0,$$

is equivalent, by Schur complement [BV04, Section A.5.5], to $\delta - (\delta - a)\delta^{-1}(\delta - a) > 0 \Leftrightarrow \delta^2 - (\delta - a)^2 > 0 \Leftrightarrow \delta > a/2$. This implies that if $\delta > a/2$ then Assumption 4.1.3.(a) is satisfied in \mathcal{H}_I . We now consider the case of $\delta = a/2$. The sufficient condition of Lemma 4.1.2 for Assumption 4.1.3.(b), is that $\frac{a}{2} \begin{bmatrix} I & -I \\ -I & I \end{bmatrix} = \frac{a}{2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \otimes I \succcurlyeq 0$, which is satisfied because $a > 0$ and $\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \otimes I$ has non-negative eigenvalues. We finally consider the case $\delta \in (0, a/2)$. From the sufficient condition in Lemma 4.1.2, we get that x^{i*} is SPC in $\mathcal{H}_{(a-\delta)I}$ if $\delta \in (0, a)$ and consequently Assumption 4.1.3.(d) is met. The result then follows from Theorem 4.2.1 and 4.2.3. \blacksquare

4.5.3 Proofs of the results stated in Section 4.3

Proof of Theorem 4.3.1 (NAG Nash equilibria)

Since, by Assumption 4.1.2, $x^{i*}(z^i)$ in (4.2) is a continuous mapping in z^i , the mapping $\mathbf{x}^*(\mathbf{z})$ in (4.20) is continuous in \mathbf{z} . Consequently, the mapping $\mathcal{A}_\nu(\mathbf{z}) = \mathcal{P}_\nu \mathbf{x}^*(\mathbf{z})$ is continuous. By Assumption 4.1.1(a), the set $\mathcal{P}_\nu \mathcal{X}_{1 \times N}$ is compact and convex. Therefore, by Lemma 3.2.1 the mapping \mathcal{A}_ν admits at least one fixed point. Consider now an arbitrary fixed point $\bar{\mathbf{z}} = [\bar{z}^1; \dots; \bar{z}^N] \in \mathbb{R}^{Nn}$ of the aggregation mapping \mathcal{A}_ν in (4.21), that is $\bar{z}^i = \sum_{j=1}^N P_{ij}^\nu x^{j*}(\bar{z}^j)$. According to Definition 4.2, the set of strategies $\{\bar{x}^i := x^{i*}(\bar{z}^i)\}_{i=1}^N$ is a Nash equilibrium if no agent i , given the strategies $\{\bar{x}^j\}_{j \neq i}^N$ of all the other agents, can improve its cost, that is $\bar{x}^i = \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, P_{ii}^\nu x^i + \sum_{j \neq i}^N P_{ij}^\nu \bar{x}^j)$. By definition of fixed point and using the fact that $P_{ii}^\nu = 0$, for all i

$$\begin{aligned} \bar{x}^i &= x^{i*}(\bar{z}^i) = \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, \bar{z}^i) = \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, \sum_{j=1}^N P_{ij}^\nu \bar{x}^j) \\ &= \arg \min_{x^i \in \mathcal{X}^i} J^i(x^i, P_{ii}^\nu x^i + \sum_{j \neq i}^N P_{ij}^\nu \bar{x}^j). \end{aligned}$$

\blacksquare

Proof of Theorem 4.3.2 (Distributed convergence to a fixed point)

We start with some preliminary results.

Lemma 4.5.1. (Regularity of the extended optimizer) *Let \mathbf{x}^* be the extended mapping in (4.20). If A. 4.1.3.(a)/ A. 4.1.3.(b)/ A. 4.1.3.(c)/ A. 4.1.3.(d) holds, then the mapping \mathbf{x}^* is a CON/NEX/FNE/AMON in $\mathcal{H}_{I_N \otimes S}$, respectively.*

Proof. a) For all $i \in \mathbb{Z}[1, N]$, by A. 4.1.3.(a) the mapping x^{i*} is a CON in \mathcal{H}_S , with some rate $\delta_i \in [0, 1)$. Therefore, for any $\mathbf{r}, \mathbf{s} \in \mathbb{R}^{Nn}$ we have $\|\mathbf{x}^*(\mathbf{r}) - \mathbf{x}^*(\mathbf{s})\|_{I_N \otimes S}^2 = \|[x^{1*}(r^1) - x^{1*}(s^1); \dots; x^{N*}(r^N) - x^{N*}(s^N)]\|_{I_N \otimes S}^2 = \|x^{1*}(r^1) - x^{1*}(s^1)\|_S^2 + \dots + \|x^{N*}(r^N) - x^{N*}(s^N)\|_S^2 \leq \delta_1^2 \|r^1 - s^1\|_S^2 + \dots + \delta_N^2 \|r^N - s^N\|_S^2 \leq \max_{i \in \mathbb{Z}[1, N]} \delta_i^2 \|\mathbf{r} - \mathbf{s}\|_{I_N \otimes S}^2$. Note that it holds $\delta := \max_{i \in \mathbb{Z}[1, N]} \delta_i < 1$ since N is finite.

- b) As in the previous point, with $\delta_i = 1$ for all i .
- c) If A. 4.1.3.(c) holds then, the mappings $x^{i\star}$ are **FNE** in \mathcal{H}_S . Therefore, by Lemma 3.1.1, for all $\mathbf{r}, \mathbf{s} \in \mathbb{R}^{Nn}$ we have $\|\mathbf{x}^\star(\mathbf{r}) - \mathbf{x}^\star(\mathbf{s})\|_{I_N \otimes S}^2 = \sum_{i=1}^N \|x^{i\star}(r^i) - x^{i\star}(s^i)\|_S^2 \leq \sum_{i=1}^N (r^i - s^i)^\top S (x^{i\star}(r^i) - x^{i\star}(s^i)) = (\mathbf{r} - \mathbf{s})^\top (I_N \otimes S) (\mathbf{x}^\star(\mathbf{r}) - \mathbf{x}^\star(\mathbf{s}))$.
- d) If A. 4.1.3.(d) holds, the mappings $-x^{i\star}$ are **MON** in \mathcal{H}_S . Therefore, for any $\mathbf{r}, \mathbf{s} \in \mathbb{R}^{Nn}$ we have $(-\mathbf{x}^\star(\mathbf{r}) + \mathbf{x}^\star(\mathbf{s}))^\top (I_N \otimes S) (\mathbf{r} - \mathbf{s}) = \sum_{i=1}^N (-x^{i\star}(r^i) + x^{i\star}(s^i))^\top S (r^i - s^i) \geq 0$.

□

Lemma 4.5.2. Consider $P \in \mathbb{R}^{N \times N}$, $\nu \in \mathbb{N}$ and $\mathcal{P}_\nu = P^\nu \otimes I_n$. For any $S \in \mathbb{R}^{n \times n}$, $S \succ 0$, let $\mathcal{S} := I_N \otimes S$. If $\|P\| \leq 1$, then $\|\mathcal{P}_\nu\|_{\mathcal{S}} \leq 1$.

Proof. The condition $\|P\| \leq 1$ implies $\|P^\nu\| \leq \|P\|^\nu \leq 1$. By Lemma 3.1.4 this implies $(P^{\nu\top} P^\nu - I_N) \preceq 0$. Moreover, by Lemma 3.1.4, $\|\mathcal{P}_\nu\|_{\mathcal{S}} \leq 1 \Leftrightarrow \mathcal{P}_\nu^\top \mathcal{S} \mathcal{P}_\nu - \mathcal{S} \preceq 0 \Leftrightarrow (P^\nu \otimes I_n)^\top (I_N \otimes S) (P^\nu \otimes I_n) - I_N \otimes S \preceq 0 \Leftrightarrow (P^{\nu\top} \otimes I_n^\top) (I_N \otimes S) (P^\nu \otimes I_n) - I_N \otimes S \preceq 0 \Leftrightarrow (P^{\nu\top} I_N P^\nu \otimes I_n^\top S I_n) - I_N \otimes S \preceq 0 \Leftrightarrow (P^{\nu\top} P^\nu) \otimes S - I_N \otimes S \preceq 0 \Leftrightarrow (P^{\nu\top} P^\nu - I_N) \otimes S \preceq 0$. Since $S \succ 0$ and $(P^{\nu\top} P^\nu - I_N) \preceq 0$, by the properties of the Kronecker product we conclude that $(P^{\nu\top} P^\nu - I_N) \otimes S \preceq 0$. Finally, by the previous equivalence, we have $\|\mathcal{P}_\nu\|_{\mathcal{S}} \leq 1$. □

Note that a single iteration of Algorithm 4 updates the signal $\mathbf{z}_{(k)} = [z_{(k)}^1; \dots; z_{(k)}^N]$ according to the mapping

$$\mathcal{A}_{\nu_1, \nu_2}(\mathbf{z}) := \mathcal{P}_{\nu_2} \mathbf{x}^\star(\mathcal{P}_{\nu_1} \mathbf{z}) \quad (4.32)$$

and $\mathcal{A}_{0, \nu} \equiv \mathcal{A}_\nu$. The following proposition characterizes the regularity properties of $\mathcal{A}_{\nu_1, \nu_2}$ for different choices of ν_1, ν_2 .

Proposition 4.5.3. (*Regularity of the update mapping*) The following statements hold.

1. If A. 4.1.3.(a) holds and $\|P\| \leq 1$, then the mapping $\mathcal{A}_{0, \nu} \equiv \mathcal{A}_\nu$ in (4.21) is a **CON** in $\mathcal{H}_{I_N \otimes S}$;
2. If A. 4.1.3.(b) holds and $\|P\| \leq 1$, then the mapping $\mathcal{A}_{0, \nu} \equiv \mathcal{A}_\nu$ in (4.21) is **NEX** in $\mathcal{H}_{I_N \otimes S}$;
3. If A. 4.1.3.(c) holds, $\nu \in 2\mathbb{N}$ and $P = P^\top$, then the mapping $\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}$ in (4.32) is **FNE** in $\mathcal{H}_{I_N \otimes S}$;
4. If A. 4.1.3.(d) holds, $\nu \in 2\mathbb{N}$ and $P = P^\top$, then the mapping $\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}$ in (4.32) is **SPC** in $\mathcal{H}_{I_N \otimes S}$. □

Proof. Let $\mathcal{S} := I_N \otimes S$.

1. By Lemma 4.5.1, \mathbf{x}^* is a CON and, by Lemmas 3.1.4 and 4.5.2, \mathcal{P}_ν is NEX. Hence the mapping $\mathcal{A}_\nu = \mathcal{P}_\nu \mathbf{x}^*$, composition of a CON mapping and a NEX one, is a CON in \mathcal{H}_S .
2. Analogous to the proof of point 1, with \mathbf{x}^* NE.
3. By Lemma 4.5.1, \mathbf{x}^* is FNE in \mathcal{H}_S . Since P is by construction a row-stochastic matrix and $P = P^\top$, P is a doubly stochastic matrix. Consequently $\|P\|_1 = \|P\|_\infty = 1$. From Hölder's inequality, we have $\|P\| \leq \sqrt{\|P\|_1 \cdot \|P\|_\infty} = 1$. Therefore, from Lemma 4.5.2, $\|\mathcal{P}_{\frac{\nu}{2}}\|_{\mathcal{S}} \leq 1$. Moreover, $P = P^\top$ implies $\mathcal{P}_{\frac{\nu}{2}} = \mathcal{P}_{\frac{\nu}{2}}^\top$ and $\mathcal{P}_{\frac{\nu}{2}} \mathcal{S} = (P_{\frac{\nu}{2}} \otimes I_n)(I_N \otimes S) = (P_{\frac{\nu}{2}} \otimes S) = (I_N \otimes S)(P_{\frac{\nu}{2}} \otimes I_n) = \mathcal{S} \mathcal{P}_{\frac{\nu}{2}}$. Therefore for any $\mathbf{r}, \mathbf{s} \in \mathbb{R}^{Nn}$, we have

$$\begin{aligned}
\|\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}(\mathbf{r}) - \mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}(\mathbf{s})\|_{\mathcal{S}}^2 &= \|\mathcal{P}_{\frac{\nu}{2}}(\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{r}))) - \mathcal{P}_{\frac{\nu}{2}}(\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{s})))\|_{\mathcal{S}}^2 \\
&\leq \|\mathcal{P}_{\frac{\nu}{2}}\|_{\mathcal{S}}^2 \|\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{r})) - \mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{s}))\|_{\mathcal{S}}^2 \leq \|\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{r})) - \mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{s}))\|_{\mathcal{S}}^2 \\
&\leq (\mathcal{P}_{\frac{\nu}{2}} \mathbf{r} - \mathcal{P}_{\frac{\nu}{2}} \mathbf{s})^\top \mathcal{S} (\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{r})) - \mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{s}))) \\
&= (\mathbf{r} - \mathbf{s})^\top \mathcal{P}_{\frac{\nu}{2}}^\top \mathcal{S} (\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{r})) - \mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{s}))) \\
&= (\mathbf{r} - \mathbf{s})^\top \mathcal{S} \mathcal{P}_{\frac{\nu}{2}} (\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{r})) - \mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{s}))) = (\mathbf{r} - \mathbf{s})^\top \mathcal{S} (\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}(\mathbf{r}) - \mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}(\mathbf{s})),
\end{aligned} \tag{4.33}$$

where the third inequality derives from \mathbf{x}^* being FNE. The inequality in (4.33) implies that $\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}(\cdot)$ is FNE in \mathcal{H}_S , by Lemma 3.1.1.

4. From Lemma 4.5.1, $-\mathbf{x}^*(\cdot)$ is MON in \mathcal{H}_S . Moreover $P = P^\top$ implies that $\mathcal{P}_{\frac{\nu}{2}}^\top \mathcal{S} = \mathcal{S} \mathcal{P}_{\frac{\nu}{2}}$. Therefore for any $\mathbf{r}, \mathbf{s} \in \mathbb{R}^{Nn}$,

$$\begin{aligned}
&(-\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}(\mathbf{r}) + \mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}(\mathbf{s}))^\top \mathcal{S} (\mathbf{r} - \mathbf{s}) \\
&= (-\mathcal{P}_{\frac{\nu}{2}}(\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{r}))) + \mathcal{P}_{\frac{\nu}{2}}(\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{s}))))^\top \mathcal{S} (\mathbf{r} - \mathbf{s}) \\
&= (-\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{r})) + \mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{s})))^\top \mathcal{P}_{\frac{\nu}{2}}^\top \mathcal{S} (\mathbf{r} - \mathbf{s}) \\
&= (-\mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{r})) + \mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\mathbf{s})))^\top \mathcal{S} (\mathcal{P}_{\frac{\nu}{2}} \mathbf{r} - \mathcal{P}_{\frac{\nu}{2}} \mathbf{s}) \\
&= (-\mathbf{x}^*(\tilde{\mathbf{r}}) + \mathbf{x}^*(\tilde{\mathbf{s}}))^\top \mathcal{S} (\tilde{\mathbf{r}} - \tilde{\mathbf{s}}) \geq 0,
\end{aligned}$$

which implies that $-\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}$ is MON. By Lemma 3.1.2, $(\mathcal{I} - \mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}})$ is SMON ($\epsilon = 1$) and thus $\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}$ is SPC in \mathcal{H}_S . □

We are now ready to prove Theorem 4.3.2. We start by noting that, since the sets \mathcal{X}^i are convex and compact for every $i \in \mathbb{Z}[1, N]$, the mapping $\mathcal{A}_{\nu_1, \nu_2} : \mathcal{P}_{\nu_2} \mathcal{X}_{1 \times N} \rightarrow \mathcal{P}_{\nu_2} \mathcal{X}_{1 \times N}$ is defined from a compact and convex set to itself. For simplicity let $\mathbf{a}_{(k)} := [\mathcal{E}_{\nu_1, 0}^1(\mathbf{z}_{(k)}); \dots; \mathcal{E}_{\nu_1, 0}^N(\mathbf{z}_{(k)})]$.

Proof of statement 1 and 2 of Theorem 4.3.2

From Proposition 4.5.3, under the assumption of statement 1 (or statement 2), $\mathcal{A}_{0,\nu}(\cdot) \equiv \mathcal{A}_\nu(\cdot)$ is a **CON** (or **NEX**) in $\mathcal{H}_{I_N \otimes S}$. Therefore, by using the suggested mappings $\mathbf{z}_{(k)}$ converges to a fixed point of the mapping \mathcal{A}_ν , see Table 3.1. The proof is immediate upon noticing that, since $\nu_1 = 0$, $\mathbf{a}_{(k)} := \mathcal{P}_{\nu_1} \mathbf{z}_{(k)} \equiv \mathbf{z}_{(k)}$. Moreover, a **CON** mapping has a unique fixed point.

Proof of statement 3 and 4 of Theorem 4.3.2

From Proposition 4.5.3, under the assumption of statement 3, the mapping $\mathbf{z}_{(k)} \mapsto \mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}(\mathbf{z}_{(k)})$ is **FNE**. Therefore, by using the suggested mappings, $\mathbf{z}_{(k)}$ converges to a fixed point $\bar{\mathbf{z}}$ of the mapping $\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}$. Under the assumption of statement 4, on the other hand, $\mathbf{z}_{(k)} \mapsto \mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}(\mathbf{z}_{(k)})$ is **SPC** with $\epsilon = 1$ and Lipschitz with constant L (immediate consequence of Assumption 4.1.2). Consequently, it has a unique fixed point by Lemma 3.2.3. Convergence to such point can be guaranteed by using Φ_k^M (Lemma 3.2.8) or, if the Lipschitz constant L of $\mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}$ is known, by using Φ^K with $\lambda \in (0, \frac{2}{1+L^2})$, see Lemma 3.2.7. Finally, in both cases, $\mathbf{a}_{(k)} := \mathcal{P}_{\nu_1} \mathbf{z}_{(k)}$ hence, for $\nu_1 = \frac{\nu}{2}$, $\mathbf{a}_{(k)}$ converges to $\bar{\mathbf{a}} := \mathcal{P}_{\frac{\nu}{2}} \bar{\mathbf{z}}$, which is a fixed point of the mapping \mathcal{A}_ν since $\bar{\mathbf{z}} = \mathcal{A}_{\frac{\nu}{2}, \frac{\nu}{2}}(\bar{\mathbf{z}}) \Rightarrow \bar{\mathbf{z}} = \mathcal{P}_{\frac{\nu}{2}} \mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\bar{\mathbf{z}})) \Rightarrow \mathcal{P}_{\frac{\nu}{2}} \bar{\mathbf{z}} = \mathcal{P}_{\nu} \mathbf{x}^*(\mathcal{P}_{\frac{\nu}{2}}(\bar{\mathbf{z}})) \Rightarrow \bar{\mathbf{a}} = \mathcal{P}_{\nu} \mathbf{x}^*(\bar{\mathbf{a}}) \Rightarrow \bar{\mathbf{a}} = \mathcal{A}_\nu(\bar{\mathbf{a}})$. ■

Proof of Corollary 4.3.3 (Distributed opinion dynamics)

The cost function in (4.27) can be rewritten, up to constant terms that do not depend on x^i , as $J^i(x^i, \sigma^i) := (1 + \theta_i) \|x^i\|^2 - 2(\sigma^i + \theta_i x_{(0)}^i)^\top x^i$, for $\sigma^i = \sum_{j \neq i} P_{ij} x^j$. Consequently, the game in (4.27) is an **NAG** game with $\nu = 1$, quadratic cost with $q_i := (1 + \theta_i)$, $Q = I_n$, $C = -I_n$, $c = -\theta_i x_{(0)}^i$ and constraint set \mathcal{X}^i . Since σ^i does not depend on x^i , $P_{ii} = 0$ for every i and Assumption 4.3.1 is satisfied. It follows from Lemma 4.1.2 that Assumption 4.1.2 holds. For every agent $i \in \mathbb{Z}[1, N]$, the following equivalence holds $M_i = \begin{bmatrix} q_i Q & -C \\ -C & q_i Q \end{bmatrix} = \begin{bmatrix} (1+\theta_i) & 1 \\ 1 & (1+\theta_i) \end{bmatrix} \otimes I_n$, hence the eigenvalues of M_i are θ_i and $2 + \theta_i$, both with multiplicity n and if $\theta_i > 0$ then $M_i \succ 0$ and Assumption 4.1.3.(a) holds; if $\theta_i \geq 0$ then $M_i \succcurlyeq 0$ and Assumption 4.1.3.(b) holds. Consequently, by Theorem 4.3.2 the given conditions guarantee convergence of Algorithm 4 to a fixed point $\bar{\mathbf{z}}$ of the aggregation mapping \mathcal{A}_1 . Finally, Theorem 4.3.1 guarantees that the set of strategies $\{x^{i*}(\bar{\mathbf{z}}^i)\}_{i=1}^N$ is a Nash equilibrium. ■

4.5.4 Proofs of the results stated in Section 4.4

Proof of Lemma 4.4.1 (Convergence of aggregation mappings)

The fact that $\mathcal{A}(z) = \mathcal{I}x^*(z)$ admits at least one fixed point can be proven as done for $\mathcal{A}_\nu(z) = \mathcal{P}_\nu x^*(z)$ in the proof of Theorem 4.3.1. Let $D_N := \max_{x \in \mathcal{X}_{1 \times N}} \|x\|$; then for any $z \in \mathbb{R}^{Nn}$, $\|\mathcal{A}_\nu(z) - \mathcal{A}(z)\| = \|\mathcal{P}_\nu x^*(z) - \mathcal{I}x^*(z)\| \leq \|\mathcal{P}_\nu - \mathcal{I}\| \|x^*(z)\| \leq \|\mathcal{P}_\nu - \mathcal{I}\| D_N$. Hence $\sup_{z \in \mathbb{R}^{Nn}} \|\mathcal{A}_\nu(z) - \mathcal{A}(z)\| \leq \|\mathcal{P}_\nu - \mathcal{I}\| D_N$. By Assumption 4.4.1, $\lim_{\nu \rightarrow \infty} \|P^\nu - \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top\| = 0$ which implies that $\lim_{\nu \rightarrow \infty} \|\mathcal{P}_\nu - \mathcal{I}\| = \lim_{\nu \rightarrow \infty} \|(P^\nu - \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top) \otimes I_n\| = 0$, therefore $\lim_{\nu \rightarrow \infty} \sup_{z \in \mathbb{R}^{Nn}} \|\mathcal{A}_\nu(z) - \mathcal{A}(z)\| = 0$. \blacksquare

Proof of Theorem 4.4.2 (AAG Nash equilibria)

We follow a similar argument as in Theorem 4.2.1. Given Assumption 4.1.4 for all N and all $i \in \mathbb{Z}[1, N]$, the mappings J^i are uniformly Lipschitz with constant $L_J > 0$. For any $(N, \nu) \in \mathbb{N}^2$, consider an arbitrary fixed point $\bar{z} = [\bar{z}^1; \dots; \bar{z}^N] \in \mathbb{R}^{Nn}$ of the aggregation mapping \mathcal{A}_ν in (4.21), that is $\bar{z}^i = \sum_{j=1}^N P_{ij}^\nu x^{j*}(\bar{z}^j)$ and let $\bar{x}^i := x^{i*}(\bar{z}^i)$, for all $i \in \mathbb{Z}[1, N]$, $D := \max_{x \in \mathcal{X}} \|x\|$. By definition of optimal response we have

$$J^i(\bar{x}^i, \bar{z}^i) \leq J^i(x^i, \bar{z}^i) \quad \text{for all } x^i \in \mathcal{X}^i.$$

Using the equivalence $\bar{z}^i = \sum_{j=1}^N P_{ij}^\nu \bar{x}^j$ we get

$$J^i(\bar{x}^i, \sum_{j=1}^N P_{ij}^\nu \bar{x}^j) \leq J^i(x^i, \sum_{j=1}^N P_{ij}^\nu \bar{x}^j) \quad \text{for all } x^i \in \mathcal{X}^i. \quad (4.34)$$

Let us compute the cost of strategy \bar{x}^i in the average aggregative game

$$\begin{aligned} J^i(\bar{x}^i, \frac{1}{N} \sum_{j=1}^N \bar{x}^j) &= J^i(\bar{x}^i, \sum_{j=1}^N P_{ij}^\nu \bar{x}^j) + [J^i(\bar{x}^i, \frac{1}{N} \sum_{j=1}^N \bar{x}^j) - J^i(\bar{x}^i, \sum_{j=1}^N P_{ij}^\nu \bar{x}^j)] \\ &\leq J^i(\bar{x}^i, \sum_{j=1}^N P_{ij}^\nu \bar{x}^j) + L_J \|\frac{1}{N} \sum_{j=1}^N \bar{x}^j - \sum_{j=1}^N P_{ij}^\nu \bar{x}^j\| \\ &\leq J^i(\bar{x}^i, \sum_{j=1}^N P_{ij}^\nu \bar{x}^j) + L_J \sum_{j=1}^N |\frac{1}{N} - P_{ij}^\nu| \|\bar{x}^j\| \\ &\leq J^i(\bar{x}^i, \sum_{j=1}^N P_{ij}^\nu \bar{x}^j) + L_J D \|P^\nu - \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top\|_\infty \\ &\leq J^i(x^i, \sum_{j=1}^N P_{ij}^\nu \bar{x}^j) + L_J D \|P^\nu - \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top\|_\infty \\ &\leq J^i(x^i, \frac{1}{N} \sum_{j=1}^N \bar{x}^j) + 2L_J D \|P^\nu - \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top\|_\infty \\ &\leq J^i(x^i, \frac{1}{N} (x^i + \sum_{j \neq 1}^N \bar{x}^j)) + 2L_J D (\|P^\nu - \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top\|_\infty + \frac{1}{N}), \end{aligned}$$

where the 4th inequality follows from (4.34) and the last one can be proven as in Theorem 4.2.1. Consequently, if we define $K := 2L_J D$ an arbitrary agent i can improve its cost at most by an amount $\varepsilon_{N, \nu} := K (\|P^\nu - \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top\|_\infty + \frac{1}{N})$ if all other strategies $\{\bar{x}^j := x^{j*}(\bar{z}^j)\}_{j \neq i}^N$ are fixed. Therefore, the set of strategies $\{\bar{x}^i\}_{i=1}^N$ is an $\varepsilon_{N, \nu}$ -Nash equilibrium for the AAG in (4.9). K is a constant that does not depend on N, P or ν and for any fixed N we have $\|P^\nu - \frac{1}{N} \mathbb{1}_N \mathbb{1}_N^\top\|_\infty \rightarrow 0$ as $\nu \rightarrow \infty$. Hence, for all $\varepsilon > 0$ and for any fixed $N > \bar{N} := \frac{K}{\varepsilon}$, there exists $\bar{\nu}$ such that for all $\nu \geq \bar{\nu}$, we have $\varepsilon_{N, \nu} < \varepsilon$. \blacksquare

Proof of Corollary 4.4.3 (ε -convergence rate)

By the proof of Theorem 4.4.2 the set of strategies $\{x^{i*}(\bar{z}^i)\}_{i=1}^N$ is an $\varepsilon_{N,\nu}$ -Nash equilibrium for the game in (4.9), with $\varepsilon_{N,\nu} = K(\frac{1}{N} + \|P^\nu - \frac{1}{N}\mathbb{1}_N\mathbb{1}_N^\top\|_\infty)$. By the properties of the matrix norm $\|P^\nu - \frac{1}{N}\mathbb{1}_N\mathbb{1}_N^\top\|_\infty \leq \sqrt{N}\|P^\nu - \frac{1}{N}\mathbb{1}_N\mathbb{1}_N^\top\|_2 = \sqrt{N}\sigma_{\max}(P^\nu - \frac{1}{N}\mathbb{1}_N\mathbb{1}_N^\top) = \sqrt{N}\max_{\lambda \in \Lambda(P^\nu - \frac{1}{N}\mathbb{1}_N\mathbb{1}_N^\top)} |\lambda| = \sqrt{N}\mu_N^\nu \leq \sqrt{N}\mu^\nu$, where we used the fact that the matrix $P^\nu - \frac{1}{N}\mathbb{1}_N\mathbb{1}_N^\top$ is symmetric and, since P is symmetric, primitive and doubly stochastic, 1 is a simple eigenvalue and it holds $\Lambda(P^\nu - \frac{1}{N}\mathbb{1}_N\mathbb{1}_N^\top) = ((\Lambda(P))^\nu \setminus \{1\}) \cup \{0\}$. Hence $\varepsilon_{N,\nu} \leq K(\frac{1}{N} + \sqrt{N}\mu^\nu)$. ■

Proof of Corollary 4.4.4 (Distributed demand-response)

Given the cost function in (4.29) with $p(\bar{\sigma}_t) := \lambda\bar{\sigma}_t + p_0$, it holds $M_i = \begin{bmatrix} q_i Q & -C \\ -C & q_i Q \end{bmatrix} = \begin{bmatrix} \rho_i & -\frac{\lambda}{2} \\ -\frac{\lambda}{2} & \rho_i \end{bmatrix} \otimes I_T$, hence the eigenvalues of M_i are $\rho_i + \frac{\lambda}{2}$ and $\rho_i - \frac{\lambda}{2}$, both with multiplicity T . It follows from Lemma 4.1.2 that Assumption 4.1.2 holds and that if $\rho_i - \frac{\lambda}{2} > 0$, then $M_i \succ 0$ and Assumption 4.1.3.(a) holds while if $\rho_i - \frac{\lambda}{2} \geq 0$ then $M_i \succcurlyeq 0$ and Assumption 4.1.3.(b) holds. Moreover, $C = \frac{\lambda}{2}I_T \succ 0$ and therefore Assumption 4.1.3.(c) always holds. Consequently, by Theorem 4.3.2 the given conditions guarantee convergence of Algorithm 4 to a fixed point \bar{z} of the aggregation mapping \mathcal{A}_ν . The conclusion follows from Theorem 4.4.2. ■

Boundedly rational agents

In this chapter we address the coordination problem for populations of *boundedly rational* agents, that is, for agents that do not have the computational capabilities to solve an optimization problem as the one in (4.2)¹. In Section 5.1 we present an equivalent characterization of Nash equilibria in terms of variational inequalities and we briefly review some decentralized and distributed algorithms that have been suggested in the literature relying on this equivalence. In Section 5.2 we introduce a more general class of games by assuming that the agents strategies are not only constrained by the local sets \mathcal{X}^i , but must also satisfy a global coupling constraint $\mathcal{C} := \{x = [x^1; \dots; x^N] \mid Ax \leq b\}$. We then review the concept of generalized Nash equilibrium and its connection with variational inequalities. Using this relation, we propose in Section 5.3 a *decentralized algorithm* to steer a population of boundedly rational agents to a generalized Nash equilibrium of an AAG by means of a central operator. A *distributed* version of this algorithm, that can be used for NAGs and AAGs with boundedly rational agents and global coupling constraints, is suggested in the concluding section 10.1.1. All the proofs are given in the Appendix (Section 5.4).

Part of the results of this chapter have been published in [PPG⁺16].

5.1 Relation among Nash equilibria and variational inequalities

Consider a generic game \mathcal{G} (i.e., not necessarily aggregative)

$$\mathcal{G} = \left\{ \begin{array}{ll} x_{\text{br}}^i(x^{-i}) := \arg \min_{x^i \in \mathbb{R}^n} \tilde{J}^i(x^i, x^{-i}) & \forall i = 1, \dots, N, \\ \text{s.t. } x^i \in \mathcal{X}^i & \end{array} \right. \quad (5.1)$$

¹It is important to stress that the distinction between *myopic* and *boundedly rational* agents is a distinction between the computational capabilities of the agents and not a distinction among games. The games that we consider in this chapter are the same games as defined in Chapter 4. The difference is in the algorithms suggested to reach the Nash equilibrium which, in the case of bounded rationality, cannot require the computation of the optimal response (4.2) as part of the strategies update step.

satisfying the following assumption.

Assumption 5.1.1 (Convexity of cost functions and constraints). *The cost functions $\tilde{J}^i(x^i, x^{-i})$ are convex and \mathcal{C}^1 in x^i for fixed x^{-i} . The local constraint sets \mathcal{X}^i are compact and convex $\forall i \in \mathbb{Z}[1, N]$.*

The set of Nash equilibria of the game \mathcal{G} can be equivalently characterized in terms of a variational inequality (see Section 3.3) that has as operator a vector whose i -th block component is the gradient of the cost function $\tilde{J}^i(x^i, x^{-i})$ of player i with respect to its own strategy x^i , that is,

$$f(x) = [\nabla_{x^1} \tilde{J}^1(x^1, x^{-1}); \dots; \nabla_{x^N} \tilde{J}^N(x^N, x^{-N})]. \quad (5.2)$$

Proposition 5.1.1 (Nash equilibria and VIs). *Suppose Assumption 5.1.1 holds. A set of strategies $\{\bar{x}^i \in \mathcal{X}^i\}_{i=1}^N$ is a Nash equilibrium of the game in (5.1) if and only if*

$$\bar{x} := [\bar{x}^1; \dots; \bar{x}^N] \in \text{SOL}(\mathcal{X}_{1 \times N}, f),$$

where $\mathcal{X}_{1 \times N} := \mathcal{X}^1 \times \dots \times \mathcal{X}^N$ and $f(x)$ is as in (5.2).

Proof. This is a classical result, we report the proof for completeness [FP03], [BT97, Section 3.5.1]. A set of strategies $\{\bar{x}^i \in \mathcal{X}^i\}_{i=1}^N$ is a Nash equilibrium if and only if it solves the N coupled optimization problems in (5.1). By Assumption 5.1.1, these are convex optimization problems and the sets \mathcal{X}^i are convex and closed, therefore, by the minimum principle (Proposition 3.3.1), this requirement is equivalent to

$$\begin{aligned} & \bar{x}^i \in \text{SOL}(\mathcal{X}^i, \nabla_{x^i} \tilde{J}^i(\cdot, \bar{x}^{-i})) \quad \text{for all } i \in \mathbb{Z}[1, N] \\ \Leftrightarrow & \nabla_{x^i} \tilde{J}^i(\bar{x}^i, \bar{x}^{-i})^\top (x^i - \bar{x}^i) \geq 0 \quad \text{for all } x^i \in \mathcal{X}^i \quad \text{and } i \in \mathbb{Z}[1, N] \\ \Leftrightarrow & \sum_{i=1}^N \nabla_{x^i} \tilde{J}^i(\bar{x}^i, \bar{x}^{-i})^\top (x^i - \bar{x}^i) = f(\bar{x})^\top (x - \bar{x}) \geq 0 \quad \text{for all } x \in \mathcal{X}_{1 \times N} \\ \Leftrightarrow & \bar{x} \in \text{SOL}(\mathcal{X}_{1 \times N}, f). \end{aligned}$$

A graphical illustration for $N = 2$ is given in Figure 5.1. □

The results on existence and uniqueness of the solution for VIs, summarized in Section 3.3.1, can be readily apply to show the existence and uniqueness of Nash equilibria.

Proposition 5.1.2 (Nash equilibria: existence and uniqueness). *Suppose that Assumption 5.1.1 holds. If the function $f(x)$ defined in (5.2) is continuous the Nash game in (5.1) has at least one Nash equilibrium. If the operator f is additionally *SMON* then the Nash equilibrium is unique.*

Proof. Under Assumption 5.1.1 $\mathcal{X}_{1 \times N}$ is convex and compact and f is continuous. The results follow from Proposition 5.1.1 and Proposition 3.3.7. □

Note that this result subsumes the one in Proposition 2.1.1.

5.1.1 Decentralized and distributed algorithms

The equivalence in Proposition 5.1.1 can be readily employed, in the case of agents with bounded rationality, to suggest coordination schemes that aim at finding a Nash equilibrium by using one of the algorithms for the solution of VIs presented in Section 3.3.2. For example, if the operator f is SMON one can use the projection scheme

$$x_{(k+1)} = \Pi_{\mathcal{X}_{1 \times N}}(x_{(k)} - \alpha f(x_{(k)})), \quad (5.3)$$

where $\alpha > 0$, $x_{(k)} := [x_{(k)}^1; \dots; x_{(k)}^N]$ and f is as in (5.2). Exploiting the cartesian structure of the set $\mathcal{X}_{1 \times N} = \mathcal{X}^1 \times \dots \times \mathcal{X}^N$ and the properties of the projection operator, it is immediate to see that the update rule (5.3) can be distributed among the agents, resulting in

$$x_{(k+1)}^i = \Pi_{\mathcal{X}^i}(x_{(k)}^i - \alpha \nabla_{x^i} \tilde{J}^i(x_{(k)}^i, x_{(k)}^{-i})) \quad \text{for all } i \in \mathbb{Z}[1, N]. \quad (5.4)$$

Alternatively, one could aim at solving the VI($\mathcal{X}_{1 \times N}, f$) by using a nonlinear algorithm, as the ones described in [BT97, Section 3.5.6]. Specifically, in the context of games, the nonlinear Gauss-Seidel algorithm and the nonlinear Jacobi algorithm for VIs [BT97, Section 3.5.6] coincide with the sequential and simultaneous BR dynamics given in (2.5) and (2.6), respectively (as introduced in Section 2.3). Conditions guaranteeing their convergence have been derived in [BT97] by using the concept of block contraction.

Definition 5.1. (Block-maximum norm and block contraction [BT97, Section 3.1.2]) *Given a vector $x := [x^1; \dots; x^N] \in \mathbb{R}^{Nn}$ with block components $x^i \in \mathcal{X}^i \subseteq \mathbb{R}^n$ we define the block-maximum norm*

$$\|x\|_B := \max_{i=1, \dots, N} \|x^i\|,$$

where $\|\cdot\|$ is the Euclidean norm in \mathbb{R}^n . A mapping $f : \mathbb{R}^{Nn} \rightarrow \mathbb{R}^{Nn}$ is a block contraction if it is a CON under the above introduced block-maximum norm.

Lemma 5.1.3 (Convergence of sequential and simultaneous BR dynamics [BT97]). *Consider the game \mathcal{G} in (5.1) and suppose that Assumption 5.1.1 holds. Suppose that there exists $\alpha > 0$ such that the mapping $R(x) = x - \alpha f(x)$ is a block contraction, where f is as in (5.2). Then the game \mathcal{G} has a unique Nash equilibrium and for any initial condition $x_{(0)}$ both the sequential and simultaneous BR dynamics converge to it, geometrically.*

A sketch of the proof is given in the Appendix. We briefly comment here on the reason why the mapping $R(x) = x - \alpha f(x)$ appears in the previous lemma. We start by noticing that the projection algorithm in (5.3) can be rewritten as $x_{(k+1)} = \Pi_{\mathcal{X}_{1 \times N}}(R(x_{(k)}))$. Consequently, if $R(\cdot)$ is a CON, the composition $\Pi_{\mathcal{X}_{1 \times N}}(R(\cdot))$ is also a CON and algorithm (5.3) is a Picard–Bannach iteration applied to a CON mapping, which converges by Lemma 3.2.4². What Lemma 5.1.3 states is that if $R(\cdot)$ is a CON in the block-maximum

²Note that a sufficient condition for the existence of $\alpha > 0$ such that $R(\cdot)$ is a CON is that f is SMON, consistently with our previous statement.

norm (instead of the Euclidean norm) then not only the projection algorithm converges, but also the sequential and simultaneous BR dynamics.

5.1.2 Average aggregative games

In order to locally implement gradient steps as the one in (5.4), each agent needs to be able to evaluate the gradient of its cost function. This is in general not a trivial task since the gradient depends in general on the strategies of the whole population. In the case of AAGs, however, the gradient depends only on the local strategy of player i , $x_{(k)}^i$, and on the average of the population $\bar{\sigma}_{(k)} = \frac{1}{N} \sum_{j=1}^N x_{(k)}^j$. In fact

$$\begin{aligned} \nabla_{x^i} J^i(x^i, \bar{\sigma}(x)) &= [\nabla_{s_1} J^i(s_1, s_2) \nabla_{x^i} x^i + \nabla_{s_2} J^i(s_1, s_2) \nabla_{x^i} \bar{\sigma}(x)] \Big|_{s_1=x^i, s_2=\bar{\sigma}(x)} \quad (5.5) \\ &= \nabla_{x^i} J^i(x^i, s_2) \Big|_{s_2=\bar{\sigma}(x)} + \frac{1}{N} \nabla_{s_2} J^i(x^i, s_2) \Big|_{s_2=\bar{\sigma}(x)}, \end{aligned}$$

where we applied the chain rule to the composite function $J^i(x^i, \bar{\sigma}(x))$. Consequently, decentralized gradient schemes relying on the presence of a central operator that computes and broadcasts the average at every step, [CLLV14, PKL16], as well as distributed schemes where the agents try to recover the average by local communications [KNS12, KNS16] have been suggested in the literature, relying on the assumption that the operator $f(x)$ is SMON (so that the projection algorithm converges). In the rest of the chapter we present an extension of these works to the case when the strategies of the agents do not only need to satisfy the local constraints \mathcal{X}^i , but also a constraint \mathcal{C} that couples the agents strategies, so that it is not possible to distribute the projection operation in (5.3) as done in (5.4). Before doing so, we briefly comment on the relation between the VI algorithms discussed in this subsection and the fixed point algorithms of Section 4.2, for the case of AAGs with quadratic cost functions.

AAGs with quadratic cost functions

The fundamental difference between the VI algorithms suggested in [CLLV14, PKL16, KNS12, KNS16] and the fixed point algorithms proposed in Section 4.2 lies in the strategies update step. Specifically, in the VI algorithms the agents update their strategies by taking a gradient step, while in the schemes of Section 4.2 they select the strategy that minimizes their cost, given the current reference. Therefore, for some applications (e.g., if the agents have limited computation capability) only the former class of algorithms can be applied while for others (e.g., if the agents are myopic and/or the game is repeated over time, as in opinion dynamics) only the latter is appropriate. Nonetheless, for some applications (e.g., the PEV charging coordination game described in Section 4.2.4) both schemes are feasible. One can then choose which algorithm to use based on its conditions for convergence, as derived in Lemma 4.1.2 and in the following Lemma

5.3.1. We briefly compare these conditions in Table 5.1, where also the sequential and simultaneous BR dynamics are considered as candidate algorithms.

Sufficient conditions for the convergence of the BR dynamics are derived in Proposition 2.3.1, for potential games, and in the following lemma.

Lemma 5.1.4 (Sufficient conditions for block contraction in the quadratic case). *Suppose that Assumption 5.1.1 holds, the cost functions are as in (4.5) and that there exists $\beta \in (0, 1)$ such that $\beta\sigma_{\min}(q_i Q) \geq \sigma_{\max}(C)$ for all $i \in \mathbb{Z}[1, N]$ and all population sizes N . Then for any $N > \frac{\beta}{1-\beta}$ there exists $\alpha > 0$ such that the mapping $R(x) = x - \alpha f(x)$ is a block contraction, where f is as in (5.2).*

	Sufficient conditions for all $i \in \mathbb{Z}[1, N]$ (under Assumption 5.1.1)	Fixed point approach (Section 4.2)	Variational approach (e.g. [PKL16])	Sequential BR dynamics (Section 2.3)	Simultaneous BR dynamics (Section 2.3)
$q_i > 0$	$M_i \succ 0$	✓			$(N \rightarrow \infty)$
	$M_i \succeq 0$	✓			$(N \rightarrow \infty)$
	$-q_i Q \preceq C = C^\top \prec 0$	✓		✓	$(N \rightarrow \infty)$
	$C = C^\top \succ 0$	✓	✓	✓	$N > \frac{\beta}{1-\beta}$
	$\beta\sigma_{\min}(q_i Q) \geq \sigma_{\max}(C), \beta < 1$			$N > \frac{\beta}{1-\beta}$	$N > \frac{\beta}{1-\beta}$
$q_i \geq 0$	$C + C^\top \succeq 0$		✓		

Table 5.1: Sufficient conditions for the convergence of the fixed point algorithms described in Section 4.2, the VI algorithm described for example in [PKL16] and the sequential and simultaneous BR dynamics, for games with quadratic cost. Note that in the case of simultaneous BR dynamics with “ $(N \rightarrow \infty)$ ” we mean that Algorithm 2 converges under the given condition and it coincides with the simultaneous BR dynamics when $N \rightarrow \infty$. Moreover, for the sequential BR dynamics we use the fact that $C = C^\top$ implies that the game is potential (see Remark 4.1).

5.2 Generalized Nash equilibria

In the rest of the chapter we consider an extension of the game in (5.1), where the strategies of the agents are constrained not only by the local sets \mathcal{X}^i , but also by a coupling affine constraint

$$\mathcal{C} := \{x = [x^1; \dots; x^N] \in \mathbb{R}^{Nn} \mid Ax \leq b\} \subset \mathbb{R}^{Nn}, \quad (5.6)$$

with $A \in \mathbb{R}^{m \times Nn}$, $b \in \mathbb{R}^m$. These constraints could model, for example, the fact that each agent has a feasible range for its usage level (\mathcal{X}^i) and, at the same time, the overall usage level cannot exceed a global capacity constraint (\mathcal{C}). For example, in the PEV

application the grid operator might want to impose an upper limit on the total energy that can be provided to the population at each time t . We define

$$\mathcal{Q} = \mathcal{X}_{1 \times N} \cap \mathcal{C} \quad \text{and} \quad \mathcal{Q}^i(x^{-i}) := \{x^i \in \mathcal{X}^i \mid Ax \leq b\}, \quad (5.7)$$

the set of global and local admissible strategies for player i , given x^{-i} , respectively.

To sum up, each agent aims at minimizing its cost function according to the game

$$\mathcal{G} := \begin{cases} \min_{x^i \in \mathbb{R}^n} & \tilde{J}^i(x^i, x^{-i}) \\ \text{s.t.} & x^i \in \mathcal{Q}^i(x^{-i}) \end{cases} \quad \forall i \in \mathbb{Z}[1, N]. \quad (5.8)$$

Note that the main difference between this game and the game in (5.1) is the fact that, in (5.8), the feasible set depends on the strategies of the other players. The definition of Nash equilibrium can be generalized for the game in (5.8) as follows.

Definition 5.2 (Generalized Nash equilibrium (GNE)). *A set of strategies $\{\bar{x}^i\}_{i=1}^N$ is called a generalized Nash equilibrium of the game \mathcal{G} in (5.8) if $\bar{x} := [\bar{x}^1; \dots; \bar{x}^N] \in \mathcal{Q}$ and $\forall i \in \mathbb{Z}[1, N]$*

$$\tilde{J}^i(\bar{x}^i, \bar{x}^{-i}) \leq \tilde{J}^i(x^i, \bar{x}^{-i}), \quad \forall x^i \in \mathcal{Q}^i(\bar{x}^{-i}).$$

In the following, we denote the set of GNEs of the game \mathcal{G} in (5.8) as $\text{GNE}(\mathcal{G})$. As for Nash equilibria also generalized Nash equilibria can be characterized in terms of VIs.

Proposition 5.2.1. *(Generalized Nash equilibria and VIs [FK07, Theorem 5]): Suppose Assumption 5.1.1 holds. Let $\bar{x} \in \mathbb{R}^{Nn}$. Then $\bar{x} \in \text{SOL}(\mathcal{Q}, f) \Rightarrow \bar{x} \in \text{GNE}(\mathcal{G})$, where f and \mathcal{Q} are as in (5.2) and (5.7), respectively, and \mathcal{G} is as in (5.8). \square*

Note that, in the generalized case, one can only show that the solutions of the VI are GNEs, but not the vice versa. An intuition of why the proof of Proposition 5.1.1 cannot be adapted to the generalized case is given in Figure 5.1.

The specific GNEs that are obtained as solution of the associated VI are called *variational* GNEs. For some specific games, one can show that the variational GNEs have some additional properties with respect to generic GNEs (as for example minimizing a suitably defined welfare cost). In general, however, the main advantage of variational GNEs is that, being the solution of a VI, they are easier to compute.

The results on existence and uniqueness of the solutions for VIs, summarized in Section 3.3.1, can be readily applied to show existence and uniqueness of variational GNEs.

Proposition 5.2.2 (Variational GNE: existence & uniqueness). *Suppose that Assumption 5.1.1 holds. If the operator f in (5.2) is continuous then the generalized Nash game in (5.8) has at least one generalized Nash equilibrium. If the operator f is additionally SMON then the variational GNE is unique.*

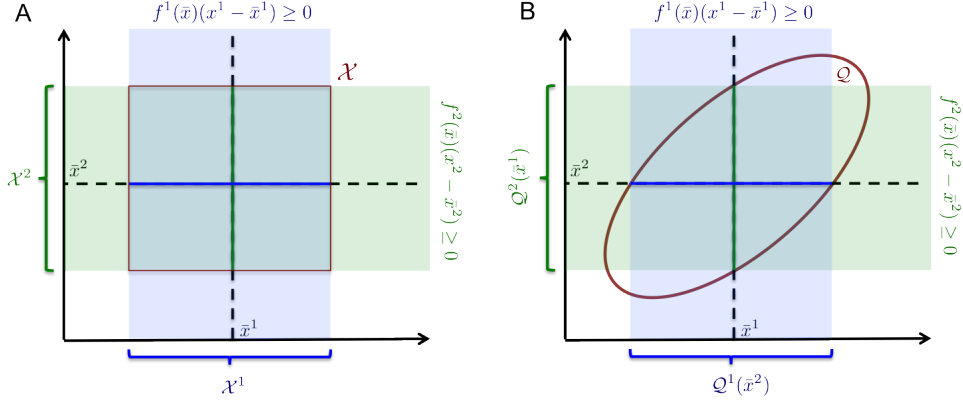


Figure 5.1: Let $f^i(x) := \nabla_{x^i} \tilde{J}^i(x^i, x^{-i})$. A) Illustration of the proof of Proposition 5.1.1. If \bar{x} is an NE, by the minimum principle, $f^1(\bar{x})(x^1 - \bar{x}^1) \geq 0$ for all $x^1 \in \mathcal{X}^1$, that is, for all x belonging to the blue area. Similarly, $f^2(\bar{x})(x^2 - \bar{x}^2) \geq 0$ in the green area. Therefore $\sum_{i=1}^2 f^i(\bar{x})(x^i - \bar{x}^i) \geq 0$ in the intersection, which coincides with $\mathcal{X}_{1 \times 2}$ and \bar{x} solves $\text{VI}(\mathcal{X}_{1 \times 2}, f)$. B) Illustration of the proof of Proposition 5.2.1. If \bar{x} is a GNE, by the minimum principle, $f^1(\bar{x})(x^1 - \bar{x}^1) \geq 0$ for all $x^1 \in \mathcal{Q}^1(\bar{x}^2)$, that is, in the blue area, and $f^2(\bar{x})(x^2 - \bar{x}^2) \geq 0$ in the green area. Clearly this is not sufficient to conclude that $\sum_{i=1}^2 f^i(\bar{x})(x^i - \bar{x}^i) \geq 0$ for all $x \in \mathcal{Q}$. The fact that a solution of the VI is an NE in case A (or a GNE in case B) can be proven as follows: if $\sum_{i=1}^2 f^i(\bar{x})(x^i - \bar{x}^i) \geq 0$ for all $x \in \mathcal{X}_{1 \times 2}$ (resp. $x \in \mathcal{Q}$), by setting $x^2 = \bar{x}^2$ we get $f^1(\bar{x})(x^1 - \bar{x}^1) \geq 0$ for all $x^1 \in \mathcal{X}^1$ ($x^1 \in \mathcal{Q}^1(\bar{x}^2)$) and by setting $x^1 = \bar{x}^1$ we get $f^2(\bar{x})(x^2 - \bar{x}^2) \geq 0$ for all $x^2 \in \mathcal{X}^2$ ($x^2 \in \mathcal{Q}^2(\bar{x}^1)$).

In the case of generalized Nash equilibrium problems the VI algorithms presented in Section 3.3.2 cannot be readily apply to guarantee convergence to the variational GNE as they all rely on projections onto the set \mathcal{Q} that, contrary to the case discussed in Section 5.1, cannot be easily decomposed into the cartesian product of local sets. Similarly to what is done in distributed optimization, one can however use a primal-dual approach that associates to each coupling constraint a dual variable (multiplier) λ and then solve a VI in the augmented space $[x; \lambda]$. We review this approach in the next subsection.

5.2.1 Primal-dual reformulation

Consider the $\text{VI}(\mathcal{Q}, f)$ and assume that \mathcal{Q} satisfies the following constraint qualification.

Assumption 5.2.1 (Slater). *There exists $g : \mathbb{R}^{Nn} \rightarrow \mathbb{R}^{m_g}$, $g \in \mathcal{C}^1$, such that $\mathcal{X}_{1 \times N} = \{x \in \mathbb{R}^{Nn} \mid g(x) \leq 0\}$. The sets $\mathcal{X}_{1 \times N} := \mathcal{X}^1 \times \dots \times \mathcal{X}^N$ and $\mathcal{Q} := \mathcal{X}_{1 \times N} \cap \mathcal{C}$, as defined in (5.7), satisfy the Slater's constraint qualification condition, as defined in Lemma 3.3.4.*

By using the KKT conditions for VIs (Proposition 3.3.6) one can immediately show the following result.

Proposition 5.2.3 (Augmented VI). *Suppose that Assumptions 5.1.1 and 5.2.1 hold. A point $\bar{x} \in \mathbb{R}^{Nn}$ solves $VI(\mathcal{Q}, f)$, with f and \mathcal{Q} as in (5.2) and (5.7), if and only if there exists $\bar{\lambda} \in \mathbb{R}^m$ such that $[\bar{x}; \bar{\lambda}]$ is a solution of the $VI(\mathcal{Y}, t)$, with $\mathcal{Y} := \mathcal{X}_{1 \times N} \times \mathbb{R}_{\geq 0}^m$ and*

$$t(x, \lambda) := \begin{bmatrix} f(x) + A^\top \lambda \\ -(Ax - b) \end{bmatrix}. \quad (5.9)$$

Proof. From the KKT conditions for $VI(\mathcal{Q}, f)$ we have that

$$\bar{x} \in \text{SOL}(\mathcal{Q}, f) \quad \Leftrightarrow \quad \exists \bar{\lambda}, \bar{\mu} > 0 \quad \text{s.t.} \quad \begin{cases} f(\bar{x}) + \nabla_x g(\bar{x})\bar{\mu} + A^\top \bar{\lambda} = 0 \\ \bar{\mu} \geq 0, g(\bar{x}) \leq 0, \bar{\mu} \perp g(\bar{x}) \\ \bar{\lambda} \geq 0, A\bar{x} - b \leq 0, \bar{\lambda} \perp b - A\bar{x} \end{cases} \quad (5.10)$$

and from the KKT conditions for $VI(\mathcal{Y}, t)$, by recalling $\mathcal{Y} = \{[\begin{smallmatrix} x \\ \lambda \end{smallmatrix}] \mid g(x) \leq 0, -\lambda \leq 0\}$,

$$[\bar{x}; \bar{\lambda}] \in \text{SOL}(\mathcal{Y}, t) \quad \Leftrightarrow \quad \exists \bar{\mu}, \bar{\eta} > 0 \quad \text{s.t.} \quad \begin{cases} \begin{bmatrix} f(\bar{x}) + A^\top \bar{\lambda} + \nabla_x g(\bar{x})\bar{\mu} \\ -(A\bar{x} - b) - \bar{\eta} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \\ \bar{\mu} \geq 0, g(\bar{x}) \leq 0, \bar{\mu} \perp g(\bar{x}) \\ \bar{\lambda} \geq 0, \bar{\eta} \geq 0, \bar{\lambda} \perp \bar{\eta} \end{cases} \quad \Leftrightarrow \quad (5.10).$$

where the last equality follows from $-(A\bar{x} - b) = \bar{\eta}$. \square

Corollary 5.2.4. *Let $VI(\mathcal{Q}, f)$ and $VI(\mathcal{Y}, t)$ be defined as in Proposition 5.2.3. The $\text{SOL}(\mathcal{Q}, f)$ is non-empty if and only if $VI(\mathcal{Y}, t)$ is non-empty.*

The main advantage of this reformulation is that the set \mathcal{Y} can now be decomposed in the cartesian product of the local sets \mathcal{X}^i and the positive orthant $\mathbb{R}_{\geq 0}^m$. We also note that solving the VI in (5.9) is equivalent (thanks to Assumption 5.1.1 and Proposition 5.1.1) to the following augmented game with $N + 1$ players, but no coupling constraints.

$$\mathcal{G}_{ext} = \begin{cases} \min_{x^i \in \mathcal{X}^i} & \tilde{J}^i(x^i, x^{-i}) + \lambda^\top ([A]_{(:,i)} x^i) & \forall i \in \mathbb{Z}[1, N] \\ \min_{\lambda \in \mathbb{R}_{\geq 0}^m} & -(Ax - b)^\top \lambda & i = N + 1 \end{cases} \quad (5.11)$$

where $[A]_{(:,i)}$ is the block of columns in the matrix A relative to player i , that is $A = [[A]_{(:,1)}, \dots, [A]_{(:,N)}]$.

Proposition 5.2.5 (Augmented game). *Suppose Assumption 5.1.1 holds. Then $[\bar{x}; \bar{\lambda}] \in \mathbb{R}^{Nn} \times \mathbb{R}^m$ is a Nash equilibrium of the game \mathcal{G}_{ext} if and only if $[\bar{x}; \bar{\lambda}]$ is a solution of the $VI(\mathcal{Y}, t)$ defined in (5.9). \square*

A summary of the relations presented in this section is given in Figure 5.2.

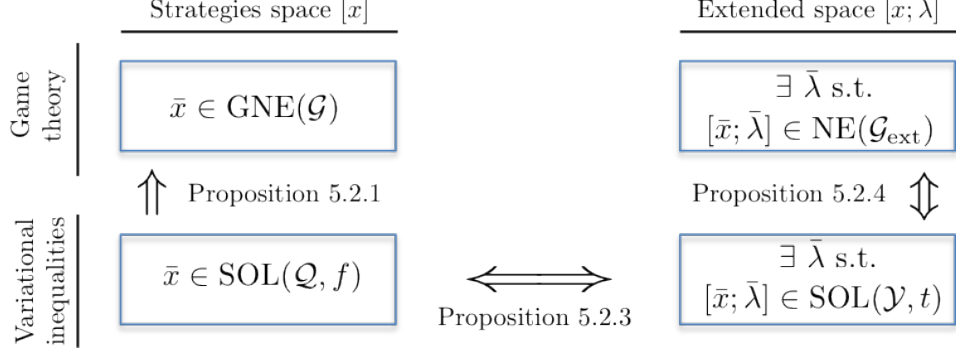


Figure 5.2: Relations between games and VIs, under Assumptions 5.1.1 and 5.2.1

In the rest of the chapter we are going to illustrate how these characterizations can be exploited to solve the coordination problem, in the case of boundedly rational agents, under the following regularity assumption.

Assumption 5.2.2 (Monotonicity). *The operator $f : \mathbb{R}^{Nn} \rightarrow \mathbb{R}^{Nn}$ as defined in (5.2) is **MON** and Lipschitz with constant L_f .*

An immediate consequence of this assumption is that the operator $t(x, \lambda)$, as defined in (5.9), is **MON** and Lipschitz as well.

Lemma 5.2.6 (Regularity of the extended operator). *Under Assumption 5.2.2, the operator $t(x, \lambda)$, as defined in (5.9), is **MON** and Lipschitz with constant $L_f + L_A$, where L_f is the Lipschitz constant of $f(x)$ and $L_A := \left\| \begin{bmatrix} 0 & A^\top \\ -A & 0 \end{bmatrix} \right\|$. \square*

The main contribution of Section 5.3 is to show that the schemes for **MON** variational inequalities, illustrated in Section 3.3.2, can be implemented in a decentralized fashion in the case of **AAGs** to solve $\text{VI}(\mathcal{Y}, t)$.

5.3 Average aggregative games: schemes with central operator

Consider the average aggregative game in (4.9) with the addition of the coupling constraint \mathcal{C} as defined in (5.6):

$$\begin{cases} \min_{x^i \in \mathbb{R}^n} & J^i(x^i, \bar{\sigma}(x)) \\ \text{s.t.} & x^i \in \mathcal{Q}^i(x^{-i}) \end{cases} \quad \forall i \in \mathbb{Z}[1, N], \quad (5.12)$$

where $\bar{\sigma}(x) = \frac{1}{N} \sum_{j=1}^N x^j$ and $\mathcal{Q}^i(x^{-i})$ is as in (5.7). Propositions 5.2.1 and 5.2.3 guarantee that any solution of the VI($\mathcal{X}_{1 \times N} \times \mathbb{R}_{\geq 0}^m, t$), with t as in (5.9), is a GNE of the game in (5.12). Thanks to Assumption 5.2.2 and Lemma 5.2.6, the operator t of such a VI is MON. Consequently, one can recover one of its solutions, and thus a generalized Nash equilibrium, by applying the iterative schemes of Section 3.3.2. We show in the next subsections how these schemes can be implemented in a decentralized fashion.

To this end, we assume the presence of a central operator that, at every iteration, computes the population average, updates the multiplier, and then broadcasts these two quantities to the agents, see Figure 5.3. Each agent, in turn, locally computes its gradient (which in the case of AAGs depends only on its strategy and on the average received by the central operator, as shown in (5.5))

$$f^i(x^i, \bar{\sigma}(x)) := \nabla_{x^i} J^i(x^i, \bar{\sigma}(x)) = \nabla_{x^i} J^i(x^i, z^i) \Big|_{z^i = \bar{\sigma}(x)} + \frac{1}{N} \nabla_{z^i} J^i(x^i, z^i) \Big|_{z^i = \bar{\sigma}(x)}$$

and consequently updates its strategy by taking a gradient step of the extended cost function in (5.11).

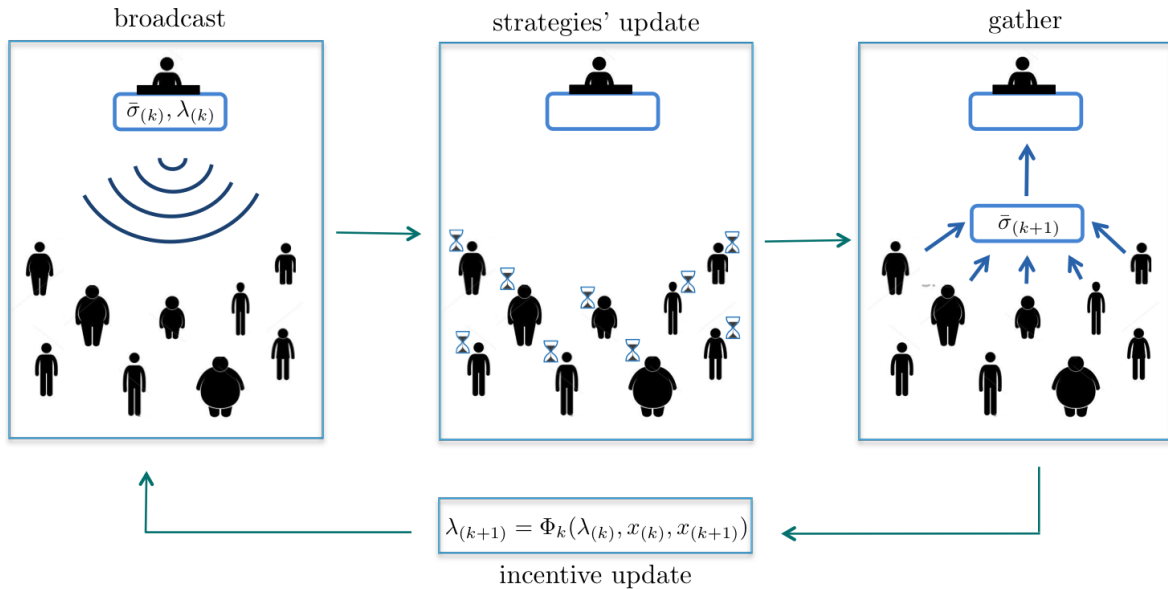


Figure 5.3: Coordination algorithm with central operator and dual variables.

5.3.1 A decentralized scheme: quadratic case

In the case of quadratic cost functions

$$J^i(x^i, z^i) := (x^i)^\top Q^i x^i + 2(Cz^i + c^i)^\top x^i, \quad (5.13)$$

where $x^i \in \mathbb{R}^n$, $Q^i = Q^{i\top} \succeq 0$, $C \in \mathbb{R}^{n \times n}$, $c^i \in \mathbb{R}^n$ and $z^i = \bar{\sigma}(x)$, the operator f in (5.2) is affine and can be explicitly characterized as follows

$$f(x) = 2 \left(\begin{bmatrix} Q^1 & & \\ & \ddots & \\ & & Q^N \end{bmatrix} + \frac{1}{N} \begin{bmatrix} C & \dots & C \\ & \ddots & \\ C & \dots & C \end{bmatrix} + \frac{1}{N} \begin{bmatrix} C^\top & & \\ & \ddots & \\ & & C^\top \end{bmatrix} \right) \begin{bmatrix} x^1 \\ \vdots \\ x^N \end{bmatrix} + 2 \begin{bmatrix} c^1 \\ \vdots \\ c^N \end{bmatrix}.$$

This structure can be exploited to derive sufficient conditions, depending only on the matrices Q^i and C , that guarantee Assumption 5.2.2.

Lemma 5.3.1 (Sufficient conditions for the quadratic case). *Suppose that the cost functions are as in (5.13), then the operator f is Lipschitz. Moreover, if $Q_i \succeq 0$ for all $i \in \mathbb{Z}[1, N]$ and $C + C^\top \succeq 0$ then f is **MON**. If additionally $C + C^\top \succ 0$ and/or $Q_i \succ 0$ for all $i \in \mathbb{Z}[1, N]$, then f is **SMON**.*

Proof. Immediate consequence of the fact that $f(x)$ is affine and of Lemma 3.1.4. \square

Since the operator f is affine, it is immediate to see that the operator t in (5.9) is affine as well. Therefore the asymmetric projection algorithm described in (3.14) can be applied to solve the $\text{VI}(\mathcal{X}_{1 \times N} \times \mathbb{R}_{\geq 0}^m, t)$, leading to Algorithm 5.

Proposition 5.3.2 (Convergence of Algorithm 5). *Suppose that Assumptions 5.1.1, 5.2.1 and 5.2.2 hold. Moreover, assume that the cost functions $J^i(x^i, \bar{\sigma}(x))$ are as in (5.13). Let τ be chosen according to Theorem 3.3.12. Then, for every initial condition $[x_{(0)}; \lambda_{(0)}]$, Algorithm 5 is guaranteed to converge to a generalized Nash equilibrium of the game in (5.12). \square*

5.3.2 A decentralized scheme: general case

Algorithm 5 works in the case of games with quadratic cost. Similarly, one can show that, for a generic operator f satisfying Assumption 5.2.2, the extragradient algorithm in (3.10) is guaranteed to converge and can again be implemented in a decentralized fashion. The only difference is that two rounds of communications among agents and central operator, as well as two primal-dual updates, are needed at every algorithmic step, see Algorithm 6.

Proposition 5.3.3 (Convergence of Algorithm 6). *Suppose that Assumptions 5.1.1, 5.2.1 and 5.2.2 hold. Let $\tau \leq \frac{1}{L_f + L_A}$, where L_A is as defined in Lemma 5.2.6. Then, for every initial condition $[x_{(0)}; \lambda_{(0)}]$, Algorithm 6 is guaranteed to converge to a generalized Nash equilibrium of the game in (5.12). \square*

Algorithm 5: AAG: Asymmetric projection algorithm (decentralized)

Initialization: Set $k \leftarrow 0$. Each agent i has initial state $x_{(0)}^i$, the central operator sets and broadcasts $\lambda_{(0)} \in \mathbb{R}_{\geq 0}^m$, $\bar{\sigma}_{(0)} = \frac{1}{N} \sum_{i=1}^N x_{(0)}^i$ and $\tau > 0$.

Iterate:

Local: strategy update

$$x_{(k+1)}^i \leftarrow \Pi_{\mathcal{X}^i} [x_{(k)}^i - \tau (f^i(x_{(k)}^i, \bar{\sigma}_{(k)}) + A_{(:,i)}^\top \lambda_{(k)})]$$

Central: multiplier and average update and broadcast

$$\begin{aligned} \lambda_{(k+1)} &\leftarrow \Pi_{\mathbb{R}_{\geq 0}^m} [\lambda_{(k)} - \tau (b - 2Ax_{(k+1)} + Ax_{(k)})] \\ \bar{\sigma}_{(k+1)} &\leftarrow \frac{1}{N} \sum_{i=1}^N x_{(k+1)}^i \end{aligned}$$

Algorithm 6: AAG: Extragradient algorithm (decentralized)

Initialization: Set $k \leftarrow 0$. Each agent i has initial state $x_{(0)}^i$, the central operator sets and broadcasts $\lambda_{(0)} \in \mathbb{R}_{\geq 0}^m$, $\bar{\sigma}_{(0)} = \frac{1}{N} \sum_{i=1}^N x_{(0)}^i$ and $\tau > 0$.

Iterate:

Local 1: strategy update

$$\tilde{x}_{(k)}^i \leftarrow \Pi_{\mathcal{X}^i} [x_{(k)}^i - \tau (f^i(x_{(k)}^i, \bar{\sigma}_{(k)}) + A_{(:,i)}^\top \lambda_{(k)})]$$

Central 1: multiplier and average update and broadcast

$$\begin{aligned} \tilde{\lambda}_{(k)} &\leftarrow \Pi_{\mathbb{R}_{\geq 0}^m} [\lambda_{(k)} - \tau (b - A\tilde{x}_{(k)})] \\ \tilde{\sigma}_{(k)} &\leftarrow \frac{1}{N} \sum_{i=1}^N \tilde{x}_{(k)}^i \end{aligned}$$

Local 2: strategy update

$$x_{(k+1)}^i \leftarrow \Pi_{\mathcal{X}^i} [x_{(k)}^i - \tau (f^i(\tilde{x}_{(k)}^i, \tilde{\sigma}_{(k)}) + A_{(:,i)}^\top \tilde{\lambda}_{(k)})]$$

Central 2: multiplier and average update and broadcast

$$\begin{aligned} \lambda_{(k+1)} &\leftarrow \Pi_{\mathbb{R}_{\geq 0}^m} [\lambda_{(k)} - \tau (b - A\tilde{x}_{(k)})] \\ \bar{\sigma}_{(k+1)} &\leftarrow \frac{1}{N} \sum_{i=1}^N x_{(k+1)}^i \end{aligned}$$

5.4 Appendix

5.4.1 Proofs of the results stated in Section 5.1

Proof of Lemma 5.1.3 (Convergence of sequential and simultaneous BR dynamics [BT97])

This proof combines known results [BT97] and is reported here for completeness. Firstly, since $\mathcal{X}_{1 \times N}$ is a cartesian product and the projection operator is NEX in the Euclidian norm [BT97, Chapter 3, Proposition 3.2 (c)] the projection operator $\Pi_{\mathcal{X}_{1 \times N}}(\cdot)$ is NEX in the block-maximum norm

$$\|\Pi_{\mathcal{X}_{1 \times N}}(x) - \Pi_{\mathcal{X}_{1 \times N}}(y)\|_B = \max_i \|\Pi_{\mathcal{X}^i}(x^i) - \Pi_{\mathcal{X}^i}(y^i)\| \leq \max_i \|x^i - y^i\| = \|x - y\|_B.$$

By assumption $R(x)$ is a block contraction. Consequently, the operator $O(x) := \Pi_{\mathcal{X}_{1 \times N}}(x - \alpha f(x)) := \Pi_{\mathcal{X}_{1 \times N}}(R(x))$ is a block contraction. The operator $O(x)$ describes one step of the projection algorithm, which is a linearized algorithm for solving the VI($\mathcal{X}_{1 \times N}, f$) with scaling matrices $A(x) = I/\alpha$. Consequently, by [BT97, Chapter 3, Proposition 5.12], VI($\mathcal{X}_{1 \times N}, f$) has a unique solution and for any initial condition $x_{(0)}$ both the sequential and simultaneous BR dynamics (i.e., the nonlinear Gauss-Seidel and Jacobi algorithms) converge to it, geometrically. The conclusion then follows from Proposition 5.1.1.

Proof of Lemma 5.1.4 (Sufficient conditions for block contraction in the quadratic case)

To prove this result we use the sufficient conditions given in [BT97, Chapter 3, Proposition 1.12] with $G_i = I_n$ for all $i \in \mathbb{Z}[1, N]$. Note that in the quadratic case $f_i(x) = 2D_i x^i + \frac{2}{N} C \sum_{j \neq i} x^j + 2c^i$, where we defined $D_i := (q_i Q + \frac{C+C^T}{N})$, and $D_i \succeq 0$ by Assumption 5.1.1. Then

$$\begin{aligned} \|f_i(x) - f_i(y)\| &\leq 2\|D_i\| \|x^i - y^i\| + \frac{2}{N} \|C\| \sum_{j \neq i} \|x^j - y^j\| \\ &\leq 2(\|D_i\| + \frac{N-1}{N} \|C\|) \|x - y\|_B \leq A_1 \|x - y\|_B, \end{aligned}$$

with $A_1 := 2 \left(\max_{i=1, \dots, N} \{ \|D_i\| \} + \frac{N-1}{N} \|C\| \right) > 0$. Moreover,

$$\begin{aligned}
(f_i(x) - f_i(y))^\top (x^i - y^i) &= (2D_i(x^i - y^i) + \frac{2}{N}C \sum_{j \neq i} (x^j - y^j))^\top (x^i - y^i) \\
&= 2(x^i - y^i)^\top D_i(x^i - y^i) + \frac{2}{N} \sum_{j \neq i} (x^j - y^j)^\top C^\top (x^i - y^i) \\
&\geq 2\sigma_{\min}(D_i) \|x^i - y^i\|^2 - \frac{2}{N} \sigma_{\max}(C) \|x^i - y^i\| \sum_{j \neq i} \|x^j - y^j\| \\
&\geq 2\sigma_{\min}(D_i) \|x^i - y^i\|^2 - 2\frac{N-1}{N} \sigma_{\max}(C) \|x - y\|_B^2 \\
&\geq 2 \left[\sigma_{\min}(q_i Q) - \sigma_{\max}\left(\frac{C+C^\top}{N}\right) \right] \|x^i - y^i\|^2 - 2\frac{N-1}{N} \sigma_{\max}(C) \|x - y\|_B^2 \\
&\geq 2 \left[\min_{i=1, \dots, N} \{q_i\} \sigma_{\min}(Q) - \sigma_{\max}\left(\frac{C+C^\top}{N}\right) \right] \|x^i - y^i\|^2 - 2\frac{N-1}{N} \sigma_{\max}(C) \|x - y\|_B^2 \\
&:= A_2 \|x^i - y^i\|^2 - A_3 \|x - y\|_B^2,
\end{aligned}$$

where we use Cauchy-Schwarz in the first inequality. Let us define $\alpha := \beta \frac{N+1}{N}$ and note that $\alpha < 1$ since $N > \frac{\beta}{1-\beta}$. By assumption $\sigma_{\min}(q_i Q) \geq \frac{1}{\beta} \sigma_{\max}(C)$ for all $i \in \mathbb{Z}[1, N]$, therefore

$$\begin{aligned}
\min_{i=1, \dots, N} \{q_i\} \sigma_{\min}(Q) &\geq \frac{1}{\beta} \sigma_{\max}(C) = \frac{1}{\alpha} \frac{N+1}{N} \sigma_{\max}(C) = \frac{1}{\alpha} \frac{N-1}{N} \sigma_{\max}(C) + \frac{1}{\alpha} \frac{2}{N} \sigma_{\max}(C) \\
&> \frac{1}{\alpha} \frac{N-1}{N} \sigma_{\max}(C) + \frac{2}{N} \sigma_{\max}(C) = \frac{1}{\alpha} \frac{N-1}{N} \sigma_{\max}(C) + \sigma_{\max}\left(\frac{C}{N}\right) + \sigma_{\max}\left(\frac{C^\top}{N}\right) \\
&\geq \frac{1}{\alpha} \frac{N-1}{N} \sigma_{\max}(C) + \sigma_{\max}\left(\frac{C+C^\top}{N}\right).
\end{aligned}$$

Consequently,

$$\frac{A_3}{2} = \frac{N-1}{N} \sigma_{\max}(C) < \alpha \left(\min_{i=1, \dots, N} \{q_i\} \sigma_{\min}(Q) - \sigma_{\max}\left(\frac{C+C^\top}{N}\right) \right) < \frac{A_2}{2}$$

and $A_3 < A_2$.

5.4.2 Proofs of the results stated in Section 5.2

Proof of Proposition 5.2.5 (Augmented game)

Let $x^{N+1} := \lambda$, $\tilde{J}^i(x^i, x^{-i}) := \tilde{J}^i(x^i, x^{-i}) + (x^{N+1})^\top ([A]_{(:,i)} x^i)$ for all $i \in \mathbb{Z}[1, N]$ and $\tilde{J}^{N+1}(x^{N+1}, x^{-(N+1)}) := -(A[x^1; \dots; x^N] - b)^\top x^{N+1}$. The statement can be proven as in Proposition 5.1.1, for the game with $N+1$ players and cost \tilde{J} , upon noticing that the minimum principle holds when the constraint set is closed and convex but not necessarily bounded. Therefore the constraint set $\mathcal{X}^{N+1} := \mathbb{R}_{\geq 0}^m$ is admissible.

Proof of Lemma 5.2.6 (Regularity of the extended operator)

The operator $t(x, \lambda)$ is the sum of the two operators $\begin{bmatrix} x \\ \lambda \end{bmatrix} \mapsto \begin{bmatrix} f(x) \\ 0 \end{bmatrix}$ and $\begin{bmatrix} x \\ \lambda \end{bmatrix} \mapsto \begin{bmatrix} 0 & A^\top \\ -A & 0 \end{bmatrix} \begin{bmatrix} x \\ \lambda \end{bmatrix} + \begin{bmatrix} 0 \\ b \end{bmatrix}$. The former operator is **MON** and Lipschitz by assumption. The latter is **MON** and Lipschitz with constant L_A because it is affine (see Lemma 3.1.4). Consequently, $t(x, \lambda)$ is **MON** and Lipschitz as well.

5.4.3 Proofs of the results stated in Section 5.3

Proof of Proposition 5.3.2 (Convergence of Algorithm 5)

Algorithm 5 can be rewritten in condensed form as

$$y^{(k+1)} := \begin{bmatrix} x^{(k+1)} \\ \lambda^{(k+1)} \end{bmatrix} = \begin{bmatrix} \Pi_{\mathcal{X}_{1 \times N}} [x^{(k)} - \tau (f(x^{(k)}) + A^\top \lambda^{(k)})] \\ \Pi_{\mathbb{R}_{\geq 0}^m} [\lambda^{(k)} - \tau (b - 2Ax^{(k+1)} + Ax^{(k)})] \end{bmatrix}$$

which coincides with one step of the asymmetric projection algorithm given in (3.14) for the $\text{VI}(\mathcal{X}_{1 \times N} \times \mathbb{R}_{\geq 0}^m, t) = \text{VI}(\mathcal{Y}, t)$. Note that $\text{SOL}(\mathcal{X}, f)$ is non-empty because \mathcal{X} is convex and compact and f is continuous. Therefore, by Corollary 5.2.4, $\text{SOL}(\mathcal{Y}, t)$ is non-empty as well. Theorem 3.3.12 then guarantees convergence of Algorithm 5 to a solution $[\bar{x}; \bar{\lambda}]$ of $\text{VI}(\mathcal{X}_{1 \times N} \times \mathbb{R}_{\geq 0}^m, t)$. Consequently, by Proposition 5.2.3, \bar{x} solves $\text{VI}(\mathcal{Q}, f)$ and is a variational GNE for the generalized AAG by Proposition 5.2.1.

Proof of Proposition 5.3.3 (Convergence of Algorithm 6)

Algorithm 6 coincides with the extragradient algorithm for the $\text{VI}(\mathcal{X}_{1 \times N} \times \mathbb{R}_{\geq 0}^m, t) = \text{VI}(\mathcal{Y}, t)$. Note that $\text{SOL}(\mathcal{X}, f)$ is non-empty because \mathcal{X} is convex and compact and f is continuous. Therefore, by Corollary 5.2.4, $\text{SOL}(\mathcal{Y}, t)$ is non-empty as well. Consequently, Proposition 3.3.10 guarantees convergence of Algorithm 6 to a solution $[\bar{x}; \bar{\lambda}]$ of $\text{VI}(\mathcal{X}_{1 \times N} \times \mathbb{R}_{\geq 0}^m, t)$. Consequently, by Proposition 5.2.3, \bar{x} solves $\text{VI}(\mathcal{Q}, f)$ and is a variational GNE for the generalized AAG by Proposition 5.2.1.

Part II

Populations of biological systems: Controlled stochastic biochemical reaction networks

Introduction

In this second part of the thesis we move our attention from populations of rational agents to populations of systems whose behavior can be described using a chemical reaction network. In full generality, a reaction network is a system comprising different “species” that interact by means of a series of “reactions” whose effect is to change the number of copies per species present in the system. Some example of applications are as follows.

- *Biochemical networks*: biochemical networks describe the evolution of chemical species inside a cell [Wil11]. A classical example, which we use as benchmark for the subsequent theory, is the gene expression process in which two species, the mRNA and the corresponding protein, interact by means of four reactions. The mRNA count is increased by transcription events (in which DNA is copied into mRNA, usually at a constant rate), the protein count is increased by translation events (in which the protein is synthesized from the mRNA copies, at a rate that is proportional to the mRNA counts) and, finally, both species are degraded by the cell, either actively or because of cell division events.
- *Pharmacokinetic networks*: pharmacokinetic models are used to study the absorption, distribution and elimination of drugs in the body. Traditionally, pharmacokinetic models aim at describing how a drug under investigation propagates in different compartments (e.g. heart, lungs, liver, brain, etc.) [Jac96]. Such compartmental models can be casted in the framework of reaction networks by associating a species to each compartment and by modeling the fluxes among different compartments as reaction events.
- *Epidemiological networks*: epidemiological networks study the spread of an infection through a population [AM92, NPP16]. Numerous publications can be found in this context both in a deterministic and stochastic setting. To illustrate the connection with reaction networks we consider the standard SIR model [MA79]. Therein, the population is divided in three classes (i.e. species): susceptible (S), infected (I) and resistant (R). Because of mutual interactions individuals that are susceptible can become infected and infected people can recover and thus become

resistant. These events can be model as reactions that happen with a rate dependent on the amount of infected people and on the infection virulence, respectively.

- *Ecological networks*: ecological networks aim at studying the interaction among different organisms living in an ecosystem. The most common model is the Lotka-Volterra model where two species, a predator and a prey, compete for survival. Also in this case the amount of predators and preys in the ecosystem can be modeled as the result of a series of reactions modeling birth and death events.

Even though many of the results developed in the following can be applied to any of the contexts described above, [GJ13, PLR15], in the rest of the thesis we focus on biochemical reaction networks. Consequently, we refer to the systems as cells and to their internal components as species.

Biochemical reaction networks have been traditionally studied using a deterministic approach according to which the amount of species (or their concentration) inside a single cell can be described as the solution of a deterministic ordinary differential equation. This approach leads to models that are computationally very efficient and has been successfully applied in many different contexts. Recent technological developments have however allowed researchers to simultaneously observe the amount of species inside thousands of cells of the same population and have shown that, even if each cell starts with the same initial configuration and is exposed to the same environment, the amount of species inside different cells of the same population can vary significantly [SS08]. These observations suggest that deterministic approaches may not always be a valid choice. This statement is particularly relevant for biochemical networks involving species present in low copy numbers, as for example the mRNA in gene expression systems. In this case, in fact, even small fluctuations may become predominant in influencing the cell behavior. Popular examples of such phenomena are stochastic switching systems, as the Enterobacteria phage λ [ARM98]. For these systems it has been observed that cells of the same population can undertake, with a certain probability, either one of two very different fates, leading to a phenomenon known as population partitioning. It has been hypothesized that population partitioning is an evolutionary mechanism adopted by cells to survive in uncertain and mutating environments (so that part of the population is always in the condition to survive) or to achieve division of labor (so that cells can perform mutually exclusive task and increase the benefit of the overall population).

The pioneering work of N.G. van Kampen and D.T. Gillespie provides the fundamental analytical and computational methods for incorporating stochasticity in the analysis of biochemical reaction networks [VK07, Gil07]. Their basic observation is that each reaction inside a cell can take place only if the necessary reactants come sufficiently close and this is a stochastic event, governed by the laws of statistical mechanics. Consequently, reaction events are stochastic and the amount of species inside each cell of a given population should be seen as a different realization of a (Markovian) stochastic

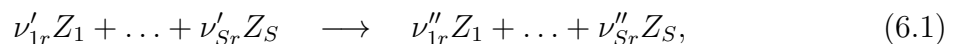
process. In the rest of the thesis we use the master equation describing the probability evolution of this process and its moments (e.g., the average amount of species, its variance etc.) to concisely describe and control the population behavior.

A fundamental assumption that we make is that the rate at which (some of the) reactions happen inside each cell can be influenced by means of an external control signal, which is applied globally to the whole population. Such a signal could for example be a light pulse, as done in [MASSO⁺11, OHL⁺14], or a concentration signal, as in [UMD⁺12, MDBDB11]. It is important to stress that, contrary to the first part of the thesis, we assume here that the cells do not interact one with the others, but are actually independent realizations of the same (controlled) stochastic process. Moreover, we assume that we can observe the cells behavior by means of population measurements, that is, measurements of the average amount of species (and possibly of the variance or higher order moments) in a sample of the population at different time instants.

Our first objective is to propose a systematic procedure to characterize these systems (i.e., to construct a model of the underlying stochastic process) from population measurements by “optimally” perturbing the system with the available external signal. We then use the identified model to investigate the range of behaviors that the population can exhibit under different choices of the external input and to design external inputs that control the population to a desired state.

6.1 The stochastic biochemical reaction network framework

A biochemical reaction network is a system comprising S molecular species Z_1, \dots, Z_S that interact through R reactions. A typical reaction $r \in \{1, 2, \dots, R\}$ can be expressed as



where $\nu'_{1r}, \dots, \nu'_{Sr} \in \mathbb{N}$ and $\nu''_{1r}, \dots, \nu''_{Sr} \in \mathbb{N}$ are the coefficients that determine how many molecules for each species are needed for the reaction to happen and how many are produced as a consequence. The net effect of each reaction can therefore be summarized with a vector $\nu_r \in \mathbb{N}^S$, whose components are $\nu''_{sr} - \nu'_{sr}$ for $s = 1, \dots, S$, which is called the *stoichiometric vector*. In the following we say that a reaction is of order k if it involves k reactants (i.e., $\sum_{s=1}^S \nu'_{sr} = k$). A standard assumption in the analysis of biochemical reaction networks is as follows.

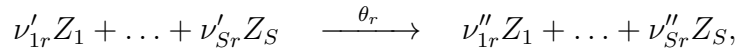
Assumption 6.1.1. *The system is well-stirred, in thermal equilibrium and with constant volume.*

It was proven by Gillespie [Gil92] that, under Assumption 6.1.1, each reaction r is a

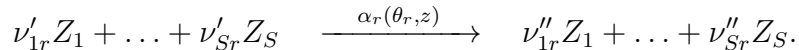
stochastic event that happens, in the infinitesimal interval $[t, t + dt]$ with probability

$$\alpha_r(\theta_r, z)dt := \theta_r \cdot h_r(z) \cdot dt, \quad (6.2)$$

where $z = [z_1, \dots, z_S]^\top$ denotes the amount of molecules per species that are present in the system at time t , $h_r(z) := \prod_{s=1}^S \binom{z_s}{\nu'_{sr}}$ is a function that counts in how many different ways one can select the required amount of reactants from the available molecules z and $\theta_r \in \mathbb{R}_{\geq 0}$ is a (usually unknown) parameter depending on binding affinities, temperature, volume, etc. In the following we refer to the time-varying quantity $\alpha_r(\theta_r, z)$ as the *propensity* of reaction r and to the constant parameter θ_r as the *reaction rate*. The formula for the propensity derived by Gillespie in [Gil92] and given in (6.2) is known as *mass action* and is based on the laws of statistical mechanics. In the following, we use the short-hand notation



to denote a reaction that follows the mass action kinetics with rate parameter θ_r . When modeling complex system, it might be useful to simplify a reaction network by lumping together several reactions or by abstracting some external influencing factor. This operation usually leads to a condensed reaction of the form given in (6.1), but with a propensity that does not follow the laws of mass action kinetics. For example, enzyme reactions are usually modeled using Michaelis-Menten kinetics [Wil11, Section 7.3]. If propensities different from mass action are allowed we use the extended notation



Let $Z(t) = [Z_1(t), \dots, Z_S(t)]^\top$ be the vector describing the number of molecules present in the network for each species at time t , that is, the state of the network at time t . As a consequence of the fact that reactions are stochastic, $Z(t)$ is a stochastic process. Moreover, since the propensity of each reaction depends only on the current state of the system, the process $Z(t)$ is Markovian. In the following, we always use the upper case to denote a process and the lower case to denote its realizations. For example, $z = [z_1, \dots, z_S]^\top$ denotes a particular realization of the state $Z(t)$ of the stochastic system at time t . Let $p(t, z) = \mathbb{P}[Z(t) = z]$ be the probability that the realization of the process Z at time t is z . It is proven in [Gil92] that

$$p(z, t + dt) = \underbrace{p(z, t) \left(1 - \sum_{r=1}^R \alpha_r(\theta_r, z) dt\right)}_{\text{term 1}} + \underbrace{\sum_{r=1}^R p(z - \nu_r, t) \alpha_r(\theta_r, z - \nu_r) dt}_{\text{term 2}} + o(dt).$$

The previous formula can be easily understood once we note that the probability of being in state z at time $t + dt$ is given by the sum of the probabilities of two mutually exclusive events: either the system was already in state z at time t and no reaction happened in

the interval $[t, t + dt]$ (term 1) or the system was not in z and the right reaction happened (term 2). Finally, the term $o(dt)$ is due to the fact that in the previous discussion we did not consider the case that more than one reaction could happen in the interval $[t, t + dt]$. By rearranging the terms and taking the limit as $dt \rightarrow 0$ one arrives at what is known as the **chemical master equation (CME)**

$$\dot{p}(z, t) = \sum_{r=1}^R [p(z - \nu_r, t)\alpha_r(\theta_r, z - \nu_r) - p(z, t)\alpha_r(\theta_r, z)]. \quad (6.3)$$

Note that typical biochemical reaction networks involve many different species, whose counts can theoretically grow unbounded. Consequently, $z \in \mathbb{N}^S$ and the **CME** in (6.3) is a system of infinitely many coupled ordinary differential equations that cannot be solved in a straightforward way, even for very simple systems. A large number of different methods, both analytical and computational, have been proposed in the literature to either solve or approximate the **CME**. We refer to [Wil11, GJ13, Rue14] for a comprehensive review. In the following we limit our discussion to the two approximate methods that are going to be used in the thesis. To better illustrate these methods we consider as running example the following renowned model of gene expression.

Example 6.1 (Gene expression reaction network). *Consider a biochemical network consisting of two species, the mRNA (M) and the corresponding protein (P), and the following reactions*



where k_r, k_p are the mRNA and protein production rate and γ_r, γ_p are the mRNA and protein degradation rate, respectively. Then $Z(t) = [M(t), P(t)]^\top$, $z = [m, p]^\top$, $\theta = [k_r, \gamma_r, k_p, \gamma_p]^\top$ and the stoichiometric matrix is

$$\nu := [\nu_1, \nu_2, \nu_3, \nu_4] = \begin{bmatrix} 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & -1 \end{bmatrix}.$$

The empty set notation is used whenever a certain species is produced or degrades without involving the other species. The **CME** for this system is

$$\begin{aligned} \dot{p}\left(\begin{bmatrix} m \\ p \end{bmatrix}, t\right) = & -\left(\alpha_1\left(k_r, \begin{bmatrix} m \\ p \end{bmatrix}\right) + \alpha_2\left(\gamma_r, \begin{bmatrix} m \\ p \end{bmatrix}\right) + \alpha_3\left(k_p, \begin{bmatrix} m \\ p \end{bmatrix}\right) + \alpha_4\left(\gamma_p, \begin{bmatrix} m \\ p \end{bmatrix}\right)\right) p\left(\begin{bmatrix} m \\ p \end{bmatrix}, t\right) + \\ & \alpha_1\left(k_r, \begin{bmatrix} m-1 \\ p \end{bmatrix}\right) p\left(\begin{bmatrix} m-1 \\ p \end{bmatrix}, t\right) + \alpha_2\left(\gamma_r, \begin{bmatrix} m+1 \\ p \end{bmatrix}\right) p\left(\begin{bmatrix} m+1 \\ p \end{bmatrix}, t\right) + \alpha_3\left(k_p, \begin{bmatrix} m \\ p-1 \end{bmatrix}\right) p\left(\begin{bmatrix} m \\ p-1 \end{bmatrix}, t\right) + \alpha_4\left(\gamma_p, \begin{bmatrix} m \\ p+1 \end{bmatrix}\right) p\left(\begin{bmatrix} m \\ p+1 \end{bmatrix}, t\right). \end{aligned} \quad (6.4)$$

In the case of mass action kinetics the propensities $\alpha_r(\theta_r, z)$ can be further specified as

$$\alpha_1(k_r, z) = k_r, \quad \alpha_2(\gamma_r, z) = \gamma_r \cdot m, \quad \alpha_3(k_p, z) = k_p \cdot m, \quad \alpha_4(\gamma_p, z) = \gamma_p \cdot p. \quad (6.5)$$

6.1.1 Solving the CME: finite state projection

Let us introduce an ordering $\{z^j\}_{j=1}^\infty$ of the possible state realizations $z \in \mathbb{N}^S$. For the system in Example 6.1, we could for instance use the mapping

$$\begin{aligned} z^1 &= [0, 0]^\top, \quad z^2 = [1, 0]^\top, \quad z^3 = [0, 1]^\top, \quad z^4 = [2, 0]^\top, \\ z^5 &= [1, 1]^\top, \quad z^6 = [0, 2]^\top, \quad z^7 = [3, 0]^\top, \quad z^8 = [2, 1]^\top, \quad \dots \end{aligned}$$

where $[m, p]^\top$ denotes the state with m mRNA copies and p proteins (see Figure 6.1).

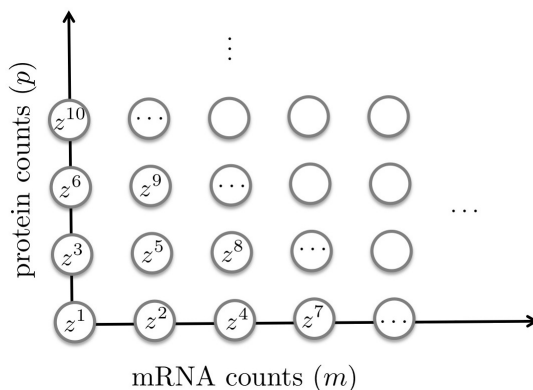


Figure 6.1: State space for the gene expression system of Example 6.1.

Setting¹ $P_j(t) := p(z^j, t)$, the CME in (6.3) can be rewritten as an infinite dimensional linear system,

$$\dot{P}(t) = FP(t), \tag{6.6}$$

where $P(t) \in [0, 1]^\infty$ and the matrix $F \in \mathbb{R}^{\infty \times \infty}$ can be derived from the reaction network and its propensities. Specifically, the element $[F]_{i,j}$ is the propensity of transitioning from state z^j to state z^i (because a reaction with stoichiometric vector $\nu_r = z^j - z^i$ happens) and $-[F]_{i,i}$ is the propensity of leaving state z^i (because any of the R reactions happens). Consequently, system (6.6) can be thought of as a Markov chain with countably many states $z^j \in \mathbb{N}^S$ and constant transition matrix F . Note that, by construction, all the off-diagonal components of F are nonnegative and thus F is, by definition, a Metzler matrix. The full derivation can be found in [MK06].

The main idea of the **finite state projection (FSP)** technique [MK06] is that the probability of visiting most of the states of the Markov chain (6.6) is practically negligible. Therefore, system (6.6) can be approximated by removing these states and keeping track only of the probability evolution for the remaining ones, that take index in a suitable

¹Not to be confused with the symbol used to denote the amount of protein in Example 6.1.

set J . For example in gene expression, even though mathematically the mRNA could grow unboundedly, it is very unlikely to observe states with more than tens of copies. To formalize this concept let us first define the reduced order system

$$\dot{\bar{P}}_J(t) = [F]_J \bar{P}_J(t), \quad \bar{P}_J(0) = P_J(0), \quad (6.7)$$

where $P_J(0)$ is the subvector of $P(0)$ corresponding to the indices in J , and $[F]_J$ denotes the sub matrix of F obtained by selecting only the rows and columns with indices in J . From now on, we denote by $P(T)$ and $\bar{P}_J(T)$ the solutions at time T of system (6.6) and system (6.7), respectively. The dependence on the initial conditions $P(0)$ and $P_J(0)$ is omitted to keep the notation compact. Intuitively, the truncated system (6.7) is a good approximation of the original system (6.6) if most of the probability mass lies in J . This intuition can be formalized as follows.

Assumption 6.1.2. *For a given finite set of state indices J , an initial condition $P_J(0)$, a given tolerance $\varepsilon > 0$ and a finite instant $T > 0$, it holds*

$$\mathbb{1}^\top \bar{P}_J(T) \geq 1 - \varepsilon.$$

□

Under Assumption 6.1.2, it is possible to show that, up to time T , the evolution of the reduced Markov chain $\bar{P}_J(t)$ is a good approximation of the original, infinite dimension, Markov chain $P(t)$.

Proposition 6.1.1 (Finite State Projection [MK06]). *If Assumption 6.1.2 holds, then*

$$P_j(T) \geq \bar{P}_j(T), \quad \forall j \in J \quad \text{and} \quad \|P_J(T) - \bar{P}_J(T)\|_1 \leq \varepsilon.$$

The FSP approach is a quite efficient approach for small-size biochemical reaction networks, however it suffers from two fundamental drawbacks.

1. Since in the original Markov chain there are transitions from states that are in the set J to states that have been discarded, asymptotically $\bar{P}_J(t) \rightarrow 0$. Therefore, Assumption 6.1.2 will not be satisfied for large T and the FSP method cannot be used to analyze the asymptotic behavior of the system;
2. For a given final time T the number of states that need to be tracked (i.e., the size of the set J) in order to guarantee Assumption 6.1.2 grows exponentially in the number of species S . Therefore, the FSP method becomes computationally intractable for large biochemical networks.

Extensions and modifications of the FSP approach that partially address these issues are discussed for example in [RMASL11], where a Kalman filter approach is suggested, or in [MK07, HMW09], where a time-varying set $J(t)$, for $t \in [0, T]$, is proposed.

6.1.2 Solving the CME: moment equations

In the previous section we have seen that, for large biochemical networks, it is usually impossible to solve the CME and consequently, it is not possible to characterize the full evolution of the probability distribution. In many practical applications however, the evolution of some low-order moments of the stochastic process $Z(t)$ is informative enough to characterize the system. In [ZRK⁺12], for example, the authors suggest a method to infer the value of the rate parameter θ using only mean and variance data. Equations describing the time evolution of the *uncentered* moments of $Z(t)$ can be easily computed from the CME, in the case of mass-action kinetics, by multiplying both sides of (6.3) by (products of) the components of z and summing over all $z \in \mathbb{N}^S$ [Rue14].

Example 6.1 (cont.) *For the gene expression system of Example 6.1, the evolution of the mean of the mRNA, $\mathbb{E}[M(t)]$, can be computed as follows*

$$\begin{aligned}
 \dot{\mathbb{E}}[M] &= \frac{d}{dt} \sum_{m=0}^{\infty} \sum_{p=0}^{\infty} m \cdot \mathbb{P} \left(\begin{bmatrix} m \\ p \end{bmatrix}, t \right) = \sum_{m=0}^{\infty} \sum_{p=0}^{\infty} m \cdot \dot{\mathbb{P}} \left(\begin{bmatrix} m \\ p \end{bmatrix}, t \right) \\
 &= \sum_{m=0}^{\infty} \sum_{p=0}^{\infty} \left[-m(k_r + \gamma_r m + k_p m + \gamma_p p) \mathbb{P} \left(\begin{bmatrix} m \\ p \end{bmatrix}, t \right) + k_r m \mathbb{P} \left(\begin{bmatrix} m-1 \\ p \end{bmatrix}, t \right) \right. \\
 &\quad \left. + \gamma_r m(m+1) \mathbb{P} \left(\begin{bmatrix} m+1 \\ p \end{bmatrix}, t \right) + k_p m^2 \mathbb{P} \left(\begin{bmatrix} m \\ p-1 \end{bmatrix}, t \right) + \gamma_p m(p+1) \mathbb{P} \left(\begin{bmatrix} m \\ p+1 \end{bmatrix}, t \right) \right] \\
 &= \sum_{m=0}^{\infty} \sum_{p=0}^{\infty} [-mk_r - \gamma_r m^2 - k_p m^2 - \gamma_p m p + k_r(m+1) + \gamma_r(m-1)m + k_p m^2 + \gamma_p m p] \cdot \mathbb{P} \left(\begin{bmatrix} m \\ p \end{bmatrix}, t \right) \\
 &= \sum_{m=0}^{\infty} \sum_{p=0}^{\infty} [k_r - \gamma_r m] \mathbb{P} \left(\begin{bmatrix} m \\ p \end{bmatrix}, t \right) = k_r - \gamma_r \mathbb{E}[M],
 \end{aligned}$$

where we used the CME expression derived in (6.4) with propensities given by the mass action kinetics in (6.5). \square

In general, the resulting system of moment equations can be written as

$$\dot{\tilde{x}}_{\infty}(t) = \tilde{A}_{\infty}(\theta) \tilde{x}_{\infty}(t), \quad (6.8)$$

where $\tilde{x}_{\infty}(t)$ contains all the uncentered moments of the probability distribution $\mathbb{P}[Z(t)]$ and $\tilde{A}_{\infty} \in \mathbb{R}^{\infty \times \infty}$ is an infinite dimensional matrix whose elements depend on the rate parameters vector $\theta := [\theta_1; \dots; \theta_R]$. From (6.8) it is possible to derive a (usually nonlinear) system describing the evolution of the *centered* moments $x_{\infty}(t)$

$$\dot{x}_{\infty}(t) = f_{\infty}(\theta, x_{\infty}(t)), \quad (6.9)$$

see for example [LKK09]. Note that the solution of either (6.8) or (6.9) completely describes the stochastic process $Z(t)$. It is therefore not surprising that, as for the CME, these systems are generally very difficult, if not impossible, to solve. Luckily, in most cases the evolution of some low-order moments is sufficiently informative to characterize

the system. Let $x_{\leq l}(t)$ be the vector of desired centered moments up to order l . Note that from (6.9) it is possible to extract the time derivative of the moments of interest, but this will in general depend also on the higher order moments $x_{>l}(t)$, that is

$$\dot{x}_{\leq l}(t) = g(\theta, x_{\leq l}(t), x_{>l}(t)). \quad (6.10)$$

An explicit formula for g is given in [LKK09] as a function of the propensities and their derivatives. This formula significantly simplifies under the assumption of affine propensity functions, as illustrated in the next proposition for the first two moments ($l = 2$).

Assumption 6.1.3. *For each reaction $r = 1, \dots, R$ the propensity $\alpha_r(\theta_r, z)$ is an affine function of z .*

Proposition 6.1.2 ([LKK09]). *Consider the CME given in (6.3) and suppose that Assumption 6.1.3 holds. Let $x_i(t) = \mathbb{E}[Z_i(t)]$, $x_{\leq 1}(t) = [x_1(t); \dots; x_S(t)]$ and $x_{ij}(t) = \mathbb{E}[(Z_i(t) - x_i(t))(Z_j(t) - x_j(t))]$. Then, for all $i, j \in \mathbb{Z}[1, S]$,*

$$\begin{aligned} \dot{x}_i &= \sum_{r=1}^R [\nu_{ir} \cdot \alpha_r(\theta_r, x_{\leq 1})] \\ \dot{x}_{ij} &= \sum_{r=1}^R \left[\nu_{ir} \cdot \sum_{l=1}^S \left(\frac{\partial \alpha_r(\theta_r, x_{\leq 1})}{\partial z_l} x_{jl} \right) + \nu_{jr} \cdot \sum_{l=1}^S \left(\frac{\partial \alpha_r(\theta_r, x_{\leq 1})}{\partial z_l} x_{il} \right) + \nu_{ir} \nu_{jr} \cdot \alpha_r(\theta_r, x_{\leq 1}) \right] \end{aligned} \quad (6.11)$$

where we omitted the time dependence for simplicity.

The most important feature of the equations in (6.11) is that the evolution of mean and variance are *closed*, in the sense that they do not depend on higher order moments,² that is

$$\dot{x}_{\leq l}(t) = f(\theta, x_{\leq l}(t)). \quad (6.12)$$

Therefore computing the evolution of the low order moments reduces to solving a system of finitely many non-linear first order differential equations. We note that, in the case of mass action kinetics, Assumption 6.1.3 is satisfied if the network contains reactions up to first order (i.e, reactions with at most one reagent). In fact if a reaction is of zero order its propensity is constant (i.e., $\alpha_r(\theta_r, z) = \theta_r$) and if a reaction is of order one then its propensity is linear (i.e., $\alpha_r(\theta_r, z) = \theta_r z_{i_r}$, for some $i_r \in \mathbb{Z}[1, S]$). Assumption 6.1.3 is violated, on the other hand, for mass action reactions of order greater than one. For example, for order two the propensity becomes $\alpha_r(\theta_r, z) = \theta_r z_{i_r} z_{j_r}$, for some $i_r, j_r \in \mathbb{Z}[1, S]$ which is quadratic and thus non-affine. If the propensities are not affine approximate methods to solve (6.10) by substituting the unknown high order moments $x_{>l}(t)$ with a non-linear function of the low order ones, that is, by imposing $x_{>l}(t) \cong h(x_{\leq l}(t))$ have

²Note that this statement is true also for $l > 2$. In general, under Assumption 6.1.3 the evolution of the moments $x_{\leq l}$, at any order l , does not depend on $x_{>l}$.

been suggested in the literature, under the name of *moment closure techniques*, resulting in the approximate dynamics

$$\dot{x}_{\leq l}(t) \cong g(\theta, x_{\leq l}(t), h(x_{\leq l}(t))) =: f(\theta, x_{\leq l}(t)).$$

Unfortunately, in this case there are no theoretical guarantees or known bounds on the approximation error. Consequently, we do not delve more into this topic and we refer the interested reader to [Hes08].

Example 6.1 (cont.) *Under the assumption of mass-action kinetics, the propensity functions of the gene expression model of Example 6.1 are affine, as shown in (6.5). Therefore, if the state of system (6.12) is ordered as $x_{\leq 2} := [x_M, x_P, x_{M^2}, x_{MP}, x_{P^2}]^\top$, then $\dot{x}_{\leq 2} = Ax_{\leq 2} + b$, where*

$$A = \begin{bmatrix} -\gamma_r & 0 & 0 & 0 & 0 \\ k_p & -\gamma_p & 0 & 0 & 0 \\ \gamma_r & 0 & -2\gamma_r & 0 & 0 \\ 0 & 0 & k_p & -(\gamma_r + \gamma_p) & 0 \\ k_p & \gamma_p & 0 & 2k_p & -2\gamma_p \end{bmatrix}, \quad b = \begin{bmatrix} k_r \\ 0 \\ k_r \\ 0 \\ 0 \end{bmatrix}. \quad (6.13)$$

Note that, since the mRNA production follows a birth and death process, $M(t)$ has a Poisson distribution and consequently $x_M = x_{M^2}$. Exploiting this information the 5-dimensional system in (6.13) can be reduced to the equivalent 4-dimensional system $\dot{x} = A_r x + b_r$, with matrices

$$A_r = \begin{bmatrix} -\gamma_r & 0 & 0 & 0 \\ k_p & -\gamma_p & 0 & 0 \\ k_p & 0 & -(\gamma_r + \gamma_p) & 0 \\ k_p & \gamma_p & 2k_p & -2\gamma_p \end{bmatrix}, \quad b_r = \begin{bmatrix} k_r \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

and state $x := [x_M, x_P, x_{MP}, x_{P^2}]^\top$. In the following we will use both these systems interchangeably.

6.2 Problem statement: controlled biochemical networks and population measurements

In the following we are interested in characterizing and possibly controlling the behavior of a population of identical cells, each containing a copy of the biochemical network of interest. To this end, we make two standing assumptions.

Population measurements: We can observe the system behavior by performing experiments in which at $H \in \mathbb{N}$ predetermined measurement times $\{t_h\}_{h=1}^H$ we:

1. collect a random sample of the population consisting of $N \in \mathbb{N}$ cells ³;
2. measure the amount of one or more species in each cell of the sample, thus constructing the dataset $\mathcal{Y}(t_h) := \{y_i(t_h)\}_{i=1}^N$;
3. discard the sample.

This is the typical process for example in flow-cytometry experiments, where one or more species of interest are tagged with fluorescence labels that are then measured via lasers and fluorescence detectors. With this measurement procedure the number N of cells measured at every sampling time is typically very large. Moreover, since cells are discarded, the recordings $\{y_i(t_h)\}_{i=1}^N$ at different measurements times are statistically independent.

Controlled reactions: We can intervene on the population by means of $M \in \mathbb{N}$ external signals, denoted by $\sigma_m(t) \in \Sigma_m \subseteq \mathbb{R}_{\geq 0}$, $m \in \{1, \dots, M\}$, which are applied to the whole population. Each external signal σ_m acts on the system by modifying the propensity of a specific reaction r_m in a multiplicative fashion, that is, it modifies its propensity from $\alpha_{r_m}(\theta_{r_m}, z(t))$ into $\alpha_{r_m}(\theta_{r_m}, z(t)) \cdot \sigma_m(t)$.

Examples of controlled stochastic biochemical reaction networks satisfying this assumption are described for example in [MASSO⁺11, RPMA⁺15, OHL⁺14, UMD⁺12, MDBDB11]. Note that the CME in case of controlled reaction results in

$$\begin{aligned} \dot{p}(z, t) = & \sum_{r=1}^M [p(z - \nu_r, t) \alpha_r(\theta_r, z - \nu_r) - p(z, t) \alpha_r(\theta_r, z)] \sigma_r(t) + \\ & + \sum_{r=M+1}^R [p(z - \nu_r, t) \alpha_r(\theta_r, z - \nu_r) - p(z, t) \alpha_r(\theta_r, z)], \end{aligned} \quad (6.14)$$

where we assumed, without loss of generality, that the reactions have been ordered so that $r_m = m$, for each $m \in \mathbb{Z}[1, M]$. In Section 8.3 and 8.4 we show how the solution methods described in Sections 6.1.1 and 6.1.2 for autonomous reaction networks can be adapted to approximate the CME of controlled reaction networks. For example, following the same steps as in Section 6.1.2, one can derive an equation describing the evolution

³For simplicity we assume that the number of cells sampled at each measurement time is the same. All the results can trivially be generalized to different sample sizes $N(t_h)$.

of the low-order moments of (6.14) as a function of the applied external signal, resulting in the controlled system

$$\dot{x}_{\leq l}(t) = f(\theta, x_{\leq l}(t), \sigma(t)). \quad (6.15)$$

To sum up, in the second part of the thesis, we analyze the behavior of a population of cells that we can influence through some external signals and that we can track via high-throughput population measurements. The experimental apparatus is thus defined by the available external signals and the species that can be measured. For a fixed experimental apparatus, an experiment e is completely specified by the applied input and by the measurements times, so that $e = \{\sigma(t), t_1, \dots, t_H\}$. The class of all possible experiments is denoted by \mathcal{E} . The objective of the second part of the thesis is to

1. select the best sequence of experiments $e_i \in \mathcal{E}$, $i = 1, \dots, I$ to be performed to “optimally” infer the rate parameters θ of a given reaction network from the available population measurements (Chapter 7);
2. once a model and its parameters have been identified, use the model to either predict the population behavior for an unseen experiment (Chapter 7) or characterize the whole range of behavior that the population can exhibit under different choices of the external signals (Chapter 8);
3. control the population behavior by designing a suitable external signal (Chapter 9).

To demonstrate the validity of our theoretical results we report in Chapter 9 an *in vivo* case study in which the proposed techniques are used to characterize and control a light inducible gene expression circuit implemented in yeast.

A systematic parameter inference procedure from population data

Systems biology rests on the idea that biological complexity can be better unraveled through the interplay of modeling and experimentation. The success of this approach however depends critically on the informativeness of the chosen experiments, which is usually unknown a priori. In this chapter, we propose a systematic scheme, based on iterations of optimal experiment design, population experiments and Bayesian parameter inference, to guide the discovery process in the case of stochastic biochemical reaction networks. Specifically, in Sections 7.1, 7.2, 7.3 we review known results on parameter inference, posterior predictive distributions and experiment design. We then combine these results to propose in Section 7.4 a sequential characterization procedure for controlled stochastic biochemical reaction networks from population data. The results of Section 7.4 have been published in [RPMA⁺15].

7.1 Parameter inference

Consider a controlled stochastic biochemical reaction network, as described in Section 6.2, consisting of R reactions with unknown rates $\theta = [\theta_1, \dots, \theta_R]$. Moreover, suppose that a specific experiment $e \in \mathcal{E}$ has been performed and that the data $\mathcal{Y}^e = \{\mathcal{Y}^e(t_h)\}_{h=1}^H$ has been recorded. The main objective of parameter inference is to use this data to infer the value of the unknown parameter vector θ . Let us denote by $p(z, t \mid \theta, e)$ the solution of the controlled CME in (6.14), to make explicit its dependence on the performed experiment e and on the specific parameter vector θ used to solve it. The maximum likelihood estimator is by definition the parameter vector θ that maximizes the probability of observing the recorded outcome \mathcal{Y}^e , that is,

$$\hat{\theta}^{\text{ML}}(\mathcal{Y}^e) := \arg \max_{\theta} L(\mathcal{Y}^e \mid \theta) := \arg \max_{\theta} \mathbb{P}(\mathcal{Y}^e \mid \theta).$$

The function $L(\mathcal{Y}^e \mid \theta)$ is the *likelihood* of observing the outcome \mathcal{Y}^e conditioned on the parameters choice θ . With the notation introduced above and recalling that, according

to our assumption, the measurements in \mathcal{Y}^e are independent one obtains

$$\hat{\theta}^{\text{ML}}(\mathcal{Y}^e) = \arg \max_{\theta} \{\prod_{h=1}^H \prod_{i=1}^N p(y_i(t_h), t_h | \theta, e)\}.$$

The main drawback of the maximum likelihood approach is that, to solve the previous optimization problem, one needs to solve the CME in (6.14) for each possible value of the parameter θ . As discussed in Section 6.1 this is an impossible task even for very simple biochemical reaction networks. To overcome this issue, we review in the next section an alternative, approximate approach, that uses the moment equations.

7.1.1 The likelihood of sample moments

While solving the CME (6.14) is in general a very difficult task, computing its low order moments is a viable alternative for many systems of practical interest. In fact, as illustrated in Section 6.1.2, under Assumption 6.1.3 or by using a suitable moment closure technique [Hes08], this problem is equivalent to solving the ordinary differential equation in (6.15). In the following, we denote by $x(t | \theta, e)$ the solution of (6.15) to emphasize its dependence on the experiment e and parameter θ . The main idea of the moment approach is to use the dataset $\mathcal{Y}^e = \{\mathcal{Y}^e(t_h)\}_{h=1}^H$ to construct an estimate of the low-order moments at each measurement time t_h and then select the parameter vector that maximizes the likelihood of observing those moments instead of the likelihood of observing each single realization. Let us assume for simplicity that only one species s is measured, so that $y_i(t_h) \in \mathbb{N}$ is the amount of species s in the i -th cell of the sample taken at time t_h ¹. Then the estimators for mean and variance of $Z_s(t_h)$ are

$$\hat{\mu}_s^e(t_h) = \frac{1}{N} \sum_{i=1}^N y_i^e(t_h)$$

$$\hat{\mu}_{s^2}^e(t_h) = \frac{1}{N-1} \sum_{i=1}^N (y_i^e(t_h) - \hat{\mu}_s^e(t_h))^2.$$

In the following we refer to $\hat{\mu}_s^e(t_h), \hat{\mu}_{s^2}^e(t_h)$ as the *sample moments*. Let the set $\mathcal{D}^e = \{\hat{\mu}_s^e(t_h), \hat{\mu}_{s^2}^e(t_h)\}_{h=1}^H$ be the sample moments dataset. Note that since \mathcal{Y}^e is a collection of random realizations, also the estimators $\hat{\mu}_s^e(t_h)$ and $\hat{\mu}_{s^2}^e(t_h)$ are random variables. Since they are obtained as sum of the independent random variables $y_i^e(t_h)$ when the sample size N tends to infinity, thanks to the central limit theorem, the distribution of $\hat{\mu}_s^e(t_h)$ and $\hat{\mu}_{s^2}^e(t_h)$ tends to a normal distribution, whose mean and variance can be computed from the moments $x(t | \theta, e)$ of the Z process from which the samples $y_i^e(t_h)$ were taken.

¹The following results can be generalized to an arbitrary number of measured species.

Specifically,

$$\hat{\mu}_{t_h}^e := \begin{bmatrix} \hat{\mu}_s^e(t_h) \\ \hat{\mu}_{s^2}^e(t_h) \end{bmatrix} \sim \mathcal{N}(\mu_{t_h|\theta}^e, \Sigma_{t_h|\theta}^e) \quad \text{where} \quad (7.1)$$

$$\mu_{t_h|\theta}^e := \begin{bmatrix} x_s(t_h | \theta, e) \\ x_{s^2}(t_h | \theta, e) \end{bmatrix}, \quad \Sigma_{t_h|\theta}^e := \frac{1}{N} \begin{bmatrix} x_{s^2}(t_h | \theta, e) & x_{s^3}(t_h | \theta, e) \\ x_{s^3}(t_h | \theta, e) & x_{s^4}(t_h | \theta, e) - \frac{N-3}{N-1} (x_{s^2}(t_h | \theta, e))^2 \end{bmatrix}$$

where $x_{s^i}(t_h | \theta, e)$ denotes the i -th central moment of $Z_s(t)$ (that is, the relative component in $x(t_h | \theta, e)$). Exploiting the fact that for large sample size N the moments are normally distributed, it is possible to compute an explicit formula for the likelihood of a parameter vector, given the estimated moments \mathcal{D}^e , that leads to the maximum likelihood estimator from sample moments

$$\hat{\theta}^{\text{ML}}(\mathcal{D}^e) = \arg \max_{\theta} L(\mathcal{D}^e | \theta) = \arg \max_{\theta} \prod_{h=1}^H p(\hat{\mu}_{t_h}^e | \theta, e) \quad (7.2)$$

$$:= \arg \max_{\theta} \left\{ \prod_{h=1}^H \frac{1}{2\pi \sqrt{|\Sigma_{t_h|\theta}^e|}} \exp \left(-\frac{1}{2} (\hat{\mu}_{t_h}^e - \mu_{t_h|\theta}^e)^\top \Sigma_{t_h|\theta}^e^{-1} (\hat{\mu}_{t_h}^e - \mu_{t_h|\theta}^e) \right) \right\}.$$

The advantage of performing inference from sample moments instead of realization is that, to solve (7.2), one has to compute the solution of the system of low-order moments given in (6.15) instead of the full CME (6.14), thus making the inference task computationally tractable. Clearly, this computational advantage comes at the expenses of discarding all the information contained in the higher order moments. Nonetheless, in many practical cases the information contained in the low order moments is still sufficient to accurately identify the unknown parameters (i.e. $\hat{\theta}^{\text{ML}}(\mathcal{D}^e) \cong \hat{\theta}^{\text{ML}}(\mathcal{Y}^e)$) [RL13].

7.1.2 Bayesian parameter inference

The maximum likelihood estimator in formula (7.2) allows one to select the parameter vector that best explains the sample moments dataset \mathcal{D}^e . In biological application, however, scientists have additional information on the reaction rates, coming for example from previous literature studies, and are interested in ranges of plausible parameters rather than in single parameter vectors. These aspects can be incorporated in the inference procedure by using a Bayesian approach, which aims at computing a probability distribution over the parameter space instead of a point estimator. In the Bayesian framework, any additional information on the parameters is encoded in a *prior parameter distribution* $p(\theta)$ which is then modified, according to the Bayes formula, to account for the measured data, leading to the *posterior parameter distribution*

$$p(\theta | \mathcal{D}^e) = \frac{L(\mathcal{D}^e | \theta) p(\theta)}{p(\mathcal{D}^e)}, \quad (7.3)$$

where $L(\mathcal{D}^e | \theta)$ is as in (7.2) and $p(\mathcal{D}^e) = \int L(\mathcal{D}^e | \theta) p(\theta) d\theta$. We note that in general it is not possible to compute an analytic expression for the posterior distribution. However,

particle approximations of such a distribution can be obtained using sampling methods as, for example, the Metropolis–Hastings algorithm or more complex sequential Monte Carlo algorithms if the parameter space is high dimensional [DMDJ06, MA13]. If a point estimate $\hat{\theta}$ of the model parameters is also desired, the maximum a posteriori (MAP) estimate

$$\hat{\theta}^{\text{MAP}}(\mathcal{D}^e) := \arg \max_{\theta} p(\theta \mid \mathcal{D}^e), \quad (7.4)$$

that is, the maximizer of the posterior distribution, can be extracted from the particle approximation. Note that, if the prior parameter distribution $p(\theta)$ is uniformly distributed, then $\hat{\theta}^{\text{MAP}}(\mathcal{D}^e)$ coincides with $\hat{\theta}^{\text{ML}}(\mathcal{D}^e)$. Finally, since the data collected in different experiments is statistically independent, it is straightforward to extend (7.2), (7.3), (7.4) to a dataset containing multiple experiments, $\mathcal{D} = \{\mathcal{D}^e\}_{e=1}^E$, by using as likelihood

$$L(\mathcal{D} \mid \theta) = \prod_{e=1}^E L(\mathcal{D}^e \mid \theta).$$

The parameter posterior distribution, computed using the Bayesian inference approach, allows one to quantify the available information on the parameter vector θ . A posterior parameter distribution which is very peaked around a single value suggests that the MAP estimator is a “reliable” parameter estimate. On the other hand, multinomial or flat posterior distributions can be an index of either one of two scenarios: i) the data that has been used in the Bayesian inference approach was not informative enough or ii) the model is practically unidentifiable, which means that there are multiple parameter vectors that lead to the same observed behavior. To distinguish among these two cases, and more in general to determine how well new experiments can be predicted by the inferred model, one can compute the *posterior predictive distribution* of the sample moments for one or more validation experiments.

7.2 Posterior predictive distributions of sample moments

The most straightforward approach to predict the response of the population in a new *validation* experiment² $v \in \mathcal{E}$, given the data \mathcal{D} recorded in previous experiments, is to solve the corresponding moment equations with the point estimate $\hat{\theta}^{\text{MAP}}(\mathcal{D})$ as parameters vector. Even though this is a valid approach, it does not allow one to quantify how certain this prediction actually is. Consider, for instance, a case where the parameter

²Note that we use the symbol $v \in \mathcal{E}$ instead of the generic $e \in \mathcal{E}$ to denote a validation experiment. Our aim is to stress the fact that the posterior predictive distributions are computed before the experiment is performed to predict its outcome, based on the dataset \mathcal{D} recorded in some previous experiments (i.e., not in v). Such predictions can be used to validate the model based on the difference between the predicted outcome and the actual data \mathcal{D}^v , which is measured a posteriori.

posterior distribution is very flat so that its density function is only marginally larger at the MAP estimate than at other parameter values. In such a case predictions computed with parameters other than the MAP estimate would be almost equally likely outcomes of the new experiment, even though they could be significantly different. To overcome this issue one should validate the model by using the complete posterior predictive distribution instead of predictions computed with point estimates only. This distribution describes how likely different measurements, in our case sample mean and variance $\hat{\mu}_t^v$, are for a new validation experiment v , given all the previously measured data \mathcal{D} . The posterior predictive distribution can be computed from the parameter posterior distribution $p(\theta | \mathcal{D})$ according to

$$p^{pred}(\hat{\mu}_t^v | \mathcal{D}, v) = \int_{\theta} p(\hat{\mu}_t^v | \theta, v) p(\theta | \mathcal{D}) d\theta,$$

where $p(\cdot | \theta, v)$ is the distribution of $\hat{\mu}_t^v$ given the selected experiment v and that θ are the model parameters, as in (7.2). The distribution $p^{pred}(\cdot, | \mathcal{D}, v)$, can be approximately computed by replacing the integral over θ with a sum over samples $\{\theta_q\}_{q=1}^Q$, drawn from the posterior distribution $p(\theta | \mathcal{D})$. Since for each θ the sample moments distribution $p(\cdot | \theta, v)$ is approximately a two-variate gaussian distribution³ we obtain a gaussian mixture approximation of the posterior predictive distribution at time t as

$$p^{pred}(\cdot | \mathcal{D}, v) \approx \frac{1}{Q} \sum_{q=1}^Q \mathcal{N} \left(\mu_{t|\theta_q}^v, \Sigma_{t|\theta_q}^v \right),$$

where $\mu_{t|\theta_q}^v$ and $\Sigma_{t|\theta_q}^v$ are as in (7.1).

7.3 Experiment design for parameter inference

For the inference process described in Section 7.1 to be successful, it is of paramount importance to design and perform experiments that yield the information required to identify the model parameters. To this end, optimal experiment design techniques for biochemical stochastic networks have been recently developed [RMAL13, ZNUK12].

7.3.1 The Fisher information

The first thing that is required for the design of optimal experiments is a way of quantifying the information that an experiment can provide about the unknown parameter θ . One way of quantifying this information is through the computation of the Fisher information matrix [KCRS11]. The entries of the Fisher information matrix, $I(\theta, e) \in \mathbb{R}^{R \times R}$,

³Due to the finite sampling noise as described in Section 7.1.1.

are given by

$$[I(\theta, e)]_{k,l} = \mathbb{E} \left[\left(\frac{\partial}{\partial \theta_k} \log L(\mathcal{D}^e | \theta) \right) \left(\frac{\partial}{\partial \theta_l} \log L(\mathcal{D}^e | \theta) \right) \right], \quad (7.5)$$

where $L(\mathcal{D}^e | \theta)$ is the likelihood of the data \mathcal{D}^e , recorded in experiment e , given that θ are the model parameters as in (7.2), and the expectation is taken with respect to all possible realizations of the data \mathcal{D}^e . If one uses as data the sample moments of the measured species Z_s , the Fisher information matrix can be approximated, for N large, using only the first four moments of the probability distribution, according to the formulas derived in [RMAL13]:

$$I(\theta, e) = \sum_{h=1}^H I_{t_h}^e(\theta) \quad \text{where} \quad (7.6)$$

$$[I_{t_h}^e(\theta)]_{k,l} = N \left(\frac{\frac{\partial x_{s,t_h}}{\partial \theta_k} \frac{\partial x_{s,t_h}}{\partial \theta_l}}{x_{s^2,t_h}} + \frac{\left(x_{s^2,t_h} \frac{\partial x_{s^2,t_h}}{\partial \theta_k} - \frac{\partial x_{s,t_h}}{\partial \theta_k} x_{s^3,t_h} \right) \left(x_{s^2,t_h} \frac{\partial x_{s^2,t_h}}{\partial \theta_l} - \frac{\partial x_{s,t_h}}{\partial \theta_l} x_{s^3,t_h} \right)}{(x_{s^2,t_h})^2 (x_{s^4,t_h} - (x_{s^2,t_h})^2) - x_{s^2,t_h} (x_{s^3,t_h})^2} \right),$$

where $x_{s^i,t_h} = x_{s^i}(t_h | \theta, e)$ is the i -th centered moment of the distribution $\mathbb{P}[Z_s(t_h) = z | \theta, e]$. To evaluate this formula, in addition to the moments themselves, partial derivatives of means and variances with respect to θ have to be computed from the model. These can be obtained by solving the population moment equations (6.12) with any solver for ordinary differential equations which also returns parameter sensitivities, such as CVODES of the SUNDIALS toolbox [HBG⁺05].

7.3.2 Designing an optimal experiment

The expectation in (7.5) is taken with respect to all possible realizations of the data. Accordingly, the Fisher information matrix does not depend on any measurement and can be used to evaluate the utility of different experiments before they are performed. This means that one can search among all possible experiments for the one which can be expected to provide the most information about the model parameters. In other words, one can aim at solving the optimization problem

$$e^* = \arg \max_{e \in \mathcal{E}} \{ \det I(\theta, e) \}, \quad (7.7)$$

where $I(\theta, e)$ is the Fisher information matrix defined in (7.6) and \mathcal{E} is the set of all possible experiments. The determinant $\det I(\theta, e)$ in (7.7) provides one way of summarizing the information of an experiment in a scalar quantity that can be maximized. This is known as D-optimality. There exist many other optimality criteria, we refer the reader to [Rue14] for a detailed discussion.

It is important to underline that the Fisher information matrix depends on the values of the parameters θ which are to be estimated. These parameters are obviously unknown

(otherwise performing an experiment for their identification would not be necessary). One way to overcome this problem is described in the next section.

7.4 Sequential experiment design for parameter inference

In the previous section we have described how, given an experiment, one can identify the “best” parameter vector θ and how, given a parameter vector, one can design an “optimal” experiment. Sequential experiment design aims at combining these two processes to infer the parameter vector with as less experiments as possible. Specifically, in sequential experiment design the parameters θ are replaced by their best currently available estimates $\hat{\theta}$, so that a new experiment is designed using $\hat{\theta}$ for the computation of the Fisher information matrix. The data collected in this experiment can then be used to improve the quality of the estimates and another experiment can be designed with the updated parameter estimates. In general, there is no guarantee that evaluating $I(\theta, e)$ at estimated values $\hat{\theta}$ will result in the design of informative experiments. Simulation studies [HWT13a], however, have shown that sequential experiment design often leads to good results, especially if many different experiments are needed to identify the model parameters. By making use of the inference and experiment design methods reviewed in Section 7.1 and 7.3, we propose here a sequential experiment design procedure for controlled stochastic biochemical reaction networks from population data.

The proposed approach, illustrated in Figure 7.1 is an iterative procedure where at every step k :

1. A new experiment e_k , optimally complementing the previously performed experiments $\{e_i\}_{i=1}^{k-1}$, is designed based on the MAP estimate $\hat{\theta}_{k-1}$. Mathematically,

$$e_k = \arg \max_{e \in \mathcal{E}} \left\{ \det I \left(\hat{\theta}_{k-1}, \{e_i\}_{i=1}^{k-1} \cup \{e\} \right) \right\}.$$

2. The experiment e_k is performed and the sample moments dataset \mathcal{D}^{e_k} is computed from the recorded measurements \mathcal{Y}^{e_k} .
3. The posterior parameter distribution is updated, by using Bayes formula,

$$p(\theta | \{\mathcal{D}^{e_i}\}_{i=1}^k) = \frac{L(\mathcal{D}^{e_k} | \theta) p(\theta | \{\mathcal{D}^{e_i}\}_{i=1}^{k-1})}{p(\mathcal{D}^{e_k})},$$

and the new MAP estimate is computed

$$\hat{\theta}_k(\{\mathcal{D}^{e_i}\}_{i=1}^k) := \arg \max_{\theta} p(\theta | \{\mathcal{D}^{e_i}\}_{i=1}^k).$$

The procedure stops when either the posterior parameter distribution is sufficiently peaked or the posterior predictive distributions for some validation experiments (different from $\{e_i\}_{i=1}^k$) are sufficiently tight. To initialize the sequential procedure one can set $p(\theta | \{\mathcal{D}^{e_i}\}_{i=1}^0)$ to be the prior distribution $p(\theta)$ and $\hat{\theta}_0 = \arg \max_{\theta} p(\theta)$, if the prior is not flat, or a value taken from the literature, otherwise.

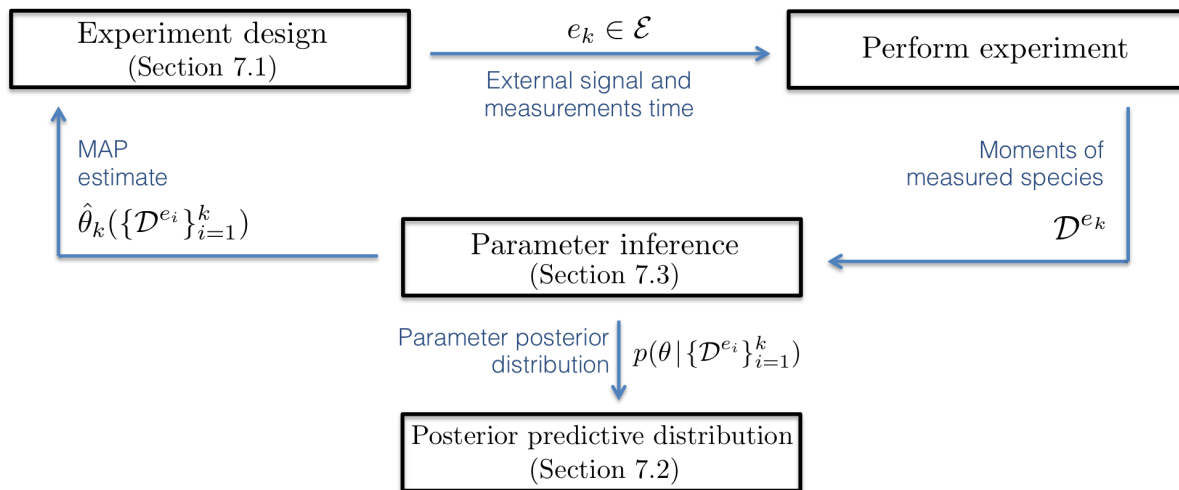


Figure 7.1: Sequential experiment design for parameter inference.

An important aspect to underline is that, for the design of the next experiment, the joint Fisher information matrix of the already performed experiments and the yet to be determined experiment has to be computed. If only the Fisher information matrix of the new experiment alone would be used for the design, it would be likely that an experiment which is similar to the first one would be designed (since the only difference to the design of the first experiment would be that the Fisher information matrix is evaluated at updated parameter values). If, on the other hand, the joint information is used, one can expect that an experiment that adds new information and in some sense complements the already performed experiments is designed at every step [HWT13b]. We also note that only the MAP estimate is used in the experiment design step instead of the full parameter posterior distribution. This is a suboptimal choice which is however mandatory for computational reasons. Selecting the best experiment according to the full parameter distribution would indeed require the computation, for each candidate experiment, of the expected Fisher information matrix over the parameter distribution. This task cannot be solved analytically but could be solved approximately by replacing the integral over θ with a sum over samples $\{\theta_{q_k}\}_{q_k=1}^Q$, drawn from the posterior distribution $p(\theta | \{\mathcal{D}^{e_i}\}_{i=1}^k)$

and then selecting as next experiment

$$e_{k+1} = \arg \max_{e \in \mathcal{E}} \left\{ \sum_{q_k=1}^Q \det I(\theta_{q_k}, \{e_i\}_{i=1}^k \cup \{e\}) \right\}.$$

This would however increase the computation time by a multiplicative factor Q .

The usefulness of sequential experiment design is verified in Chapter 9, where the proposed approach is applied to the in vivo characterization of a light inducible gene expression system.

Reachability analysis

A fundamental question in the study of stochastic controlled biochemical reaction networks is to what extent is noise an intrinsic property of the system and consequently what are the fundamental limits in noise suppression [LVP10]. In this chapter we tackle this question by rephrasing it in terms of moments of the underlying stochastic process. Specifically, we aim at investigating what values of mean and variance (or higher order moments) of the species present in the network are obtainable by perturbing the system with the available external signal. To this end, we adapt the hyperplane method, originally proposed in [GK91], to compute inner and outer approximations of the reachable set of the system describing the moments evolution. A remarkable feature of this approach is that it allows one to easily compute projections of the reachable set for pairs of moments of interest, without first requiring the computation of the full reachable set, thus making our method scalable.

The results of this chapter have been published in [PVL14, PVL15, PVL16].

8.1 Problem formulation

One of the most impressive results achieved by synthetic biology in the last decade is the synthesis and introduction of the externally controllable modules described in Section 6.2 in pre-existing biochemical reaction networks, allowing researchers to influence and possibly control the behavior of a cell *in vivo*. A fundamental prerequisite for the successful application of these modules, however, is understanding what range of behaviors they can exhibit under different choices of the external signal. For deterministic systems, this refers to the problem of computing the set of states that can be reached by the system trajectories starting from a known initial configuration [BDJPG08, CF03]. As discussed in Chapter 6, however, biochemical reaction networks are inherently stochastic and, for any fixed control input, many different evolutions may arise with different likelihood. Therefore, the analysis of these systems has to be formulated in a probabilistic setting. One may be interested, for example, in computing the probability of entering a prespecified subset of the state space, given a fixed initial condition or an initial probability

distribution [BHHK03, AKLP10, ESPP+06, MK06].

Here we analyze controlled stochastic biochemical reaction networks by posing a related but different question: instead of focusing on the behavior of the different realizations we are interested in characterizing how the external input influences the mean, variance (and possibly higher order moments) of the stochastic process. This approach is motivated for example by biotechnology applications, where one would like to control the average production of the cells in large populations, instead of each cell individually. More on the theoretical side, this different perspective can be useful to investigate fundamental questions on noise suppression in biochemical reaction networks [LVP10]. In [PVL14], for example, a similar approach is employed to study a simple gene expression network and to derive bounds on the variability of protein levels that can be observed in a population of identical cells. Finally, information regarding the species' moments can be useful to derive computationally efficient algorithms to simulate the behavior of the full Markov chain distribution.

Note that while the trajectories of the system are stochastic the evolution of the moments of the process is deterministic, as detailed in Section 6.1.2. Consequently, for the class of networks described in Chapter 6, the above question can be rephrased as a standard reachability problem in the moments space. Many different methods have been proposed in the literature to compute the reachable set of a deterministic system, among which level set methods [MT05], ellipsoidal methods [KV97] and sensitivity based methods [DM07]. Biochemical networks typically have an high-dimensional state space, therefore the choice of a method that scales well with the system size is essential. Here, we opted for the hyperplane method introduced in [GK91] for linear systems. The reason is that, in biological applications, researchers are often interested in analyzing the behavior of only a few chemical species of the possibly many involved in the network. Consequently, one is typically interested in the projection of the reachable set (which is a high-dimensional object) on some low-dimensional space of interest. The hyperplane method stands out in this respect, since it allows one to compute directly the projection of interest, without requiring the computation of the full high-dimensional reachable set first. Finally, intersections of hyperplanes are very easy to handle and to visualize.

The chapter is organized as follows. In Section 8.2 we present the main theoretical results on reachability analysis. Specifically, in Section 8.2.2 we review the hyperplane method for linear systems with bounded input, while in Section 8.2.3 we extend it to switched affine systems. In Section 8.3 we apply these theoretical results to biochemical networks satisfying Assumption 6.1.3. Under this assumption the system describing the moment evolution is, in fact, closed and linear with bounded input (Section 8.3.1) or switched affine (Section 8.3.2). In Section 8.4, we consider generic reaction networks for which the system of moment equations is not necessarily closed. Since moment closure methods would lead to nonlinearities, the theoretical results of Section 8.2 cannot be applied. To obviate to this problem we derive in Section 8.4.2 an extension of the FSP

method for controlled biochemical networks and we show that the reduced Markov chain is switched affine. We then apply the results of Section 8.2.3 to this new system. Figure 8.1 provides a summary of the main contributions of this chapter. All the proofs are given in the Appendix.

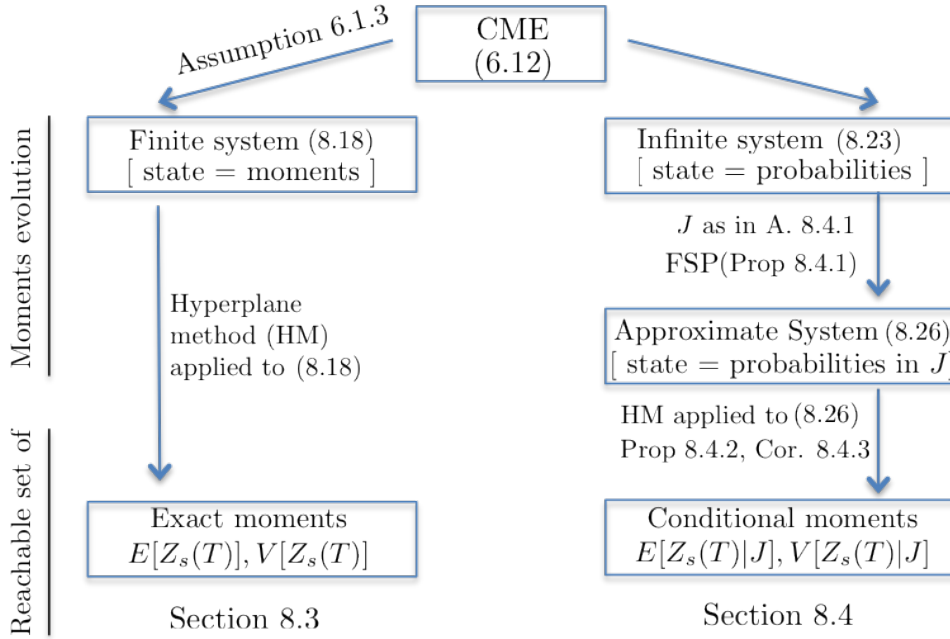


Figure 8.1: Contributions of Chapter 8.

8.2 Reachability tools

8.2.1 The reachable set

Consider the n -dimensional nonlinear controlled system

$$\dot{x}(t) = f(x(t), \sigma(t)) \quad t \geq 0. \quad (8.1)$$

In the following we assume that, for every initial condition $x(0) \in \mathbb{R}^n$ and every input function $\sigma(\cdot)$ belonging to *the set of admissible control laws* \mathcal{S} , the solution of (8.1), denoted by $x(t; x(0), \sigma(\cdot))$, is well defined and unique at every time $t > 0$. We also assume that the admissible laws $\sigma \in \mathcal{S}$ are uniformly bounded.

Remark 8.1. *The above assumptions are satisfied, for example, by a controlled linear system with bounded input, as discussed in Section 8.2.2, or by an autonomous switched affine system that can switch at most K times within a finite set of I possible modes*

$\Sigma := \{\bar{\sigma}^1, \dots, \bar{\sigma}^I\}$ at preassigned switching instants $0 = t_0 < \dots < t_{K+1} = T$, that is

$$\mathcal{S}_K := \{\sigma(\cdot) \mid \sigma(t) = \bar{\sigma}^{i_k} \in \Sigma, \quad \forall t \in [t_k, t_{k+1}), \forall k = 0, \dots, K\},$$

as discussed in Section 8.2.3.

The reachable set of system (8.1) at time T is defined as follows.

Definition 8.1 (Reachable set at time T). *The reachable set $\mathcal{R}_T(x_0)$ from x_0 at time $T > 0$, for the system (8.1) with admissible control set \mathcal{S} , is defined as the set of all states $x \in \mathbb{R}^n$ that are reachable at time T from $x(0) = x_0$, using an admissible control law:*

$$\mathcal{R}_T(x_0) := \{x \in \mathbb{R}^n \mid x = x(T; x_0, \sigma(\cdot)), \exists \sigma \in \mathcal{S}\}.$$

□

As motivated in the previous section, biologists are usually interested in characterizing the behavior of only a few species of the many involved in the network. In mathematical terms, this means that one is interested in the projection of the reachable set on a plane of interest. For example one may be interested in the relation between the mean behavior of two species or between mean and variance of a single species. For any $p, q \in \{1, \dots, n\}, p \neq q$, the projection of the reachable set $\mathcal{R}_T(x_0)$ in the (x_p, x_q) -plane is

$$\mathcal{R}_T^{p,q}(x_0) := \{y \in \mathbb{R}^2 \mid y = [x_p, x_q]^\top = \begin{bmatrix} e_p^\top \\ e_q^\top \end{bmatrix} x, \quad \exists x \in \mathcal{R}_T(x_0)\},$$

where e_p, e_q are canonical vectors with entry 1 in position p, q , respectively. This definition can be generalized to any pair of linear combinations C_1x and C_2x of the state components, where $C_1, C_2 \in \mathbb{R}^n$, by identifying the linear combinations C_1x and C_2x as the system output, and then investigating the set of all values that the output of the system can reach. Specifically, given a matrix $C := \begin{bmatrix} C_1 \\ C_2 \end{bmatrix} \in \mathbb{R}^{2 \times n}$, with C_j being the j th row of matrix C , we define the output of system (8.1) as

$$y(t) = Cx(t) = \begin{bmatrix} C_1x(t) \\ C_2x(t) \end{bmatrix} \in \mathbb{R}^2 \quad (8.2)$$

and the output reachable set as follows.

Definition 8.2 (Output reachable set at time T). *The output reachable set $\mathcal{R}_T^C(x_0)$ from x_0 at time $T > 0$, for system (8.1) with admissible control set \mathcal{S} and output given in (8.2), is defined as the set of all output values $y := Cx \in \mathbb{R}^2$ that are obtainable at time T from $x(0) = x_0$, using an admissible control law. That is,*

$$\mathcal{R}_T^C(x_0) := \{y \in \mathbb{R}^2 \mid y = Cx, \exists x \in \mathcal{R}_T(x_0)\}.$$

□

Computing the (output) reachable set for nonlinear systems is in general a very difficult task. In the next section we revise a method to approximate this set in the case of controlled linear systems with bounded input, which we then extend in Section 8.2.3 to the case of switched affine autonomous systems.

8.2.2 The hyperplane method for linear systems with bounded input

Consider a time-invariant linear system

$$\dot{x}(t) = Ax(t) + B\sigma(t), \quad (8.3)$$

where $x(t) \in \mathbb{R}^n$, $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$ and $\sigma(t) = [\sigma_1(t); \dots; \sigma_m(t)] \in \Sigma := \Sigma_1 \times \dots \times \Sigma_m$. If $\Sigma = \mathbb{R}^m$, that is, if the control input is unbounded the following classical result of control theory allows one to explicitly characterize the reachable set, as reported for example in [Son98, Lemma 3.2.2].

Proposition 8.2.1. *Consider the linear system in (8.3) and assume that $\Sigma = \mathbb{R}^m$. Then, for any $T > 0$, we get*

$$\mathcal{R}_T(x_0) = e^{AT}x_0 + \text{Im}(R), \quad R := [B, AB, A^2B, \dots, A^{n-1}B].$$

Consequently, $\mathcal{R}_T(x_0) = \mathbb{R}^n$ if and only if the pair (A, B) is reachable, that is, if $\text{rank}(R) = n$. \square

According to the previous proposition if the control input is unbounded the reachable set \mathcal{R}_T is an affine subspace and $\mathcal{R}_T(0)$ is independent on the final time T . This result does not hold anymore if the input is bounded, that is, if $\Sigma \subset \mathbb{R}^m$. Nonetheless, by assuming sufficient structure of Σ it is possible to prove that the reachable set has the following properties [Son98, Lib11].

1. If $0 \in \Sigma$, then for any $\hat{T} > T > 0$, it holds $\mathcal{R}_T(0) \subseteq \mathcal{R}_{\hat{T}}(0)$.
2. If the constraint set Σ is convex, the reachable set $\mathcal{R}_T(x_0)$ is convex, for any $T > 0$.

Since biological signals are non-negative and bounded, we assume in the following that $\Sigma_r = [0, \bar{\sigma}_r]$, so that Σ satisfies both the conditions above and according to the first statement, the longer the horizon T is the more states can be reached. The limiting set, that is, the set of states that can be reached in any arbitrarily long but finite time, is the infinite-time reachable set.

Definition 8.3 (Infinite-time reachable set). *The infinite-time reachable set $\mathcal{R}(x_0)$, for system (8.3) with control constraint set Σ , is defined as the set of all states $x \in \mathbb{R}^n$ that are reachable from $x(0) = x_0$, using an admissible control law. That is*

$$\mathcal{R}(x_0) := \{x \in \mathbb{R}^n \mid x \in \mathcal{R}_T(x_0), \quad \exists T > 0\}.$$

\square

The main idea of the hyperplane method, introduced in [GK91], is to construct an outer approximation of the reachable set \mathcal{R}_T as the intersection of a family of half-spaces defined by hyperplanes that are tangent to the boundary of \mathcal{R}_T , see Figure 8.2. For any given direction $c \in \mathbb{R}^n$, let us define two supporting hyperplanes that are tangent to $\partial\mathcal{R}_T$, one from above $H_T^a(c) := \{x \in \mathbb{R}^n \mid c^\top x = v_T^a(c)\}$ and one from below $H_T^b(c) := \{x \in \mathbb{R}^n \mid c^\top x = v_T^b(c)\}$. The two values $v_T^a(c), v_T^b(c) \in \mathbb{R}$ can be computed by solving the following optimization problems

$$v_T^a(c) := \max_{x \in \mathcal{R}_T(x_0)} c^\top x, \quad \text{and} \quad v_T^b(c) := \min_{x \in \mathcal{R}_T(x_0)} c^\top x, \quad (8.4)$$

for simplicity we omit the dependence of $v_T^a(c), v_T^b(c)$ on the initial condition x_0 . By reformulating these problems as finite time optimal control problems and using the Maximum Principle [Lib11], one can explicitly compute two control laws $\sigma^{a,*}(t)$ and $\sigma^{b,*}(t)$ that lead the state $x(0) = x_0$ to two points $x(T; x_0, \sigma^{a,*}) =: x_T^{a,*}(c)$ and $x(T; x_0, \sigma^{b,*}) =: x_T^{b,*}(c)$ satisfying

$$x_T^{a,*}(c) \in H_T^a(c) \cap \mathcal{R}_T(x_0), \quad x_T^{b,*}(c) \in H_T^b(c) \cap \mathcal{R}_T(x_0), \quad (8.5)$$

respectively. In other words, the two hyperplanes $H_T^a(c), H_T^b(c)$ are tangent to $\mathcal{R}_T(x_0)$ in $x_T^{a,*}(c), x_T^{b,*}(c)$. Note that these points are unique if the set $\mathcal{R}_T(x_0)$ is strictly convex.

Proposition 8.2.2 (Tangent hyperplanes [GK91]). *Given system (8.3), with control constraint sets $\Sigma_r := [0, \bar{\sigma}_r] \subset \mathbb{R}_{\geq 0}$, define the following admissible control laws, expressed component-wise for any $r = 1, \dots, m$,*

$$\sigma_r^{a,*}(t) := \begin{cases} \bar{\sigma}_r & \text{if } c^\top e^{A(T-t)} b_r > 0 \\ 0 & \text{if } c^\top e^{A(T-t)} b_r < 0 \\ 0 \leq \sigma_r^a \leq \bar{\sigma}_r & \text{if } c^\top e^{A(T-t)} b_r = 0 \end{cases}$$

$$\sigma_r^{b,*}(t) := \begin{cases} 0 & \text{if } c^\top e^{A(T-t)} b_r > 0 \\ \bar{\sigma}_r & \text{if } c^\top e^{A(T-t)} b_r < 0 \\ 0 \leq \sigma_r^b \leq \bar{\sigma}_r & \text{if } c^\top e^{A(T-t)} b_r = 0, \end{cases}$$

where b_r denotes the r^{th} column of B . Then

$$v_T^a(c) = c^\top e^{AT} x_0 + \sum_{r=1}^m \bar{\sigma}_r \int_0^T [c^\top e^{A(T-t)} b_r]_+ dt \quad (8.6)$$

$$v_T^b(c) = c^\top e^{AT} x_0 - \sum_{r=1}^m \bar{\sigma}_r \int_0^T [c^\top e^{A(T-t)} b_r]_- dt. \quad (8.7)$$

Suppose additionally that the pair (A, b_r) is reachable, that is $\text{rank}[b_r, Ab_r, \dots, A^{n-1}b_r] = n$, for all $r = 1, \dots, m$. Then there are no singular arcs, that is, there exists no interval $[\tau_1, \tau_2]$, with $0 \leq \tau_1 < \tau_2 \leq T$ such that $c^\top e^{A(T-t)} b_r = 0$ for all $t \in [\tau_1, \tau_2]$. Consequently, the laws $\sigma_r^{a,*}(t), \sigma_r^{b,*}(t)$ are completely specified and the two intercepts can be obtained as

$$x_T^{a,*}(c) = e^{AT} x_0 + \int_0^T e^{A(T-t)} B \sigma^{a,*}(t) dt \quad \text{and} \quad x_T^{b,*}(c) = e^{AT} x_0 + \int_0^T e^{A(T-t)} B \sigma^{b,*}(t) dt. \quad (8.8)$$

□

Using Proposition 8.2.2 it is possible to construct approximations both of the finite-time reachable set $\mathcal{R}_T(x_0)$ and of the infinite-time reachable set from the origin $\mathcal{R}(0)$. In fact, let us define the half-spaces

$$\mathcal{H}_T^a(c) := \{x \in \mathbb{R}^n \mid c^\top x \leq v_T^a(c)\}, \quad \mathcal{H}_T^b(c) := \{x \in \mathbb{R}^n \mid c^\top x \geq v_T^b(c)\}, \quad (8.9)$$

where $v_T^a(c), v_T^b(c)$ are defined as in (8.6) and (8.7). By construction, the reachable set $\mathcal{R}_T(x_0)$ is contained in both half-spaces and therefore in their intersection. Moreover, since $x_T^{a,*}(c), x_T^{b,*}(c) \in \mathcal{R}_T(x_0)$ and $\mathcal{R}_T(x_0)$ is convex, $\text{conv}(x_T^{a,*}(c), x_T^{b,*}(c)) \subseteq \mathcal{R}_T(x_0)$. In the following proposition we generalize these intuitions to a set of different directions $\mathcal{C} := \{c^1, \dots, c^D\}$, as illustrated in Figure 8.2 for $D = 2$.

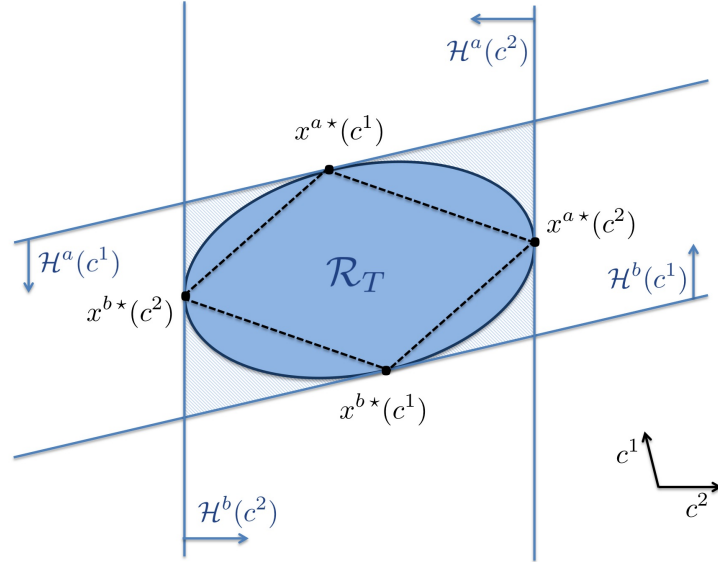


Figure 8.2: Illustration of the hyperplane method. The reachable set $\mathcal{R}_T(x_0)$ is in blue. The lined region is the outer approximation, the region in between the dotted lines is the inner approximation.

Proposition 8.2.3 (The hyperplane method [GK91]). *Given system (8.3), a fixed time $T > 0$, an integer number $D \geq 1$, and a set of D directions $\mathcal{C} := \{c^1, \dots, c^D\}$, define the half-spaces $\mathcal{H}_T^a(c^d), \mathcal{H}_T^b(c^d)$ as in (8.9), for all $d = 1, \dots, D$.*

1. *The following two sets are an outer and inner approximation of the reachable set $\mathcal{R}_T(x_0)$ at time T ,*

$$\mathcal{R}_T^{\text{out}}(x_0) := \bigcap_{d=1}^D \{\mathcal{H}_T^a(c^d) \cap \mathcal{H}_T^b(c^d)\}, \quad \mathcal{R}_T^{\text{in}}(x_0) := \text{conv} \left(\bigcup_{d=1}^D \{x_T^{a,*}(c^d), x_T^{b,*}(c^d)\} \right),$$

where $x_T^{a,}(c^d), x_T^{b,*}(c^d)$ are as defined immediately before (8.5).*

2. Suppose that $x_0 = 0$, A is Hurwitz stable and diagonalizable. Let $\{u_h\}_{h=1}^n$ and $\{\hat{u}_h^\top\}_{h=1}^n$ be n linearly independent right and left eigenvectors of A corresponding to the eigenvalues $\{\lambda_h\}_{h=1}^n$, respectively, and satisfying $\hat{u}_h^\top u_h = 1$ for all $h = 1, \dots, n$. For every $d = 1, 2, \dots, D$, set

$$\varepsilon_T(c^d) := \sum_{r=1}^m \left(\bar{\sigma}_r \sum_{h=1}^n \|d_{h,r}(c^d)\| \frac{e^{\operatorname{Re}(\lambda_h)T}}{|\operatorname{Re}(\lambda_h)|} \right) \geq 0,$$

where $d_{h,r}(c^d) := c^{d\top} u_h \hat{u}_h^\top b_r \in \mathbb{C}$, and set

$$\begin{aligned} \mathcal{H}^a(c^d) &:= \{x \in \mathbb{R}^n \mid c^{d\top} x \leq v_T^a(c^d) + \varepsilon_T(c^d)\}, \\ \mathcal{H}^b(c^d) &:= \{x \in \mathbb{R}^n \mid c^{d\top} x \geq v_T^b(c^d) - \varepsilon_T(c^d)\}, \end{aligned}$$

where $v_T^a(c^d), v_T^b(c^d)$ are as in (8.4). Then $\mathcal{R}^{\text{out}}(0) := \bigcap_{d=1}^D \{\mathcal{H}^a(c^d) \cap \mathcal{H}^b(c^d)\}$ and $\mathcal{R}^{\text{in}}(0) := \mathcal{R}_T^{\text{in}}(0)$ are an outer and inner approximation of $\mathcal{R}(0)$, respectively.

3. Set $\bar{\lambda} := \max\{\operatorname{Re}(\lambda) \mid \lambda \in \Lambda(A)\} < 0$, $\bar{\sigma} := \max\{\bar{\sigma}_r \mid r = 1, \dots, m\}$ and $l := \max_{d=1, \dots, D} \{\|c^{d\top} [u_1, \dots, u_n]\|_\infty \cdot \|\hat{u}_1, \dots, \hat{u}_n\|_\infty\}$. Then

$$\varepsilon_T(c^d) \leq l n m \bar{\sigma} \cdot \frac{e^{\bar{\lambda}T}}{|\bar{\lambda}|}, \quad \forall d = 1, \dots, D. \quad (8.10)$$

□

In words, the hyperplane method allows one to construct inner and outer approximations of the reachable set as intersections of the half-spaces described by the set of directions $\mathcal{C} := \{c^1, \dots, c^D\}$. The higher the dimension n of the state space, the higher would in general be the number of directions D required to obtain a good characterization of the reachable set. In [GK91, Figure 4] an iterative algorithm that adds one direction at a time, given the current estimate of the reachable set, is proposed. If, on the other hand, one is interested in the output reachable set only, then inner and outer approximations of the set \mathcal{R}_T^C can be easily computed, via the hyperplane method, by selecting only hyperplanes that are perpendicular to the plane of interest, as detailed in the following corollary.

Corollary 8.2.4 (Projection on a two dimensional subspace). *Consider system (8.3) with output (8.2). Choose D values $\gamma_d \in \mathbb{R}$, set $c^d := C_2^\top - \gamma_d C_1^\top \in \mathbb{R}^n$ and*

$$\mathcal{H}_T^{C,a}(\gamma_d) := \{y \in \mathbb{R}^2 \mid y_2 \leq \gamma_d y_1 + v_T^a(c^d)\}, \quad \mathcal{H}_T^{C,b}(\gamma_d) := \{y \in \mathbb{R}^2 \mid y_2 \geq \gamma_d y_1 + v_T^b(c^d)\},$$

where $v_T^a(c^d), v_T^b(c^d)$ are as in (8.6) and (8.7). Moreover, let

$$x_T^{C,a}(\gamma_d) := C x_T^{a,*}(c^d), \quad x_T^{C,b}(\gamma_d) := C x_T^{b,*}(c^d),$$

where $x_T^{a,*}(c^d), x_T^{b,*}(c^d)$ are defined as in (8.5). The sets

$$\mathcal{R}_T^{C,out}(x_0) := \cap_{d=1}^D \{\mathcal{H}_T^{C,a}(\gamma_d) \cap \mathcal{H}_T^{C,b}(\gamma_d)\}, \quad \mathcal{R}_T^{C,in}(x_0) := \text{conv} \left(\cup_{d=1}^D \{x_T^{C,a}(\gamma_d), x_T^{C,b}(\gamma_d)\} \right)$$

are an outer and inner approximation of the output reachable set $\mathcal{R}_T^C(x_0)$ at time T , respectively. \square

The previous results have been derived under the assumption that the inputs are bounded in the compact and convex sets Σ_r , $r = 1, \dots, m$. In the following proposition we show that the derived approximations are valid also if the control inputs take values in a subset Σ_r^d of Σ_r of finite cardinality, as long as 0 and $\bar{\sigma}_r$ are among the possible choices.

Proposition 8.2.5 (Finite cardinality constraint set). *Consider system (8.3) and suppose that the control constraint sets have finite cardinality, that is, for any r there exists $2 \leq K_r < +\infty$ such that $\Sigma_r^d := \{0 = \sigma_r^1 < \sigma_r^2 < \dots < \sigma_r^{K_r} = \bar{\sigma}_r\} \subset \mathbb{R}_{\geq 0}$. Then $\mathcal{R}_T^{out}(x_0)$ and $\mathcal{R}_T^{in}(x_0)$ are outer and inner approximation of $\mathcal{R}_T(x_0)$, respectively, $\mathcal{R}^{out}(0)$ is an outer approximation of $\mathcal{R}(0)$ and $\mathcal{R}_T^{C,out}(x_0)$ and $\mathcal{R}_T^{C,in}(x_0)$ are outer and inner approximation of $\mathcal{R}_T^C(x_0)$, respectively. \square*

8.2.3 The hyperplane method for switched affine systems

In this section, we propose an extension of the hyperplane method reviewed in the previous section to the case of a switched affine system of the form

$$\dot{x}(t) = A_{\sigma(t)}x(t) + b_{\sigma(t)}, \quad (8.11)$$

where the switching sequence $\sigma(\cdot)$ satisfies the following assumption.

Assumption 8.2.1. *The input sequence $\sigma(t)$ can switch at most K times within a finite set of I possible modes $\Sigma := \{\bar{\sigma}^1, \dots, \bar{\sigma}^I\}$ at preassigned switching instants $0 = t_0 < \dots < t_{K+1} = T$, that is, $\sigma \in \mathcal{S}_K$*

$$\mathcal{S}_K := \{\sigma(\cdot) \mid \sigma(t) = \bar{\sigma}^{i_k} \in \Sigma \quad \forall t \in [t_k, t_{k+1}), \quad \forall k = 0, \dots, K\}. \quad (8.12)$$

Under this assumption we show that an outer approximation of the desired output reachable set can be obtained by solving a collection of **mixed integer linear programs (MILPs)**. The main advantage of the proposed approach is that the number of MILPs to be solved does not scale with the system dimension but with the dimension of the output reachable set, which is considerably smaller. To this end, we rewrite the two

optimization problems in (8.4) as finite time optimal control problems, namely as:

$$v_T^a(c) := \max_{\sigma \in \mathcal{S}} c^\top x(T) \quad (8.13)$$

$$\begin{aligned} \text{s.t. } \dot{x}(t) &= A_{\sigma(t)}x(t) + b_{\sigma(t)} \quad \forall t \in [0, T] \\ x(0) &= x_0 \end{aligned}$$

$$v_T^b(c) := \min_{\sigma \in \mathcal{S}} c^\top x(T) \quad (8.14)$$

$$\begin{aligned} \text{s.t. } \dot{x}(t) &= A_{\sigma(t)}x(t) + b_{\sigma(t)} \quad \forall t \in [0, T] \\ x(0) &= x_0. \end{aligned}$$

The main advantage of this reformulation is that the solution of finite time optimal control problems for switched systems has been already addressed in the literature (see e.g. [SCGB06, BCV15]). We consider the case when the total number K and time $\{t_1, \dots, t_K\}$ of the switchings is fixed and known, that is, when the input set is \mathcal{S}_K as defined in (8.12). It is important to remark that, in this case, the reachable set consists of a finite number of points that can be computed by solving the moment equations for each possible input sequence. Since the cardinality of the set \mathcal{S}_K grows exponentially with K , this approach is however computationally infeasible even for small systems. Following the same procedure as in [SCGB06, Section IV.A], we show that, on the other hand, the hyperplane constants defined in (8.13) and (8.14) can be computed by solving an MILP. Let us consider the maximization problem in (8.13) and define $x_k := x(t_k)$, $\bar{A}_i^k := e^{A_{\bar{\sigma}^i}(t_{k+1}-t_k)}$ and $\bar{b}_i^k = [\int_0^{(t_{k+1}-t_k)} e^{A_{\bar{\sigma}^i}\tau} d\tau] b_{\bar{\sigma}^i}$ for every k and $i \in 1, \dots, I$. Then for the switched affine system in (8.11) we get

$$v_T^a(c) := \max_{i_k \in \{1, \dots, I\}} c^\top x_{K+1} \quad (8.15)$$

$$\begin{aligned} \text{s.t. } x_{k+1} &= \bar{A}_{i_k}^k x_k + \bar{b}_{i_k}^k \quad k = 0, \dots, K \\ x_0 &\in \mathbb{R} \text{ assigned.} \end{aligned}$$

We introduce the binary variables $\gamma_i^k \in \{0, 1\}$ defined so that, for each $i \in \{1, \dots, I\}$ and $k \in \{0, \dots, K\}$, $\gamma_i^k = 1$ if and only if the value of $\sigma(t)$ in the time interval $[t_k, t_{k+1})$ is $\bar{\sigma}^i$, namely the system is in mode i . Moreover, we introduce a copy of the state vector for each possible mode of the system: $z_i^{k+1} = (\bar{A}_i^k x_k + \bar{b}_i^k) \gamma_i^k$. Problem (8.15) can then be rewritten as

$$v_T^a(c) := \max_{x_k, z_i^k, \gamma_i^k} c^\top x_{K+1} \quad (8.16)$$

$$\text{s.t. } z_i^{k+1} = (\bar{A}_i^k x_k + \bar{b}_i^k) \gamma_i^k, \quad \forall i, \quad (8.17)$$

$$\sum_{i=1}^I \gamma_i^k = 1, \quad k = 0, \dots, K,$$

$$x_k = \sum_{i=1}^I z_i^k, \quad k = 1, \dots, K+1,$$

$$x_0 \in \mathbb{R} \text{ assigned.}$$

The constraint in (8.17) can be equivalently replaced by

$$\begin{aligned} z_i^{k+1} &\leq (\bar{A}_i^k x_k + \bar{b}_i^k) + B(1 - \gamma_i^k), & z_i^{k+1} &\geq -B\gamma_i^k, \\ z_i^{k+1} &\geq (\bar{A}_i^k x_k + \bar{b}_i^k) - B(1 - \gamma_i^k), & z_i^{k+1} &\leq B\gamma_i^k, \end{aligned}$$

where $B \in \mathbb{R}_{\geq 0}^n$ is an upper bound on the absolute value of state vector, that is, $B \geq \max_k \{|x_k|\}$ component-wise (see the big-M method in [BM99, Eq. (5b)]). With this substitution Problem (8.16) becomes an MILP. Similar reasonings can be applied to the minimization problem (8.14). We note that finding an a priori upper bound B on $|x(t_k)|$, for all $k \in \{0, \dots, K+1\}$ and all $\sigma \in \mathcal{S}_K$, is not trivial in general. We will however show in Sections 8.3 and 8.4 how to solve this problem for controlled biochemical reaction networks. The desired outer approximation of the reachable set can then be computed as in Section 8.2.2. Note that in the case of switched affine systems it is not possible to recover an inner approximation, since there is no guarantee in general that the reachable set would be convex. The outer approximation can, on the other hand, be derived exactly as done in [GK91] for linear systems.

Proposition 8.2.6 (The hyperplane method for switched affined systems). *Given system (8.11), a fixed time $T > 0$ and a set of $D \geq 1$ directions $\mathcal{C} := \{c^1, \dots, c^D\}$, define the half-spaces*

$$\begin{aligned} \mathcal{H}_T^a(c^d) &:= \{x \in \mathbb{R}^n \mid c^{d\top} x \leq v_T^a(c^d)\}, \\ \mathcal{H}_T^b(c^d) &:= \{x \in \mathbb{R}^n \mid c^{d\top} x \geq v_T^b(c^d)\}, \end{aligned}$$

for $d = 1, \dots, D$, where $v_T^a(c^d), v_T^b(c^d)$ are defined as in (8.13) and (8.14). The set $\mathcal{R}_T^{\text{out}}(x_0) := \bigcap_{d=1}^D \{\mathcal{H}_T^a(c^d) \cap \mathcal{H}_T^b(c^d)\}$ is an outer approximation of the reachable set $\mathcal{R}_T(x_0)$. \square

Similarly, an outer approximation of the output reachable set can be easily computed with the hyperplane method by selecting as set of directions \mathcal{C} the set of vectors $c^d := C_2^\top - \gamma_d C_1^\top \in \mathbb{R}^n$, $d = 1, 2, \dots, D$, for arbitrary choices of the real parameter γ_d .

Corollary 8.2.7. *Consider system (8.11) with output (8.2). Choose D values $\gamma_d \in \mathbb{R}$, set $c^d := C_2^\top - \gamma_d C_1^\top \in \mathbb{R}^n$ and*

$$\begin{aligned} \mathcal{H}_T^{C,a}(\gamma_d) &:= \{y \in \mathbb{R}^2 \mid y_2 \leq \gamma_d y_1 + v_T^a(c^d)\}, \\ \mathcal{H}_T^{C,b}(\gamma_d) &:= \{y \in \mathbb{R}^2 \mid y_2 \geq \gamma_d y_1 + v_T^b(c^d)\}, \end{aligned}$$

where $v_T^a(c^d), v_T^b(c^d)$ are computed as in (8.13) and (8.14). The set $\mathcal{R}_T^{C,\text{out}}(x_0) := \bigcap_{d=1}^D \{\mathcal{H}_T^{C,a}(\gamma_d) \cap \mathcal{H}_T^{C,b}(\gamma_d)\}$ is an outer approximation of the output reachable set $\mathcal{R}_T^C(x_0)$. \square

8.3 Reachable set of networks with affine propensities

To apply the previous theoretical results to controlled biochemical reaction networks let us consider the controlled CME given in (6.14) and reported here for convenience.

$$\begin{aligned} \dot{p}(z, t) = & \sum_{r=1}^M [p(z - \nu_r, t)\alpha_r(\theta_r, z - \nu_r) - p(z, t)\alpha_r(\theta_r, z)] \sigma_r(t) + \\ & + \sum_{r=M+1}^R [p(z - \nu_r, t)\alpha_r(\theta_r, z - \nu_r) - p(z, t)\alpha_r(\theta_r, z)]. \end{aligned} \quad (6.14)$$

Following the same steps as in Proposition 6.1.2, in the case of affine propensities it is possible to derive an explicit characterization of the moment equations for mean and variance in the presence of controlled reactions.

Proposition 8.3.1. *Consider the controlled CME given in (6.14) and assume that the propensity functions $\{\alpha_r(\theta_r, z)\}_{r=1}^R$ are affine functions of the state z (Assumption 6.1.3). Let $x_i(t) = \mathbb{E}[Z_i(t)]$, $x_{\leq 1}(t) = [x_1(t); \dots; x_S(t)]$ and $x_{ij}(t) = \mathbb{E}[(Z_i(t) - x_i(t))(Z_j(t) - x_j(t))]$. Then, for all $i, j \in \mathbb{Z}[1, S]$,*

$$\begin{aligned} \dot{x}_i = & \sum_{r=1}^M [\nu_{ir} \cdot \alpha_r(\theta_r, x_{\leq 1})] \sigma_r + \sum_{r=M+1}^R [\nu_{ir} \cdot \alpha_r(\theta_r, x_{\leq 1})] \\ \dot{x}_{ij} = & \sum_{r=1}^M \left[\nu_{ir} \cdot \sum_{l=1}^S \left(\frac{\partial \alpha_r(\theta_r, x_{\leq 1})}{\partial z_l} x_{jl} \right) + \nu_{jr} \cdot \sum_{l=1}^S \left(\frac{\partial \alpha_r(\theta_r, x_{\leq 1})}{\partial z_l} x_{il} \right) + \nu_{ir} \nu_{jr} \cdot \alpha_r(\theta_r, x_{\leq 1}) \right] \sigma_r \\ & + \sum_{r=M+1}^R \left[\nu_{ir} \cdot \sum_{l=1}^S \left(\frac{\partial \alpha_r(\theta_r, x_{\leq 1})}{\partial z_l} x_{jl} \right) + \nu_{jr} \cdot \sum_{l=1}^S \left(\frac{\partial \alpha_r(\theta_r, x_{\leq 1})}{\partial z_l} x_{il} \right) + \nu_{ir} \nu_{jr} \cdot \alpha_r(\theta_r, x_{\leq 1}) \right] \end{aligned}$$

where we omitted the time dependence for simplicity. \square

Corollary 8.3.2. *Let $x_{\leq 2}(t)$ be a vector whose components are the moments of $Z(t)$ up to second order. Then*

$$\dot{x}_{\leq 2}(t) = A(\sigma_{\text{first}}(t))x_{\leq 2}(t) + B\sigma_{\text{zero}}(t) + b \quad (8.18)$$

where $\sigma_{\text{zero}}(t)$ is a vector comprising all the external signals influencing a reaction with constant propensity, that is, $\alpha_r(\theta_r, z) = \theta_r$ and $\sigma_{\text{first}}(t)$ is a vector comprising the external signals influencing a reaction with linear propensity in z , that is, $\alpha_r(\theta_r, z) = \theta_r \cdot \nu_r' z$. \square

Remark 8.2. *We note that the affine term b in (8.18) is present if and only if there are uncontrolled zero order reactions. One could however include this term in the input matrix by introducing an artificial control signal that is always equal to one, so that*

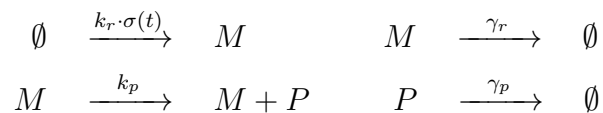
$$\dot{x}_{\leq 2}(t) = A(\sigma_{\text{first}}(t))x_{\leq 2}(t) + [B, b] \begin{bmatrix} \sigma_{\text{zero}}(t) \\ 1 \end{bmatrix}.$$

8.3.1 Linear moment equations

According to Corollary 8.3.2 if only reactions with constant propensities are controlled then the system of moment equations is linear with bounded input¹. Therefore the theoretical results of Section 8.2.2 can be applied. For mass-action kinetics, this case corresponds to the situation when only zero-order reactions are controlled. To illustrate this scenario we consider again the gene expression model of Example 6.1.

Gene expression case study: one external signal

Consider the gene expression model of Example 6.1 and assume that an external control signal influencing the first reaction, that is the mRNA production, is available [MASSO⁺11, RPMA⁺15, OHL⁺14, UMD⁺12, MDBDB11]. The corresponding biochemical network is thus



where $\sigma(t) \in \mathbb{R}_{\geq 0}$ is the external control signal, while the parameters $k_r, k_p, \gamma_r, \gamma_p > 0$ are the mRNA and protein production and degradation rates. Let the mean and the covariance of the two species be

$$\begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} \triangleq \mathbb{E} \begin{bmatrix} M(t) \\ P(t) \end{bmatrix} \quad \begin{bmatrix} x_1(t) & x_3(t) \\ x_3(t) & x_4(t) \end{bmatrix} \triangleq \mathbb{V} \begin{bmatrix} M(t) \\ P(t) \end{bmatrix}.$$

Note that we impose $\mathbb{E} [M(t)] = \mathbb{V} [M(t)] = x_1(t)$ since the mRNA follows a birth-death process, as explained in Example 6.1. The moments evolution over time is given by

$$\dot{x}(t) = Ax(t) + B\sigma(t), \tag{8.19}$$

$$A \triangleq \begin{bmatrix} -\gamma_r & 0 & 0 & 0 \\ k_p & -\gamma_p & 0 & 0 \\ k_p & 0 & -(\gamma_r + \gamma_p) & 0 \\ k_p & \gamma_p & 2k_p & -2\gamma_p \end{bmatrix}, \quad B \triangleq \begin{bmatrix} k_r \\ 0 \\ 0 \\ 0 \end{bmatrix}.$$

We note that, since the rate parameters are always positive, the matrix A in (8.19) is by construction Hurwitz stable. Moreover, apart from singular cases in the parameters space, the diagonal elements of A are distinct, leading to a diagonalizable matrix. Therefore, the assumptions of Proposition 8.2.3 are satisfied and one can use the results in Section 8.2.2 to study the reachable set of system (8.19). Note that the signal $\sigma(t)$ is a switching sequence taking values in the set $\Sigma^d = \{0, 1\}$ if the control input is of

¹See Remark 8.2 if uncontrolled zero order reactions are present in the network.

the ON-OFF type [MASSO+11, RPMA+15, OHL+14, MDBDB11] and in the interval $\Sigma^c = [0, 1]$ if the control input is continuous [UMD+12]. Proposition 8.2.5 guarantees the validity of the following results both for Σ^d and Σ^c .

For this biochemical network it is of particular interest to characterize what combinations of the protein mean and variance are achievable, that is, the output reachable set obtained by selecting as output matrix

$$C := \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

for $x_0 = 0$. The problem of computing an outer approximation of this set was studied in [PVL14] using ad hoc methods. In Figure 8.3 we compare the outer approximation obtained therein (magenta line) with the inner (red) and outer (blue) approximations that we obtained using the hyperplane method. For this case study we used the parameters identified in [MASSO+11] and we set a final horizon of $T = 10^4$ minutes, which, according to formula (8.10), leads to $\varepsilon_T(c^d) \leq 5.7 \cdot 10^{-20}$, hence for all practical purposes $\mathcal{R}_T^{out}(0) \equiv \mathcal{R}^{out}(0)$. This figure clearly shows that the outer approximation computed using the hyperplane method is more accurate than the one previously obtained in the literature. Moreover, since inner and outer approximations practically coincide, this method allows one to effectively recover the reachable set.

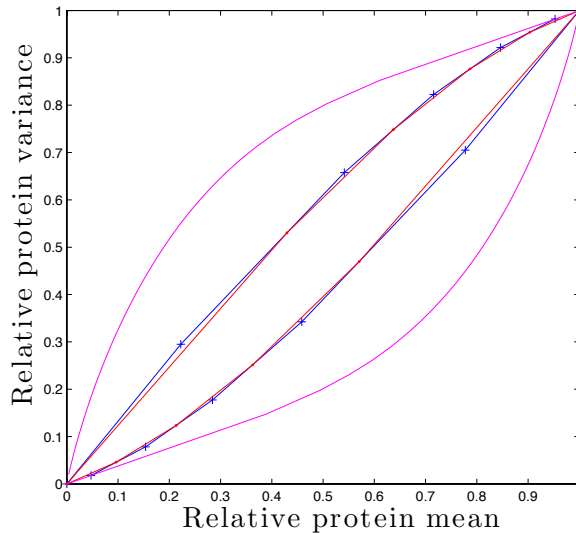


Figure 8.3: Comparison of the inner (red) and outer (blue) approximations of the reachable set for the protein mean and variance, according to model (8.19), computed using the hyperplane method and the outer approximation computed according to [PVL14] (magenta). The axis are normalized with respect to the maximum reachable value of protein mean and variance, respectively.

8.3.2 Switched affine moment equations

If the control inputs affect also reactions with affine propensities and if these inputs satisfy Assumption 8.2.1 then the moment equations in (8.18) of Corollary 8.3.2 can be equivalently rewritten as the autonomous switched system given in (8.11) and the theoretical tools described in Section 8.2.3 can be applied. Note that Assumption 8.2.1 is equivalent in biological applications to the assumption that the number of switchings and their timing during a given experiment is fixed a priori. This assumption can be motivated by the fact that changes in the external stimulus are costly and/or stressful for the cells. Moreover, it is trivially satisfied if the stimulus can only be changed simultaneously with some fixed events, such as culture dilution or measurements.

Gene expression case study: two external signals and fluorescent protein

We consider again the gene expression system of Example 6.1, but we now assume that

1. both mRNA production and degradation can be controlled, so that the vector of propensities is $\alpha(z) = [k_r \cdot \sigma_1(t), \gamma_r \cdot m \cdot \sigma_2(t), k_p \cdot m, \gamma_p \cdot p]^\top$ and $\sigma(t) := \begin{bmatrix} \sigma_1(t) \\ \sigma_2(t) \end{bmatrix}$;
2. the protein P can mature into a fluorescent protein F according to the additional maturation and degradation reactions



where $\alpha_5(k_f, z) := k_f \cdot p$, $\alpha_6(\gamma_p, z) := \gamma_p \cdot f$ and $k_f > 0$ is the maturation rate. The degradation rate of the fluorescent protein F is assumed to be the same as that of the non-mature protein P ;

3. the fluorescence intensity $I(t)$ of each cell can be measured by flow-cytometry and is proportional to the amount of fluorescence proteins, that is, $I(t) = rF(t)$ for a fixed scaling parameter $r > 0$.

The system describing the evolution of mean and variances of this augmented network

$$\dot{x}(t) = A_{\sigma(t)}^f x(t) + b_{\sigma(t)}^f, \quad (8.20)$$

$$A_{\sigma(t)}^f = \begin{bmatrix} -\gamma_r \sigma_2(t) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ k_p & -(\gamma_p + k_f) & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \gamma_p & -k_f & 0 & 0 & 0 & 0 & 0 & 0 \\ k_p & 0 & 0 & -(\gamma_r \sigma_2(t) + \gamma_p + k_f) & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \gamma_p & -(\gamma_r \sigma_2(t) + k_f) & 0 & 0 & 0 & 0 \\ k_p & (\gamma_p + k_f) & 0 & 2k_p & 0 & -2(\gamma_p + k_f) & 0 & 0 & 0 \\ 0 & -\gamma_p & 0 & 0 & k_p & \gamma_p & -(2k_f + \gamma_p) & 0 & 0 \\ 0 & \gamma_p & k_f & 0 & 0 & 0 & 2\gamma_p & -2k_f & 0 \end{bmatrix},$$

$$b_{\sigma(t)}^f = [k_r \sigma_1(t) \quad 0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0 \quad 0]^\top,$$

is 8-dimensional, with state vector

$$x = [\mathbb{E}[M], \mathbb{E}[P], \mathbb{E}[F], \mathbb{V}[M, P], \mathbb{V}[M, F], \mathbb{V}[P], \mathbb{V}[P, F], \mathbb{V}[F]]^\top,$$

and it depends on the parameter vector $\theta = [k_r, \gamma_r, k_p, \gamma_p, k_f, r]$ (for more details see [RMAL13, SI pg. 16]). These parameters can be identified from real data using the method described in Chapter 7, see also Chapter 9. For our case study we use the MAP estimate identified in [PVL15], that are, in min^{-1}

$$\begin{aligned} k_r &= 0.0236 & \gamma_r &= 0.0503 & k_p &= 178.398 \\ k_f &= 0.0212 & \gamma_p &= 0.0121 & r^{-1} &= 646.86 \end{aligned} \quad (8.21)$$

and we set

$$C^f := \begin{bmatrix} 0 & 0 & r & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & r^2 \end{bmatrix}, \quad (8.22)$$

to compute the mean and variance reachable set for the fluorescence intensity.

Our aim is to compare the reachable set of the gene expression system when both mRNA production and degradation are controlled (as in [BK12]), with the corresponding set for the case when only the mRNA production can be controlled (as studied in the previous section and in [PVL15]). To this end, we use the hyperplane method as described in Section 8.2.3 with possible modes $\bar{\sigma}^i := \begin{bmatrix} \bar{\sigma}_1^i \\ \bar{\sigma}_2^i \end{bmatrix}$ belonging to

$$\begin{aligned} \Sigma^{2\text{in}} &:= \left\{ \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0.5 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 0.5 \end{bmatrix} \right\} \\ \Sigma^{1\text{in}} &:= \left\{ \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 1 \end{bmatrix} \right\}, \end{aligned}$$

respectively. Note that we set the minimum input for the mRNA degradation to $0.5 > 0$ to avoid unboundedness. We set $x(0) = 0$ as initial state. With these choices of inputs it is intuitive that the highest possible state is reached when mRNA production is at its maximum and mRNA degradation is at its minimum. Therefore, in the MILPs we can use the bounds $B = x(T; 0, \sigma(t) = \begin{bmatrix} 1 \\ 0.5 \end{bmatrix} \forall t)$ for the case of two inputs and $B = x(T; 0, \sigma(t) = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \forall t)$, for the case of one input. We set $T = 300$ min and we assume that switches can occur every 20 min, so that $K = 15$. Figure 8.4 shows the output reachable set for the case of two inputs. The black crosses are obtained simulating the output of the system for 5000 randomly constructed input sequences. The simulation time for computing the outer approximation with the hyperplane method was 5.6 hrs. Computing the exact reachable set by simulating all the possible switching sequences, assuming that one simulation takes 10^{-4} sec and neglecting the time needed to enumerate all the possible sequences, would take 29.8 hrs. Figure 8.5 shows the comparison of the output reachable set obtained for the cases of one and two inputs.

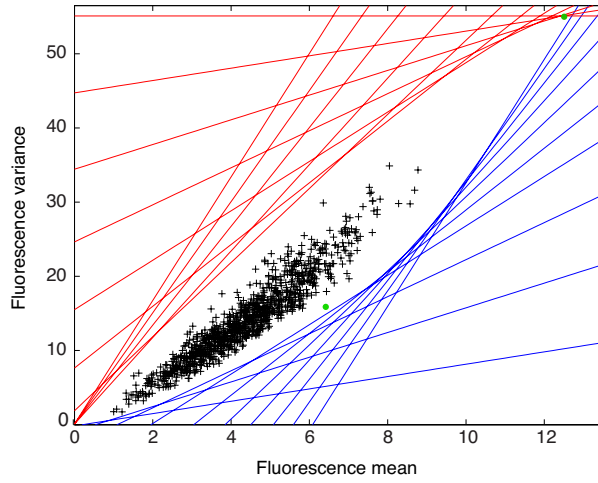


Figure 8.4: Outer approximation of the output reachable set of system (8.20), with output as in (8.22) and parameters as in (8.21), when the set of possible modes is $\Sigma^{2\text{in}}$. The red lines are the hyperplanes tangent from above, the blue the ones from below. The two green dots represent the outputs when $\sigma(t) = [1, 0.5]^\top \forall t$ and $\sigma(t) = [1, 1]^\top \forall t$, respectively. The black crosses represent the output for arbitrary sequences with $\sigma(t) \in \Sigma^{2\text{in}}$.

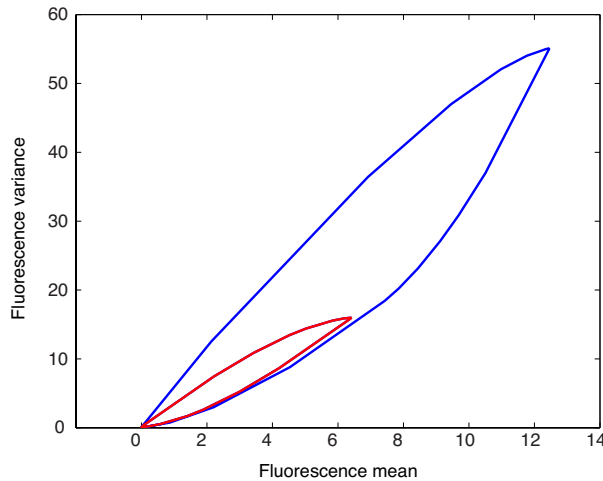


Figure 8.5: Output reachable set of system (8.20), with output as in (8.22) and parameters as in (8.21). The blue line is the outer approximation obtained for $\sigma(t) \in \Sigma^{2\text{in}}$ (see also Fig. 8.4), the red line the one for $\sigma(t) \in \Sigma^{1\text{in}}$.

8.4 Reachable set of networks with generic reactions

If the network contains reactions of order higher than one or if the reactions do not follow the laws of mass action kinetics, then the propensities might be non-affine. In such cases, the arguments illustrated in the previous section cannot be applied. To overcome this problem we first approximate the infinite dimensional controlled CME (6.14) with a finite dimensional system, similarly to what done for uncontrolled reactions in the FSP method described in Section 6.1.1. Then we show how this system can be used to construct an outer approximation of the desired projected reachable set. For this method it is actually easier to work with the uncentered moments $\{\mathbb{E}[Z_s^k(t)]\}_{k=2}^{\infty}$ instead of the centered moments $\{\mathbb{E}[(Z_s(t) - \mathbb{E}[Z_s(t)])^k]\}_{k=2}^{\infty}$. Consequently, we present our results in the former case. Nonetheless, we note that, given the bijective relation between the two sets of moments, the reachable set of the centered moments can be directly recovered from the reachable set of the uncentered moments. For example, if one is interested in mean and variance the formula $\mathbb{V}[Z_s(t)] = \mathbb{E}[(Z_s(t) - \mathbb{E}[Z_s(t)])^2] = \mathbb{E}[Z_s^2(t)] - \mathbb{E}[Z_s(t)]^2$ can be used.

8.4.1 An equivalent formulation of the controlled CME

Let us introduce an ordering $\{z^j\}_{j=1}^{\infty}$ of the possible state realizations $z \in \mathbb{N}^S$ as done in Section 6.1.1. Following the same steps as in [MK06] and setting² $P_j(t) := p(z^j, t)$, the CME in (6.14) can be rewritten as an infinite dimensional linear autonomous *switched* system that has $I = |\Sigma|$ modes,

$$\dot{P}(t) = F_{\sigma(t)}P(t), \quad (8.23)$$

where $P(t) \in [0, 1]^{\infty}$ and, for each value $\sigma(t) = \bar{\sigma}^i \in \Sigma$, $F_{\bar{\sigma}^i} \in \mathbb{R}^{\infty \times \infty}$ is a Metzler matrix. Note that system (8.23) can be thought of as a Markov chain with countably many states $z^j \in \mathbb{N}^S$ and time-varying transition matrix $F_{\sigma(t)}$ or as an infinite dimensional switched affine autonomous system. From system (8.23) one can compute the evolution of the statistical (uncentered) moments of $Z(t)$, as a linear function of $P(t)$. For instance if z_s^j denotes the amount of species Z_s in the state z^j , then the mean can be obtained as $C_s^m \cdot P(t)$, where

$$C_s^m := [z_s^1, z_s^2, \dots] \quad (8.24)$$

and the second moment can be obtained as $C_s^v \cdot P(t)$, where

$$C_s^v := [(z_s^1)^2, (z_s^2)^2, \dots]. \quad (8.25)$$

Therefore the original problem can be restated as: Find the output reachable set of the infinite dimensional linear autonomous switched system (8.23), with output $y(t) = \begin{bmatrix} C_s^m \\ C_s^v \end{bmatrix} P(t)$.

²Not to be confused with the symbol used to denote the amount of protein.

Example 6.1 (cont.) *With the ordering introduced in Section 6.1.1, the moments of mRNA and protein for the gene expression system can be computed as the output of (8.23) by setting*

$$\begin{aligned} C_M^m &= [0 \ 1 \ 0 \ 2 \ 1 \ 0 \ 3 \ 2 \ \dots], \\ C_M^v &= [0 \ 1 \ 0 \ 4 \ 1 \ 0 \ 9 \ 4 \ \dots], \\ C_P^m &= [0 \ 0 \ 1 \ 0 \ 1 \ 2 \ 0 \ 1 \ \dots], \\ C_P^v &= [0 \ 0 \ 1 \ 0 \ 1 \ 4 \ 0 \ 1 \ \dots]. \end{aligned}$$

□

8.4.2 The FSP reduction technique for the controlled CME

As in the FSP method for autonomous CME we can try to approximate the behavior of the infinite Markov chain in (8.23) by constructing a reduced Markov chain that keeps track of the probability of visiting only the states in a suitable set J . Let us define the reduced order system

$$\dot{\bar{P}}_J(t) = [F_{\sigma(t)}]_J \bar{P}_J(t), \quad \bar{P}_J(0) = P_J(0), \quad (8.26)$$

where $P_J(0)$ is the subvector of $P(0)$ corresponding to the indices in J , and $[F]_J$ denotes the sub matrix of F obtained selecting only the rows and columns with indices in J . From now on, we denote by $P(T; \sigma)$ and $\bar{P}_J(T; \sigma)$ the solutions at time T of system (8.23) and system (8.26) when input σ is applied. The dependence on the initial conditions $P(0)$ and $P_J(0)$ is omitted to keep the notation compact. As in the uncontrolled case, the truncated system (8.26) is a good approximation of the original system (8.23) if most of the probability mass lies in J . However in the controlled case we need to guarantee that this happens for all possible input sequences. This intuition can be formalized as follows.

Assumption 8.4.1. *For a given finite set of state indices J , an initial condition $P_J(0)$, a given tolerance $\varepsilon > 0$ and a finite instant $T > 0$, it holds*

$$\mathbb{1}^\top \bar{P}_J(T; \sigma) \geq 1 - \varepsilon, \quad \forall \sigma(\cdot) \in \mathcal{S}.$$

□

Note that Assumption 8.4.1 holds if and only if

$$\begin{aligned} 1 - \varepsilon &\leq \min_{\sigma \in \mathcal{S}} \mathbb{1}^\top \bar{P}_J(T; \sigma) \\ \text{s.t.} \quad &\dot{\bar{P}}_J(t; \sigma) = [F_{\sigma(t)}]_J \bar{P}_J(t; \sigma), \quad \bar{P}_J(0) = P_J(0). \end{aligned}$$

Therefore, for $\mathcal{S} = \mathcal{S}_K$ this problem is equivalent to an MILP and Assumption 8.4.1 can be checked using the method illustrated in Section 8.2.3, with $c = \mathbb{1}$ and $B = \mathbb{1}$.

Under Assumption 8.4.1, the following relation between the solutions of (8.23) and (8.26) holds.

Proposition 8.4.1 (FSP for switched CME). *If Assumption 8.4.1 holds, then for every input sequence $\sigma \in \mathcal{S}_K$, it holds*

$$\begin{aligned} P_j(T; \sigma) &\geq \bar{P}_j(T; \sigma), \quad \forall j \in J \\ \|P_J(T; \sigma) - \bar{P}_J(T; \sigma)\|_1 &\leq \varepsilon. \end{aligned}$$

□

8.4.3 Relation between projected reachable sets

Let $C_s^m(j)$ and $C_s^v(j)$ be the j -th components of the vectors C_s^m and C_s^v , respectively, as defined in (8.24) and (8.25). For a given species of interest s , we denote by

$$\begin{aligned} y_1(t; \sigma) &:= \mathbb{E} [Z_s(t) \mid Z(t) \in J, \sigma(\cdot)] = \frac{\sum_{j \in J} C_s^m(j) \cdot P_j(t; \sigma)}{\sum_{j \in J} P_j(t; \sigma)} \\ y_2(t; \sigma) &:= \mathbb{E} [Z_s^2(t) \mid Z(t) \in J, \sigma(\cdot)] = \frac{\sum_{j \in J} C_s^v(j) \cdot P_j(t; \sigma)}{\sum_{j \in J} P_j(t; \sigma)} \end{aligned} \quad (8.27)$$

the mean and second order moment of $Z(t)$ conditioned on the fact that $Z(t)$ is in J and the input σ is applied. The aim of this section is to obtain an outer approximation of the output reachable set of the infinite system (8.23) with the nonlinear output (8.27), by using computations involving only the finite dimensional system (8.26). To this end, we define the two entries of the linear output of the finite dimensional system as

$$\begin{aligned} \bar{y}_1(t; \sigma) &:= \sum_{j \in J} C_s^m(j) \cdot \bar{P}_j(t; \sigma) =: \bar{C}_1 \bar{P}_J(t; \sigma) \\ \bar{y}_2(t; \sigma) &:= \sum_{j \in J} C_s^v(j) \cdot \bar{P}_j(t; \sigma) =: \bar{C}_2 \bar{P}_J(t; \sigma). \end{aligned} \quad (8.28)$$

Proposition 8.4.2. *Suppose Assumption 8.4.1 holds. Let $\mathcal{R}_T^C(x_0)$ and $\bar{\mathcal{R}}_T^{\bar{C}}(x_0)$ be the output reachable sets at time T of systems (8.23) and (8.26), with output (8.27) and (8.28), respectively. For given $\gamma, \bar{v}^a, \bar{v}^b \in \mathbb{R}$, consider the half-spaces*

$$\bar{\mathcal{H}}^a := \{\bar{y} \in \mathbb{R}^2 \mid \bar{y}_2 \leq \gamma \bar{y}_1 + \bar{v}^a\}, \quad \bar{\mathcal{H}}^b := \{\bar{y} \in \mathbb{R}^2 \mid \bar{y}_2 \geq \gamma \bar{y}_1 + \bar{v}^b\},$$

and assume that $\bar{\mathcal{R}}_T^{\bar{C}}(x_0) \subseteq \bar{\mathcal{H}}^a \cap \bar{\mathcal{H}}^b$. Set $M_1 := \max_{j \in J}(C_s^m(j))$, $M_2 := \max_{j \in J}(C_s^v(j))$,

$$\begin{aligned} [\gamma]_+ &:= \begin{cases} \gamma & \text{if } \gamma > 0 \\ 0 & \text{if } \gamma \leq 0 \end{cases}, & [\gamma]_- &:= \begin{cases} 0 & \text{if } \gamma \geq 0 \\ -\gamma & \text{if } \gamma \leq 0 \end{cases}, \\ \delta^a(\gamma) &:= \frac{2\varepsilon}{1-\varepsilon} \cdot ([\gamma]_- \cdot M_1 + M_2), & \delta^b(\gamma) &:= \frac{2\varepsilon}{1-\varepsilon} \cdot ([\gamma]_+ \cdot M_1), \\ \mathcal{H}^a &:= \{y \in \mathbb{R}^2 \mid y_2 \leq \gamma y_1 + \bar{v}^a + \delta^a(\gamma)\}, & \mathcal{H}^b &:= \{y \in \mathbb{R}^2 \mid y_2 \geq \gamma y_1 + \bar{v}^b - \delta^b(\gamma)\}, \end{aligned} \quad (8.29)$$

with ε as in Assumption 8.4.1. Then $\mathcal{R}_T^C(x_0) \subseteq \mathcal{H}^a \cap \mathcal{H}^b$. \square

Corollary 8.4.3 (The hyperplane method for system (8.23)). *Suppose Assumption 8.4.1 holds. Let $\mathcal{R}_T^C(x_0)$ be the output reachable set at time $T > 0$ of system (8.23) with output (8.27). Choose D values $\gamma_d \in \mathbb{R}$ and set $c^d := (\bar{C}_2)^\top - \gamma_d(\bar{C}_1)^\top \in \mathbb{R}^n$, with \bar{C}_1, \bar{C}_2 as in (8.28). Set*

$$\begin{aligned} \mathcal{H}_T^{C,a}(\gamma_d) &:= \{y \in \mathbb{R}^2 \mid y_2 \leq \gamma_d y_1 + \bar{v}_T^a(c^d) + \delta^a(\gamma_d)\}, \\ \mathcal{H}_T^{C,b}(\gamma_d) &:= \{y \in \mathbb{R}^2 \mid y_2 \geq \gamma_d y_1 + \bar{v}_T^b(c^d) - \delta^b(\gamma_d)\}, \end{aligned}$$

where $\bar{v}_T^a(c^d), \bar{v}_T^b(c^d)$ are computed as in (8.13), using the finite dimension system (8.26), and $\delta^a(\gamma_d), \delta^b(\gamma_d)$ are computed as in (8.29). Then the set

$$\mathcal{R}_T^{C,out}(x_0) := \bigcap_{d=1}^D \{\mathcal{H}_T^{C,a}(\gamma_d) \cap \mathcal{H}_T^{C,b}(\gamma_d)\}$$

is an outer approximation of $\mathcal{R}_T^C(x_0)$.

Proof. This is a direct consequence of Corollary 8.2.4 and Proposition 8.4.2. \square

8.5 Appendix

8.5.1 Proofs of the results stated in Section 8.2

Proof of Proposition 8.2.2

The proof follows the same lines as [GK91, Lemma 2.1 and Theorem 2.1], where the same reachability problem is studied when the constraint sets are subsets of \mathbb{R} instead of $\mathbb{R}_{\geq 0}$ and symmetric, that is $\Sigma_r := [-\bar{\sigma}_r, \bar{\sigma}_r]$. We notice that if (A, b_r) is reachable for any $r \in \mathbb{Z}[1, R]$ then there are no singular arcs [Lib11, Section 4.4.2], therefore the control laws $\sigma^{a,*}(t)$ and $\sigma^{b,*}(t)$ are bang-bang with finitely many discontinuities. Consequently, formula (8.8) allows one to explicitly compute $x_T^{a,*}(c)$ and $x_T^{b,*}(c)$.

Proof of Proposition 8.2.3

This proof follows the lines of Theorem 2.2 in [GK91]. $\mathcal{R}_T(x_0) \subseteq \mathcal{R}_T^{out}(x_0)$ comes from the fact that, by definition of $v_T^a(c^d)$ and $v_T^b(c^d)$, for any direction c^d , $\mathcal{R}_T(x_0) \subseteq \mathcal{H}_T^a(c^d)$ and $\mathcal{R}_T(x_0) \subseteq \mathcal{H}_T^b(c^d)$. Viceversa, each element in $\mathcal{X} := \cup_{d=1}^D \{x_T^{a,*}(c^d), x_T^{b,*}(c^d)\}$ belongs to $\mathcal{R}_T(x_0)$, which is a convex set. Therefore, $\mathcal{R}_T^{in}(x_0) = \text{conv}(\mathcal{X}) \subseteq \mathcal{R}_T(x_0)$.

To approximate the infinite-time reachable set *from the origin*, consider an instant $\hat{T} > T$ and define $\varepsilon^a(T, \hat{T}, c) = v_{\hat{T}}^a(c) - v_T^a(c)$ and $\varepsilon^b(T, \hat{T}, c) = v_{\hat{T}}^b(c) - v_T^b(c)$, so that the values $v_{\hat{T}}^{a/b}(c)$ at time $\hat{T} > T$ can be computed by adding $\varepsilon^{a/b}(T, \hat{T}, c)$ to the values $v_T^{a/b}(c)$ at time T . Note that $\mathcal{R}_T(0) \subseteq \mathcal{R}_{\hat{T}}(0)$ implies $\varepsilon^a(T, \hat{T}, c) \geq 0$ and $\varepsilon^b(T, \hat{T}, c) \leq 0$ (see also [GK91, Lemma 2.2]). Moreover,

$$\begin{aligned} \varepsilon^a(T, \hat{T}, c) &= \sum_{r=1}^m \bar{\sigma}_r \int_0^{\hat{T}-T} \left[c^\top e^{A(\hat{T}-t)} b_r \right]_+ dt \leq \sum_{r=1}^m \bar{\sigma}_r \int_0^{\hat{T}-T} |c^\top e^{A(\hat{T}-t)} b_r| dt =: \epsilon(T, \hat{T}), \\ -\varepsilon^b(T, \hat{T}, c) &= \sum_{r=1}^m \bar{\sigma}_r \int_0^{\hat{T}-T} \left[c^\top e^{A(\hat{T}-t)} b_r \right]_- dt \leq \sum_{r=1}^m \bar{\sigma}_r \int_0^{\hat{T}-T} |c^\top e^{A(\hat{T}-t)} b_r| dt =: \epsilon(T, \hat{T}). \end{aligned}$$

From [GK91, Theorem 2.2] we get $\epsilon(T, \hat{T}) \leq \varepsilon_T(c)$. Therefore, for any $\hat{T} > T$, $v_{\hat{T}}^a(c) = v_T^a(c) + \varepsilon^a(T, \hat{T}, c) \leq v_T^a(c) + \varepsilon_T(c)$ and $v_{\hat{T}}^b(c) = v_T^b(c) + \varepsilon^b(T, \hat{T}, c) \geq v_T^b(c) - \varepsilon_T(c)$. Consequently, for any $d = 1, \dots, D$, $\mathcal{R}_{\hat{T}}(0) \subseteq \mathcal{H}_{\hat{T}}^a(c^d) := \{x \in \mathbb{R}^n \mid c^{d\top} x \leq v_{\hat{T}}^a(c^d)\} \subseteq \mathcal{H}^a(c^d)$, $\mathcal{R}_{\hat{T}}(0) \subseteq \mathcal{H}_{\hat{T}}^b(c^d) := \{x \in \mathbb{R}^n \mid c^{d\top} x \geq v_{\hat{T}}^b(c^d)\} \subseteq \mathcal{H}^b(c^d)$. Since these equations hold for any $\hat{T} \geq T$ and $\mathcal{H}^a(c^d)$ and $\mathcal{H}^b(c^d)$ do not depend on \hat{T} , we get $\mathcal{R}(0) \subseteq \mathcal{H}^a(c^d)$ and $\mathcal{R}(0) \subseteq \mathcal{H}^b(c^d)$ for all $d = 1, \dots, D$, which yields $\mathcal{R}(0) \subseteq \mathcal{R}^{out}(0)$. Finally $\mathcal{R}_T^{in}(0) \subseteq \mathcal{R}_T(0) \subseteq \mathcal{R}(0)$. \square

Proof of Corollary 8.2.4

For the given set of γ_d s the results are direct consequences of Proposition 8.2.3 with the c^d s defined as in the statement.

Proof of Proposition 8.2.5

The proof follows from the fact that the finite-time reachable set $\mathcal{R}_T(x_0)$, obtained using Σ , and the finite-time reachable set $\mathcal{R}_T^d(x_0)$, obtained using Σ^d , coincide. To prove this, let $\mathcal{R}_T^{bb}(x_0)$ be the finite-time reachable set obtained using $\Sigma_r^{bb} := \{0, \bar{\sigma}_r\}$ for any r , that is, the set of vertices of Σ . Since Σ is a convex polyhedron, by [Sus83, Theorem 8.1.2], System (8.3) with input set Σ has the bang-bang with bound of number of switching (BBNS) property. Consider now an arbitrary point $\bar{x} \in \mathcal{R}_T(x_0)$. By definition there exists an admissible control law in Σ that steers x_0 to \bar{x} . Consequently, by the BBNS property there exists a bang-bang control law in Σ^{bb} that achieves the same result with a finite number of discontinuities. Thus $\bar{x} \in \mathcal{R}_T^{bb}(x_0)$. Since this is true for any $\bar{x} \in \mathcal{R}_T(x_0)$, we get $\mathcal{R}_T(x_0) \subseteq \mathcal{R}_T^{bb}(x_0)$. From the fact that $\Sigma^{bb} \subseteq \Sigma^d \subseteq \Sigma$ we get $\mathcal{R}_T(x_0) \subseteq \mathcal{R}_T^{bb}(x_0) \subseteq \mathcal{R}_T^d(x_0) \subseteq \mathcal{R}_T(x_0)$, which concludes the proof.

Proof of Proposition 8.2.6 and Corollary 8.2.7

As done in Proposition 8.2.3 and Corollary 8.2.4. Note that the only difference among the hyperplane method for linear and switched affine systems is the method used to solve the optimization problems in (8.4) efficiently.

8.5.2 Proofs of the results stated in Section 8.4

Proof of Proposition 8.4.1

This result has been proven in [MK06] for linear systems. We extend it here to the case of switched systems with K switchings. Note that for any $\bar{\sigma}^i \in \Sigma$, $F_i := F_{\bar{\sigma}^i}$ has non-negative off diagonal elements. Hence, using the same argument as in [MK06, Theorem 2.1] it can be shown that for any index set J , and any $\tau \geq 0$

$$[\exp(F_i \tau)]_J \geq \exp([F_i]_J \tau) \geq 0, \quad \forall i \in 1, \dots, I.$$

Consider an arbitrary input sequence $\sigma(\cdot) \in \mathcal{S}_K$. We have

$$\begin{aligned} P_J(T; \sigma) &= [\prod_{k=0}^K \exp(F_{i_k}(t_{k+1} - t_k)) \cdot P(0)]_J \geq \prod_{k=0}^K [\exp(F_{i_k}(t_{k+1} - t_k))]_J \cdot P_J(0) \\ &\geq \prod_{k=0}^K \exp([F_{i_k}]_J(t_{k+1} - t_k)) \cdot P_J(0) = \bar{P}_J(T; \sigma). \end{aligned} \quad (8.30)$$

Moreover, from $1 = \sum_{y=1}^{\infty} P_y(T; \sigma) \geq \sum_{j \in J} P_j(T; \sigma) = \mathbb{1}^\top P_J(T; \sigma)$ and Assumption 8.4.1, we get

$$\mathbb{1}^\top \bar{P}_J(T; \sigma) \geq 1 - \varepsilon \geq \mathbb{1}^\top P_J(T; \sigma) - \varepsilon. \quad (8.31)$$

Combining (8.30) and (8.31) yields $0 \leq \mathbb{1}^\top P_J(T; \sigma) - \mathbb{1}^\top \bar{P}_J(T; \sigma) \leq \varepsilon$, thus $\|P_J(T; \sigma) - \bar{P}_J(T; \sigma)\|_1 \leq \varepsilon$.

Proof of Proposition 8.4.2

Let us first focus on the conditional mean. Choose a sequence $\sigma(\cdot) \in \mathcal{S}$. For simplicity we will omit the dependence on $(T; \sigma)$ in P_j, \bar{P}_j, y and \bar{y} . By taking into account the following conditions: (1) $C_s^m(j) \geq 0$ for all $j \in J$; (2) $P_j \geq \bar{P}_j$ for all $j \in J$, due to Proposition 8.4.1, and (3) $\sum_{j \in J} P_j \leq 1$, we get $y_1 \geq \bar{y}_1$. Consequently, at time $t = T$ we have

$$\begin{aligned} |y_1 - \bar{y}_1| &= y_1 - \bar{y}_1 = \frac{\sum_{j \in J} C_s^m(j) \cdot P_j}{\sum_{j \in J} P_j} - \sum_{j \in J} C_s^m(j) \cdot \bar{P}_j \leq \frac{\sum_{j \in J} C_s^m(j) \cdot P_j}{1 - \varepsilon} - \sum_{j \in J} C_s^m(j) \cdot \bar{P}_j \\ &= \left(1 + \frac{\varepsilon}{1 - \varepsilon}\right) \sum_{j \in J} C_s^m(j) \cdot P_j - \sum_{j \in J} C_s^m(j) \cdot \bar{P}_j \\ &= \frac{\varepsilon}{1 - \varepsilon} \sum_{j \in J} C_s^m(j) \cdot P_j + \sum_{j \in J} C_s^m(j) \cdot (P_j - \bar{P}_j) \\ &\leq M_1 \left(\frac{\varepsilon}{1 - \varepsilon} \sum_{j \in J} P_j + \sum_{j \in J} (P_j - \bar{P}_j) \right) \\ &\leq M_1 \left(\frac{\varepsilon}{1 - \varepsilon} + \|P_J - \bar{P}_J\|_1 \right) \leq M_1 \left(\frac{\varepsilon}{1 - \varepsilon} + \varepsilon \right) \leq M_1 \frac{2\varepsilon}{1 - \varepsilon}, \end{aligned}$$

where we used $\sum_{j \in J} P_j \geq \sum_{j \in J} \bar{P}_j \geq 1 - \varepsilon$ (due to Assumption 8.4.1), and $P_j \geq \bar{P}_j$, $\|P_J - \bar{P}_J\|_1 \leq \varepsilon$ (following from Proposition 8.4.1). Summing up, $\bar{y}_1 \leq y_1 \leq \bar{y}_1 + M_1 \frac{2\varepsilon}{1-\varepsilon}$. Similarly, it can be proven that $\bar{y}_2 \leq y_2 \leq \bar{y}_2 + M_2 \frac{2\varepsilon}{1-\varepsilon}$. The previous relations imply that if $\gamma \geq 0$ and $(\bar{y}_1, \bar{y}_2) \in \bar{\mathcal{H}}^a$,

$$y_2 \leq \bar{y}_2 + M_2 \frac{2\varepsilon}{1-\varepsilon} \leq \gamma \bar{y}_1 + \bar{v}^a + M_2 \frac{2\varepsilon}{1-\varepsilon} \leq \gamma y_1 + \bar{v}^a + M_2 \frac{2\varepsilon}{1-\varepsilon} = \gamma y_1 + \bar{v}^a + \delta^a(\gamma)$$

while for $(\bar{y}_1, \bar{y}_2) \in \bar{\mathcal{H}}^b$ we have

$$y_2 \geq \bar{y}_2 \geq \gamma \bar{y}_1 + \bar{v}^b \geq \gamma y_1 + \bar{v}^b - \gamma M_1 \frac{2\varepsilon}{1-\varepsilon} = \gamma y_1 + \bar{v}^b - \delta^b(\gamma).$$

On the other hand, when $\gamma < 0$ and $(\bar{y}_1, \bar{y}_2) \in \bar{\mathcal{H}}^a$ then

$$y_2 \leq \bar{y}_2 + M_2 \frac{2\varepsilon}{1-\varepsilon} \leq \gamma \bar{y}_1 + \bar{v}^a + M_2 \frac{2\varepsilon}{1-\varepsilon} \leq \gamma y_1 + \bar{v}^a + (M_2 - \gamma M_1) \frac{2\varepsilon}{1-\varepsilon} = \gamma y_1 + \bar{v}^a + \delta^a(\gamma),$$

while for $(\bar{y}_1, \bar{y}_2) \in \bar{\mathcal{H}}^b$ we have

$$y_2 \geq \bar{y}_2 \geq \gamma \bar{y}_1 + \bar{v}^b \geq \gamma y_1 + \bar{v}^b - \delta^b(\gamma).$$

Therefore for every sequence σ it holds

$$\gamma y_1(T; \sigma) + \bar{v}^b - \delta^b(\gamma) \leq y_2(T; \sigma) \leq \gamma y_1(T; \sigma) + \bar{v}^a + \delta^a(\gamma)$$

and consequently $[y_1(T; \sigma); y_2(T; \sigma)]^\top \in \mathcal{H}^a \cap \mathcal{H}^b$.

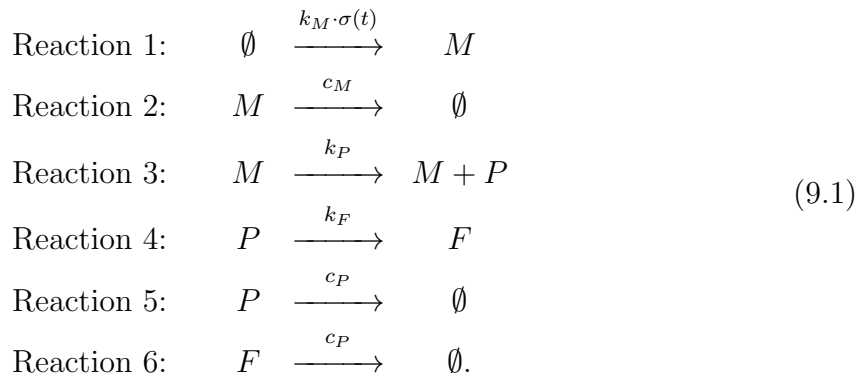
From theory to practice: a light-inducible gene expression circuit

Methods for parameter inference (see Section 7.1) and optimal experiment design (see Section 7.3) for stochastic models have been developed and applied to a number of biological systems. However, a systematic characterization procedure as the one described in Section 7.4, that exploits the information gained from each performed experiment, has not been fully developed or experimentally validated yet. In this chapter, we provide the first study in which a noisy biochemical reaction network is characterized, and ultimately also controlled, through iterations of optimally designed flow-cytometry experiments and Bayesian inference. Specifically, we consider a gene expression circuit that has been engineered in yeast so that the expression of the gene can be induced and inhibited by exposure of the cells to red and far-red light [SSHTQ02, MASSO⁺11]. We use the sequential experiment design procedure proposed in Chapter 7 to ensure that the light-induction pattern yielding the most informative output is iteratively administered to the cells and that the most informative measurement times are chosen, until the outcome of future experiments can be predicted with low uncertainty. Ultimately, we obtain a stochastic model that is capable of predicting the response of the entire cell population to arbitrary light-induction patterns with high precision. We then use the obtained model to study the fundamental limits on noise suppression for the protein production, using the techniques discussed in Chapter 8. This allows us to *in silico* plan light-induction patterns that regulate statistics of the protein distribution to desired profiles. Our experimental results show that different reference profiles can be successfully tracked over long time horizons. In contrast to previous studies, the use of a stochastic model allows us not only to regulate population averages as done in [MASSO⁺11, OHL⁺14, OT14] or individual cells as in [UMD⁺12], but also the variability across the population.

The results of this chapter have been published in [RPMA⁺15]. Earlier results in this direction were also reported in [Rue14].

9.1 The system

We consider the engineered gene expression circuit presented in [MASSO+11]. The main component of this system is a light-responsive phytochrome/phytochrome-interaction-factor (Phy/PIF) module [SSHTQ02] that can be used to drive the expression of a yellow fluorescent protein (YFP) reporter by shining red and far-red light on a population of yeast cells. We model this system with the biochemical reaction network illustrated in Figure 9.1, that consists of the following reactions:



This reaction network is similar to the one used in [MASSO+11] but contains a crucial difference in the way the external signal $\sigma(t)$ (i.e. the light input) is incorporated. Here, we assume that the mRNA production rate is multiplied by a signal $\sigma(t) = \sigma(t; L)$ of the form

$$\sigma(t; L) = U \frac{e^{-d_r(t-t_c)}}{e^{-d_r(t-t_c)} + h},$$

where L is the applied light-pattern. When a red pulse is applied, t_c is set to the current time and U is set to one, so that mRNA transcription takes place with maximum rate. The unknown parameters d_r and h capture the natural decay of the signal after a red pulse due to dark reversion [MASSO+11]. When a far-red pulse is applied U is set to zero, so that transcription is arrested until a new red pulse is applied. See [RPMA+15, SI Section S.2.1.1] for a detailed derivation.

Finally, we assume that the fluorescent intensity $I(t)$ of a sample of the population can be measured via flow-cytometry. Analogously to what done in Section 8.3.2 we assume that the fluorescence intensity $I(t)$ is proportional to the amount of fluorescent proteins $F(t)$ via a deterministic and unknown scaling parameter r . In other words, we assume $I(t) = rF(t)$.

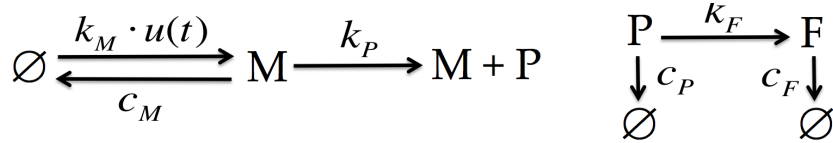
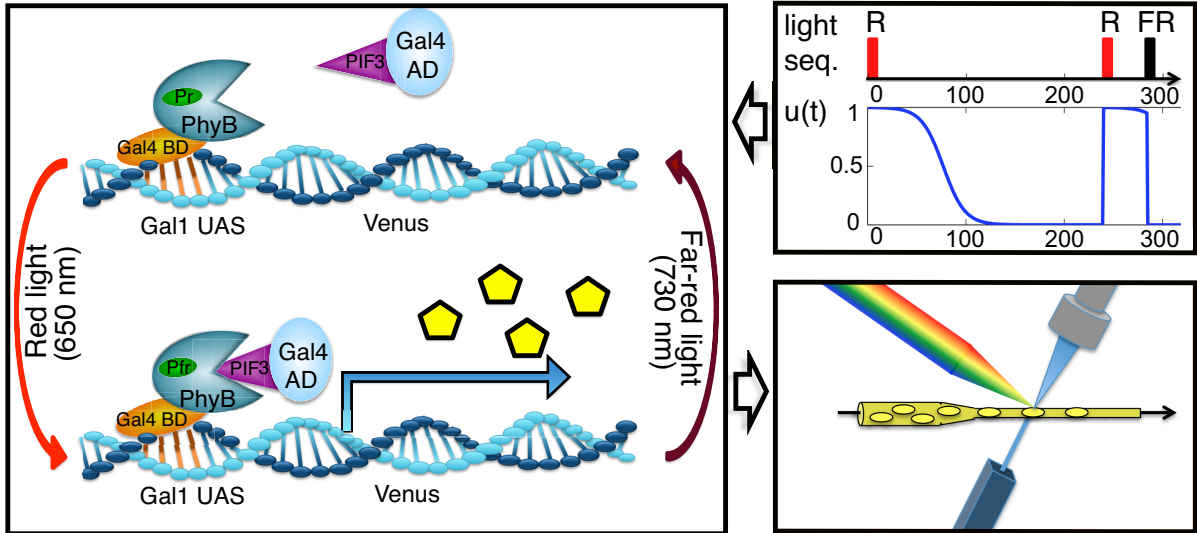


Figure 9.1: Stochastic modeling of the light-inducible gene expression circuit. The binding and unbinding of PhyB-PIF3 (left panel), caused by the light pulses, is modeled by multiplying the mRNA production rate by the signal $\sigma(t; L)$ (upper right panel). When the signal is active, mRNA M is produced with a rate $k_M \cdot \sigma(t; L)$. To capture cell-to-cell variability in the light responsive module we assume that k_M varies between different cells according to a gamma distribution P_{k_M} with unknown mean M_{k_M} and variance V_{k_M} . When mRNA is present, protein P is produced with rate k_P and then becomes fluorescent with rate k_F . All the species degrade: the mRNA (M) with rate c_M and the dark (P) and fluorescent (F) protein with rate $c_P = c_F$, as detailed in the bottom part of the figure. We assume that each fluorescent protein molecule emits an unknown but deterministic amount r of fluorescence. The fluorescence distribution in the cell population is recorded over time using flow cytometry (lower right panel). In total, the model (bottom part of the figure) comprises 3 species, 6 reactions, and 9 unknown parameters $\theta = [M_{k_M} \ V_{k_M} \ k_P \ k_F \ c_M \ c_P \ d_r \ h \ r]^T$.

9.2 The CME for heterogeneous populations

To capture variability in the light-responsive module we model the mRNA production rate k_M as a random variable that is distributed according to a gamma distribution¹ P_{k_M} with unknown mean M_{k_M} and variance V_{k_M} (see [RMAL13]). This assumption means that we allow the parameter k_M to vary among cells, but not in time. Under our assumption, the time evolution of the amount of molecules in an individual cell can be described by a continuous time Markov chain conditioned on the value of k_M in that cell. Consequently, we obtain a conditional controlled CME

$$\begin{aligned} \frac{d}{dt}p(z, t|k_M) &= -p(z, t|k_M)\alpha_1(\bar{\theta}_1, z)\sigma(t; L) + p(z - \nu_k, t|k_M)\alpha_1(\bar{\theta}_1, z - \nu_1)\sigma(t; L) \\ &\quad - \sum_{r=2}^6 p(z, t|k_M)\alpha_r(\bar{\theta}_r, z) + \sum_{r=2}^6 p(z - \nu_r, t|k_M)\alpha_r(\bar{\theta}_r, z - \nu_r), \end{aligned}$$

where $p(z, t|k_M)$ is the probability that z molecules are present at time t conditional on the value of the parameter k_M in the cell, $\nu_r \in \mathbb{Z}^3$, $r = 1, \dots, 6$ are the stoichiometric transition vectors of the 6 chemical reactions in (9.1) and $\bar{\theta}_r$ are the reaction rates $\bar{\theta} := [k_M \ k_P \ k_F \ c_M \ c_P \ c_P]^\top$. By integrating (9.2) over all possible values of k_M with respect to the probability measure P_{k_M} and simultaneously multiplying by different polynomials in z and summing over all possible values of z , similarly to what done in Section 6.1.2, we can derive a system of population moment equations from (9.2) (see [ZRK⁺12]) that depends on those rate parameters that are fixed for all cells in the population and on the moments of the distribution P_{k_M} . Since we assumed that P_{k_M} is an unknown gamma distribution parametrized by its mean and variance, we obtain a system of population moment equations, that depends on a new population parameter vector $\bar{\theta}^P = [M_{k_M} \ V_{k_M} \ k_P \ k_F \ c_M \ c_P \ d_r \ h]^\top$, of the form

$$\dot{x}(t \mid \bar{\theta}^P, e) = A(\bar{\theta}^P, \sigma(t; \bar{\theta}^P, L_e)) x(t \mid \bar{\theta}^P, e) + B(\bar{\theta}^P, \sigma(t; \bar{\theta}^P, L_e)), \quad (9.2)$$

where L_e is the light-induction pattern applied in experiment e and $x(t \mid \bar{\theta}^P, e)$ is a vector which comprises moments up to a desired order (in our case four) of the joint distribution of $Z(t)$ and the parameter k_M in experiment e assuming the vector of parameters $\bar{\theta}^P$. Note that Corollary 8.3.2 does not directly apply here because we assumed that k_M is a stochastic parameter. Nonetheless, the moment equation system (9.2) turns out to be closed and with a switched affine structure similar to system (8.18). Hence, its solution

¹The gamma distribution is a generalization of the exponential distribution and can be used to approximate a wide variety of shapes.

²Note that $\bar{\theta}$ contains a fixed reaction rate k_M , while $\bar{\theta}^P$ contains the first two moments of the distribution of k_M across the population, since we assumed it to be heterogeneous. Consequently, one may think of $\bar{\theta}$ as the parameter vector characterizing a single cell of the population (which is different from cell to cell if these are heterogeneous) and $\bar{\theta}^P$ as the unique parameter vector characterizing the whole population.

can be computed numerically using standard solvers for ordinary differential equations. Finally, the moments of the measured fluorescence intensity $I(t)$ can be easily derived from the moments of $F(t)$, since $x_{Fi}(t | \theta, e) = r^i x_{Fi}(t | \bar{\theta}^p, e)$, where $\theta = [\frac{\bar{\theta}^p}{r}]$ is the final vector of parameters to be identified, that are, the population parameters as in $\bar{\theta}^p$ and the deterministic parameter r denoting the fluorescence intensity emitted by a single fluorescent protein, as explained in Section 8.3.2.

9.3 Systematic parameter inference procedure

To optimally identify the model parameters we use the iterative characterization procedure described in Section 7.4 and illustrated for the light inducible gene expression circuit in Figure 9.2A. Specifically, we designed a first optimal experiment (Figure 9.2B), based on an initial estimate of the parameters $\hat{\theta}^0$ taken from the literature (Table 9.1). We then administered the resulting light-induction pattern to the cells, using a custom-built LED-based light delivery system, and measured by flow cytometry the resulting fluorescence intensity at the optimal measurement times. Subsequently, the data was processed and used in the inference algorithm to determine the parameter posterior distribution. Figure 9.2B shows that the model output computed using the corresponding maximum a posteriori (MAP) estimates $\hat{\theta}^1$ agrees well with the means and variances of the measured fluorescence distributions. This, however, does not guarantee that the MAP estimates can be used to predict the outcome of new experiments. Indeed, the parameter posterior distribution (Figure 9.3 top row) is flat in some dimensions, indicating that some of the parameters are practically unidentifiable from the data measured in the first experiment only. Based on these considerations, we concluded that one experiment is not sufficient to characterize the system. Consequently, we designed a second experiment that, according to the Fisher information matrix computed with the MAP estimates $\hat{\theta}^1$, optimally complements the already performed one. The resulting light-induction pattern and measurement times are shown in Figure 9.2C. After performing the second experiment, we again used Bayesian moment-based inference to update the parameter posterior distribution. The resulting distribution shows (Figure 9.3 bottom row) that additional certainty about the model parameters was gained from the second experiment.

Table 9.1: Initial estimates $\hat{\theta}^0$ of the model parameters and MAP estimates obtained from the first optimal experiment $\hat{\theta}^1$ and from both optimal experiments $\hat{\theta}^2$

	γ_M	k_P	k_F	γ_P	d_r	h	M_{k_M}	V_{k_M}	r^{-1}
$\hat{\theta}^0$	0.03	0.22	0.0419	0.0066	0.0155	0.5	0.9	0.27	10^{-3}
$\hat{\theta}^1$	0.0124	3159.9	0.0022	0.0465	0.2677	$1.34 \cdot 10^{-10}$	0.1185	0.0045	2583.1
$\hat{\theta}^2$	0.0322	731.1953	0.0300	0.0114	0.2741	$1.74 \cdot 10^{-10}$	0.1484	0.0080	17171

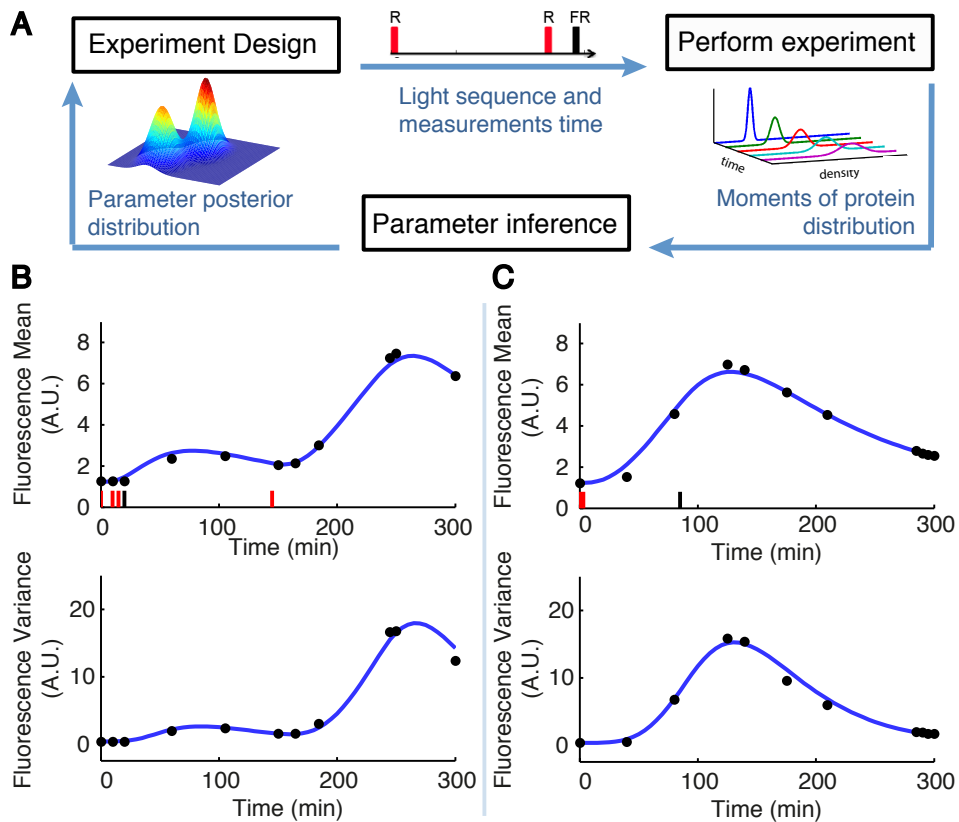


Figure 9.2: Optimal characterization of the light-inducible gene expression circuit. (A) Illustration of the iterative experiment design scheme. (B) Applied light-induction pattern (red and black bars) and measured means and variances in the first optimal experiment (black dots). The blue line is the model output with the MAP estimates $\hat{\theta}^1$ obtained from the data of this experiment. (C) Applied light-induction pattern and measured means and variances in the second optimal experiment. The blue line is the model output with the MAP estimates $\hat{\theta}^2$ obtained from the data of the two optimal experiments.

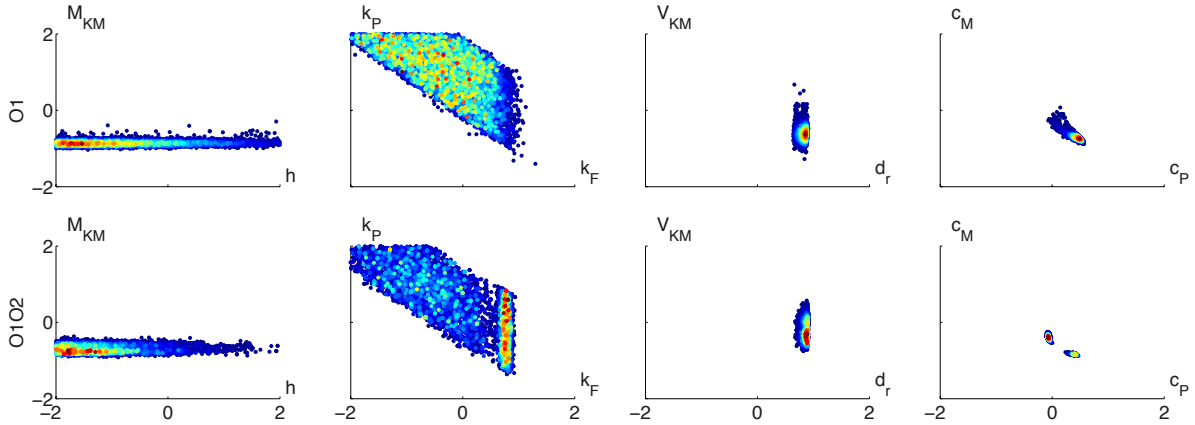


Figure 9.3: Comparison of the posterior distributions computed from the data collected in the first optimal experiment O1 (top row) and from both the optimal experiments O1O2 (bottom row). The different panels show some of the two dimensional marginals of the full posterior distribution of the parameters θ . The color is an index of the likelihood of each particle: blue for the particles with lower likelihood and red for the particles with higher likelihood.

To determine whether the residual prediction uncertainty is sufficiently small to terminate the iterative procedure, we used the obtained model to predict the outcome of a 10-hour experiment, with a randomly chosen light pattern. In particular, to quantify how the uncertainty in the posterior distribution of the model parameters influences the prediction of future experiments, we computed the posterior predictive distribution (as in Section 7.2). Figure 9.4A shows the 98% confidence region, for both the fluorescence mean and variance, computed from the obtained posterior predictive distributions. We judged these confidence regions to be sufficiently tight to terminate the iterative procedure.

To validate the obtained model we performed the experiment of Figure 9.4A and verified that the means and variances of the fluorescence distributions, measured every 30 minutes, lie within or very close to the precomputed confidence regions. We further validated the model by comparing the entire predicted fluorescence distribution to the measured one at different times (Figure 9.4B); the model predictions were obtained by simulating the system using Gillespie’s stochastic simulation algorithm [Gil76] with the MAP estimates $\hat{\theta}_2$. The results agree very well with the experimentally measured distributions, indicating that the model is capable of predicting entire population distributions, even though only sample means and variances were used in the identification.

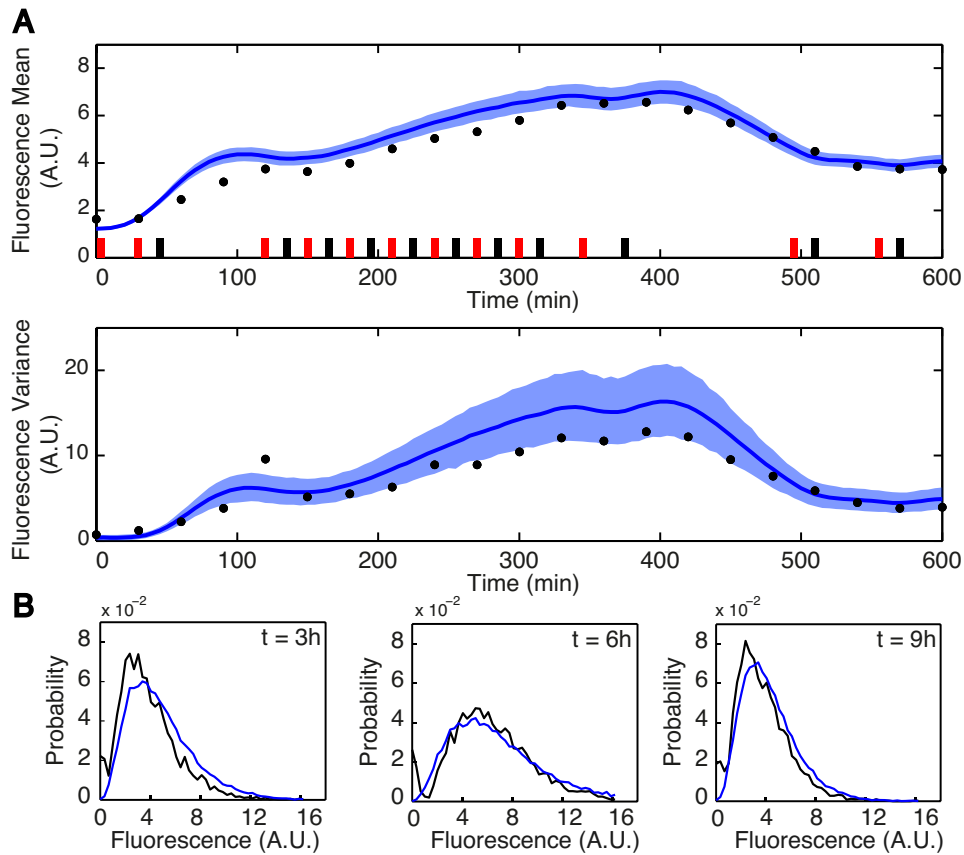


Figure 9.4: Validation of the identified model. (A) Measured and predicted mean (top) and variance (bottom) of the fluorescence distribution in a validation experiment. Model predictions are visualized in terms of the means (solid line) and 98% confidence regions (shaded region) of the posterior predictive distributions. (B) The measured fluorescence distributions (black) agree very well with simulated distributions (blue) obtained with the MAP estimates $\hat{\theta}^2$.

9.4 Comparison with random and experience-based experiments

Random experiments

The results of the previous section show that our iterative characterization procedure leads to a predictive model after only two experiments. To demonstrate that optimal experiment design is indeed necessary to obtain this result, we performed two experiments of the same duration and with the same number of measurements (equally spaced) as the optimal experiments but with randomly chosen light-induction patterns ([RPMA⁺15, SI Section S.7.1]). The parameter posterior distribution, computed from the resulting data, shows that the random experiments convey much less information than the optimal ones, leading to large residual uncertainty about the parameter values ([RPMA⁺15, SI Figure S.8]).

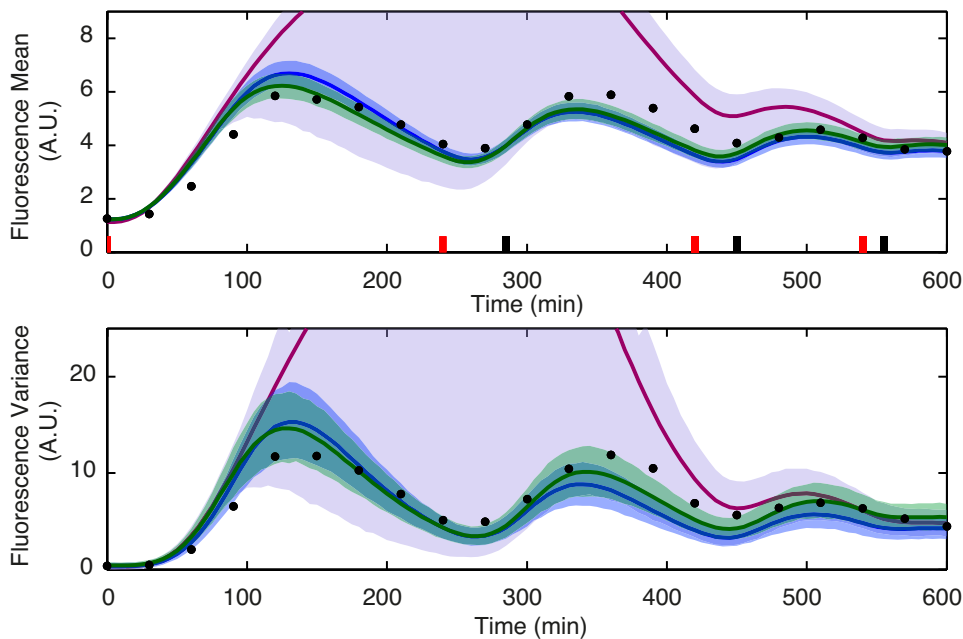


Figure 9.5: Comparison of the posterior predictive distributions computed using the parameter posterior distribution obtained from the two optimally designed (blue), the two best experience-based (green) and two random (magenta) experiments, for a second validation experiment. Model predictions are visualized in terms of the means (solid line) and 98% confidence regions (shaded region) of the posterior predictive distributions. The light sequence was chosen to produce damped oscillations of the mean fluorescence. The means and variances of the measured fluorescence distributions are shown with black dots.

Figure 9.5 shows that the model identified from the two random experiments cannot adequately predict the outcome of a validation experiment, that is, the large uncertainty remaining in the parameter posterior distribution propagates to the predictive distributions of the fluorescence mean and variance. Adding a third random experiment improves the situation only marginally ([RPMA⁺15, SI Figures S.8 and S.9]). According to our experience, a large number of random experiments would be required to obtain an accurate model of this system.

Experience-based experiments

Because the light-inducible gene expression circuit is a relatively simple system, it is also reasonable to design experiments based on intuition/experience only. It is obviously a subjective matter which experiments should be termed intuitively good for the characterization of this system. We decided that the most objective choice was to use the experiments performed for the identification of this system in [MASSO⁺11]. Hence, we chose three of the experiments shown in [MASSO⁺11, Figure 1] (one from each panel), applied the corresponding light-induction patterns to the cell population, and measured the fluorescence for 5 hours every 30min, as done for the random experiments. The posterior distribution computed from the resulting data ([RPMA⁺15, SI Figure S.11]) and the corresponding model predictions ([RPMA⁺15, SI Figure S.12]) show that any combination of only two experience-based experiments leads on average to worse results than the two optimal experiments.

Table 9.2 gives a summarizing comparison of how well a number of different experiments are predicted by the models obtained from the optimal, the random and two different pairs of experience-based experiments. From this table it can be seen that the performance of the experience-based approach depends strongly on the particular choice of the pair of experiments. Furthermore, the model identified from the two optimal

	Optimal	Best experience-based	Worst experience-based	Random
Figure 9.4	-28.8*	-123.0	-335.0	-206.2
Figure 9.5	-69.6	30.6*	20.2	-115.9
Figure 9.6A	-44.7*	-125.1	-356.5	-225.6
Figure 9.6B	-174.4*	-210.0	-342.3	-310.7
Figure 9.6C	-67.7*	-135.9	-315.1	-261.4
Figure 9.6D	-67.1*	-71.2	-131.5	-143.6

Table 9.2: Comparison of optimal, random and experience-based experiments. For each performed experiment the log of the mean likelihood of the measured data, according to the different parameters posterior distributions, is computed. The best model, that is the one with highest expected likelihood, is labeled with a star.

experiments outperforms the one identified from the best pair of experience-based experiments in 5 cases out of 6. We conclude that, for this system, experimental effort can be saved if optimal experiment design is used.

9.5 Regulating gene expression

Our final model of the gene expression circuit appears to be sufficiently accurate to predict moments of the fluorescence distribution for any light-induction pattern. Consequently, we can use it to regulate statistics of the amount of fluorescent protein in the population.

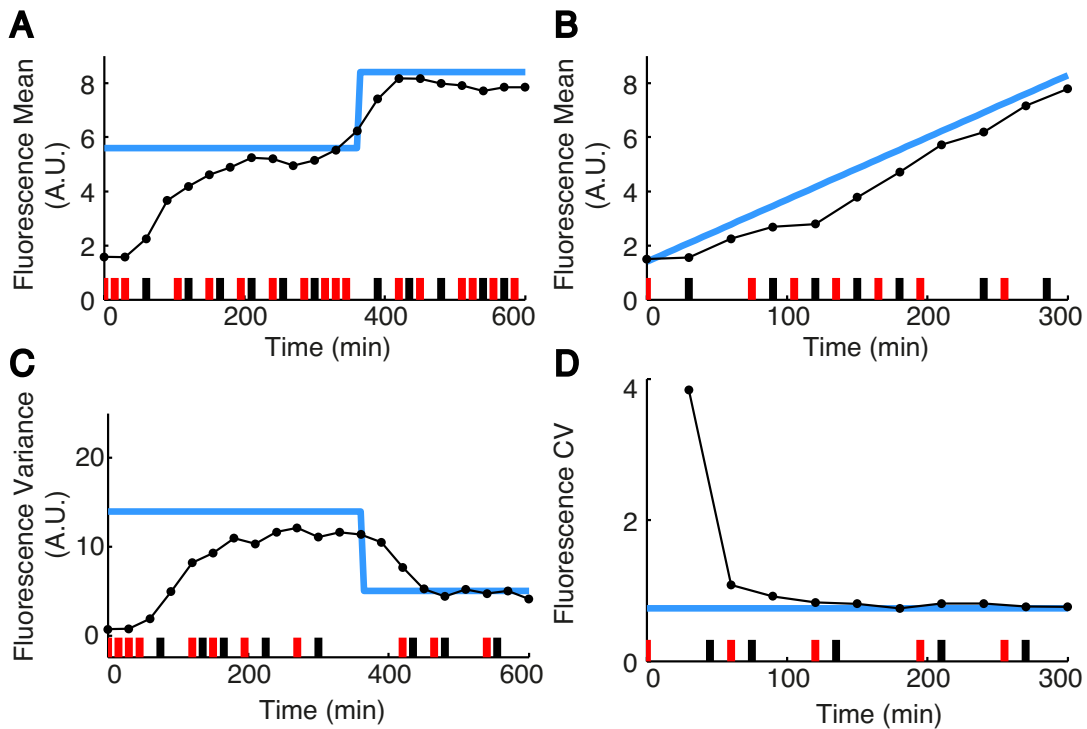


Figure 9.6: Regulation of population statistics. In all panels the light blue line is the reference time course and the black dots are the measured data. The red and black bars show the applied light pattern, which was a priori designed to make statistics of the population follow the desired reference (in light blue). (A) As a consequence of the applied light pattern, the mean of the fluorescence distribution follows the desired piecewise constant reference. (B) The mean of the fluorescence distribution follows a ramp. (C) The variance of the fluorescence distribution follows a piecewise constant reference. (D) The coefficient of variation of the fluorescence distribution is regulated to a constant value.

To illustrate this point, we used the MAP estimates $\hat{\theta}^2$, identified from the two optimal experiments, to compute two light-induction patterns that, when applied to the population, make the mean of the fluorescence distribution (seeing as an output) follow two different reference profiles. Figure 9.6A and Figure 9.6B show that a very good tracking of the reference is achieved in both experiments.

Given that our stochastic model can be used to predict higher order statistics of the fluorescence distribution, we can also choose reference time courses for other population statistics. Figure 9.6C and 9.6D show two experiments in which references for the variance and the coefficient of variation of the fluorescence distribution are tracked. For the model under consideration we found that, with the red and far-red light as the only control inputs, it is practically impossible to independently regulate the mean and the variance of the fluorescence distribution. This observation motivates us to study the reachable set of fluorescent protein mean and variance.

9.6 Reachability analysis

To investigate the reachability properties of the light inducible gene expression system we assume that the external input sequence $\sigma(t)$ takes value in $\Sigma = \{0, 1\}$. In other words, for the purpose of this section, we disregard the dark reversion process (i.e., the natural decay of the signal after a red pulse is applied, as described in Section 9.1). Under this assumption the system of population moment equations given in (9.2) becomes a switched affine system (with $I = 2$), allowing us to apply the theory developed in Section 8.2.3. For our investigation we use as parameter the final MAP estimator $\hat{\theta}^2$. Note that this choice results in a 60% Coefficient of Variation (CV) for k_M (i.e., $\frac{\sqrt{V_{k_M}}}{E_{k_M}} = 0.6$).

Figure 9.7A shows that the predicted outer approximation of the reachable set is in agreement with the measured data. An interesting aspect to notice is that, most of the measured data lies in the bottom left region of the estimated reachable set. This phenomenon can be explained by noticing that:

1. in the performed experiments the input signal $\sigma(t)$ is at the zero value (OFF) for most of the time. This means that the gene expression system is only mildly induced, thus leading to low protein production levels (low mean). For the system under investigation, input sequences that are mostly ON could not be applied since continuous red light administration would damage the cells.
2. the outer approximation of the reachable set obtained with the method of Section 8.3 is by construction convex. On the other hand, the reachable set of a switched system does not need to be. The fact that the hyperplanes approximating the reachable set from above (red lines) are all tangent at either one of the two extremes is a strong evidence that for this model the reachable set is actually

not convex. We therefore conclude that the outer approximation is not tight, in line with the experimental results suggesting that mean and variance cannot be controlled independently.

In Figure 9.7B, we investigate how the reachable set changes if we change the level of heterogeneity in the population, that is if we assume that the mean E_{k_M} is fixed and we consider different values for V_{k_M} . As to be expected, this additional source of stochasticity results in a shift of the reachable set towards regions with higher variance. Moreover, while for CV=0% we know that the reachable set is convex, the higher the CV is the less tight the outer approximation becomes.

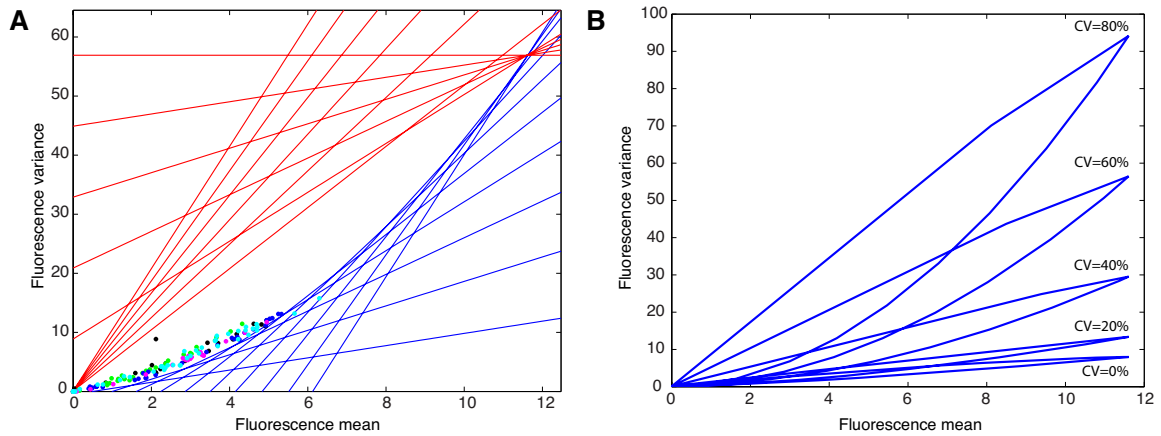


Figure 9.7: Reachable set of the light inducible gene expression model using as output the fluorescence intensity mean and variance. A) Comparison of the real data with the outer approximation of the reachable set obtained using the hyperplane method and $\hat{\theta}^2$ as parameter vector. The data are those collected in the designed (blue), random (magenta), intuitive (green), validating (black) and control (cyan) experiments. B) Comparison of the output reachable set when the parameters are as in $\hat{\theta}^2$, except for V_{k_M} which is chosen in order to obtain different levels of heterogeneity (as defined in Section 9.2).

Part III

Conclusions

Conclusions and possible extensions

10.1 Part I: populations of rational agents

In the first part of the thesis we have presented a framework to describe populations of rational agents that are profit maximizing and whose cost functions depend in an aggregative way on the strategies of the other players. Specifically, we have focused on **AAGs** (in which the cost of each player depends on the average of the strategies of the rest of the population) and on **NAGs** (in which each player has a different aggregator function that is a convex combination of the strategies of its neighbors). In both cases, motivated by real applications, we have considered the case of heterogeneous agents with multidimensional strategies and personalized constraint sets. To guarantee scalability and privacy, we have devised control solutions that rely only on global signals or local communications, while keeping the decision making process and the knowledge of the cost functions and constraint sets local.

We have discussed under which conditions dynamics that are traditionally investigated in the economic literature, as the simultaneous and sequential **BR** dynamics, are guaranteed to converge to a Nash equilibrium. We have then proposed other types of dynamics, which can be seen as an extension of the simultaneous **BR** dynamics in the case of large populations, and derived different conditions that guarantee convergence for **AAGs** (both using a central operator or local communications) and for **NAGs** (by local communications). These schemes are suitable for myopic agents that at every iterative step select the strategy that minimizes their cost function.

We have then reviewed known gradient schemes for populations of boundedly rational agents, that is agents for which the central operator can also design the strategy update law, and we have extended these algorithms to generalized Nash problems, where the agents need to satisfy an additional constraint \mathcal{C} coupling their strategies.

10.1.1 Direct extensions

Some immediate extensions of the previous results have been omitted in the interest of space and are briefly reported here.

NAGs with coupling constraints: a distributed scheme

Consider the network aggregative game with ν communications given in (4.19) with the addition of a set of coupling constraints in the neighbors aggregate states

$$\mathcal{C}_\nu := \{x \in \mathcal{R}^{Nn} \mid A^i \sigma_\nu^i(x) \leq b^i \quad \forall i \in \mathbb{Z}[1, N]\} =: \{x \in \mathcal{R}^{Nn} \mid A_\nu x \leq b_\nu\},$$

where $A^i \in \mathbb{R}^{m_i \times n}$, $b^i \in \mathbb{R}^{m_i}$ for all $i \in \mathbb{Z}[1, N]$, $A_\nu := \text{diag}(A^1, \dots, A^N) \cdot \mathcal{P}_\nu$ and $b_\nu = [b^1; \dots; b^N]$, resulting in

$$\begin{cases} \min_{x^i \in \mathbb{R}^n} & J^i(x^i, \sigma_\nu^i(x)) \\ \text{s.t.} & x^i \in \mathcal{Q}_\nu^i(x^{-i}) \end{cases} \quad \forall i \in \mathbb{Z}[1, N], \quad (10.1)$$

where $\sigma_\nu^i(x) = \sum_{j=1}^N P_{ij}^\nu x^j$, $\mathcal{Q}_\nu := \mathcal{X}_{1 \times N} \cap \mathcal{C}_\nu$ and $\mathcal{Q}_\nu^i(x^{-i}) := \{x^i \in \mathcal{X}^i \mid A_\nu x \leq b_\nu\}$. Exactly as in Section 5.3, one can use the iterative schemes illustrated in Section 3.3.2 to solve the associated VI($\mathcal{X}_{1 \times N} \times \mathbb{R}_{\geq 0}^m, t$), with t defined as in (5.9) using

$$f_\nu(x) := [\nabla_{x^i} J^i(x^i, \sigma_\nu^i(x))]_{i=1}^N$$

instead of $f(x)$, and hence find a variational GNE of the game in (10.1). We note that, in this case, the operator $f_\nu(x)$ can be evaluated locally, without the need of a central operator, since each agent can compute its component by knowing only its strategy and the neighbors aggregate state. Moreover, each agent can update locally the dual variable $\lambda^i \in \mathbb{R}^{m_i}$ associated with the coupling constraint on its neighbors aggregate state (i.e., $A^i \sigma_\nu^i(x) \leq b^i$). These two features allows one to implement the schemes of Section 3.3.2 in a totally distributed fashion in the case of NAGs, provided that $P = P^\top$.

AAGs with coupling constraints: a distributed scheme

The algorithm suggested in the previous section for NAGs can also be used to guarantee distributed convergence to almost generalized Nash equilibria in AAGs. The details are omitted in the interest of space. The main idea, however, is to follow the same procedure as in Section 4.4 and allow a sufficient number of communications in between two strategies update, over a sparse network P satisfying Assumption 4.4.1. Under suitable regularity conditions on the cost functions and on the constraints, one can in fact use [Mos69, Theorem A (b)] to show that the solution of the VI(\mathcal{Q}_ν, f_ν) associated with the NAG with ν communications converges to the solution of the VI(\mathcal{Q}, f) associated with the AAG. Consequently, one can use the scheme for NAGs to approximate the Nash equilibrium of an AAG with arbitrary precision, by selecting ν large enough.

10.1.2 Further research directions

The use of distributed algorithms to coordinate a population of agents to configurations that have a meaningful interpretation in terms of game theory (e.g. to Nash equilibria)

is a relatively new field. We conclude this chapter by highlighting possible extensions to the presented framework that may be of interest as future work.

Social optimality

Throughout this thesis our objective was to coordinate the strategies of the agents to a Nash equilibrium, which is a configuration of strategies that is fair in the sense that no agent can improve its cost by unilateral deviations. We did not investigate what are the properties of the obtained equilibrium. Specifically, concepts as social global optimality, efficiency and price of anarchy [TX14, HCM12] were not considered in this thesis. The easiest case to analyze is that of games with SMON integrable operator $f(x) := [\nabla_{x^1} J^1(x^1, x^{-1}); \dots; \nabla_{x^N} J^N(x^N, x^{-N})]$. That is, games for which there exists a strongly convex function $S(x) : \mathbb{R}^{Nn} \rightarrow \mathbb{R}$ such that $\nabla_x S(x) = f(x)$. In this case one can show that any Nash equilibrium obtained with the gradient schemes described in Chapter 5 is a minimizer of $S(x)$, which can therefore be thought of as a welfare function [PPG⁺16]. In general however there is no guarantee that the Nash equilibrium would be socially optimal. Note that, in the case of boundedly rational agents, the central operator could directly design the strategies update laws to guarantee convergence to the minimum of the social function by using standard distributed optimization tools. This would however give no guarantees at all on the performance of the single agents. In other words, it might be that the performance of a single agent is greatly sacrificed for the good of the rest of the population. Consequently, in the literature there are schemes that guarantee convergence to socially optimal solutions that might not be fair and schemes that guarantee convergence to fair solutions that might not be socially optimal. I believe that one of the most challenging future direction for this field is overcoming this dicotomy. For example, in the case of games with multiple equilibria, a first step could be the development of algorithms that steer the agents to the Nash equilibrium that minimizes an arbitrarily chosen social function or that minimizes the price of anarchy. Similarly, the problem of how to modify the original game by mechanism or utility design (i.e., by using incentives, tolls, etc.) so that the Nash equilibria of the new game have a desired social property is an interesting field of research [MS14, LM13, JT05].

Cheating

An important assumption of our work is that the agents communicate truthfully with the central operator or with their neighbors. Note that the latter do not need to be spatial neighbors. A neighborhood could be defined as the set of players that an agent trust and with whom he is willing to communicate truthfully. Moreover, in technical applications (as e.g. demand response schemes) one may envision these communications to be performed by an automatic device. Nonetheless, a future research direction is the analysis of the robustness of our schemes when part of the agents are allowed to cheat.

In this sense, the algorithms presented in Chapter 4 for myopic agents are less prone to the risk of cheating since, at every step, each agent selects the strategy that minimizes its cost. Consequently, any cheating behavior would be suboptimal at the single algorithmic step, but might still be beneficial in the long run.

Stochasticity

In this thesis we considered games that are deterministic. However, many of the applications described in Chapter 2 are characterized by uncertainty and noise. In the case of demand response applications, for example, both the demand and production profiles are affected by noise coming, for example, from the uncertainties in the weather forecasts. Moreover, the number of players could fluctuate during the game, adding additional sources of variability. These aspects could be taken into account by using a stochastic version of the games presented in this thesis where, for example, the aggregator function depends on a random variable. One could then examine games where each agent tries to minimize its expected cost. The easiest way to tackle this problem is to follow the same lines as in Section 5.1, using however schemes to solve stochastic VIs (as e.g. the ones described in [KNS13, YNS16a]) and extend the results of [YNS16b] for standard Nash games to generalized Nash games, as done in Section 5.2 for the deterministic case.

Asynchronicity

Our technical results are derived for agents that update their strategies synchronously and over a fixed network. As future work, we believe it would be interesting to study whether similar convergence results can be achieved via asynchronous updates and time-varying or random communications [BT97, FZ08, ZC10]. This was done for example for AAGs with boundedly rational agents in [KNS16], by using gossip-based communication schemes. For AAGs with myopic agents and or coupling constraints, known results from consensus theory could be applied to guarantee that the convergence requirement $\sigma_{\nu^{(k)}}^i \rightarrow \bar{\sigma}$ (see Assumption 4.4.1 and Section 4.4) is achieved asynchronously. Even more interesting would be the development of algorithms where both communications and strategies updates are performed asynchronously. Once again one could rely on asynchronous distributed algorithms for convex optimization or solution of VIs to tackle this task [BT97]. Finally, it would be interesting to study the performance of the proposed schemes in the presence of communication noise or delays, as for example done in [ZF16] in the case of monotone games.

Other types of equilibria

In this work we focused on the notion of Nash equilibrium as it encodes the concept of “fairness” and “stability” in the sense that no agent has interest in unilateral deviations.

Nonetheless, other types of equilibria have been studied in game theory and, depending on the application, may offer valid alternatives. In the case of average aggregative games, for example, an important class is that of Wardrop equilibria (also known as competitive equilibria in the economic literature). Loosely speaking a Wardrop equilibrium is a set of strategies $\{\bar{x}^j\}_{j=1}^N$ where no agent has incentive in changing its own strategy if we fix the average to $\frac{1}{N} \sum_{j=1}^N \bar{x}^j$, that is, if for all agent $i \in \mathbb{Z}[1, N]$ it holds

$$J^i(\bar{x}^i, \frac{1}{N} \sum_{j=1}^N \bar{x}^j) \leq J^i(x^i, \frac{1}{N} \sum_{j=1}^N \bar{x}^j), \quad \forall x^i \in \mathcal{X}^i.$$

The main difference between Nash and Wardrop is thus that in the latter the agent is not allowed to change its contribution to the average, but only the first argument of $J^i(\cdot, \cdot)$. Wardrop equilibria were originally defined in traffic applications, where the population is very large and consequently the contribution of a single agent to the average is negligible. One can actually show that, under mild assumptions, any Wardrop equilibrium is an ε_N -Nash equilibrium with $\varepsilon_N \rightarrow 0$ as $N \rightarrow \infty$. This is the fundamental property at the core of the proof of Theorem 4.2.1. Other types of equilibria that one may consider are Stackelberg equilibria (where the agents are divided into minor and major player that select their actions sequentially [BO99, NCMH12, MGP+16]) or correlated equilibria.

10.2 Part II: populations of biological systems

In the second part of the thesis we have presented a framework to describe populations of stochastic biochemical reaction networks. The fundamental feature that we exploited to characterize these systems is that every cell in the population can be considered as an independent realization of an underlying stochastic process. Recent technological advancements have allowed researchers to measure thousands of such cells at the same time, thus providing large samples from which reliable estimates of the moments of such a stochastic process can be computed.

Based on such moments estimates, we have proposed a sequential approach, composed by experiment design, flow cytometry experiments and parameter inference, to systematically infer the unknown parameters of controlled stochastic biochemical reaction networks. We have then proposed a framework, based on reachability theory, to analyze the behavior of such control networks and answer fundamental questions on noise suppression, for example in gene expression circuits.

10.2.1 Direct extensions

A direct extension of the presented theory is briefly discussed in this section.

Sequential experiment design for model selection

The main objective of Chapter 7 was to suggest an iterative procedure to identify the unknown parameters of a given model of the stochastic biochemical reaction network of interest. More in general one might be interested in using the measured data to discriminate among a finite number of different models $m = 1, \dots, M$, all with unknown parameters. In this case the sequential procedure described in Section 7.4 can be extended as shown in Figure 10.1.

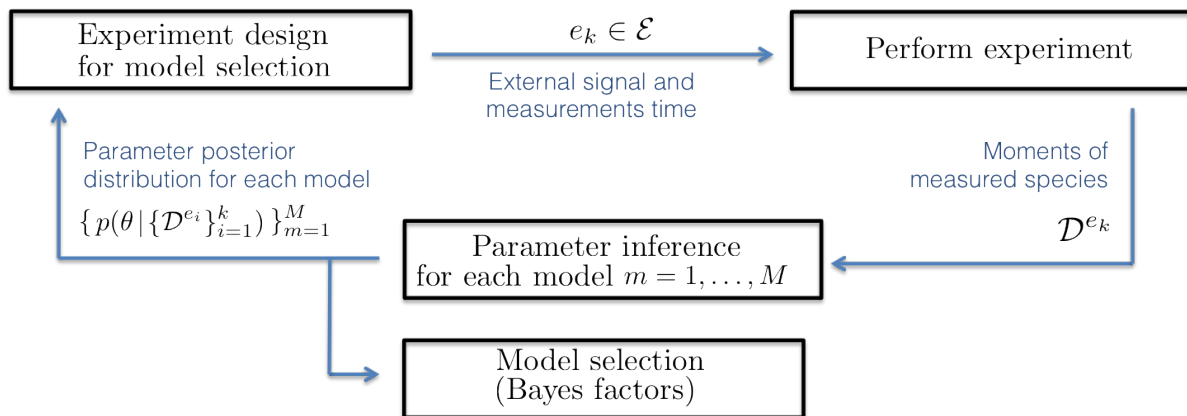


Figure 10.1: Sequential experiment design for model selection.

Note that all the steps in Figure 10.1 are the same as in Section 7.4 except for the experiment design procedure which, in this case, should aim at selecting the experiment that allows to better discriminate among the M models. To this end, one can select the experiment $e \in \mathcal{E}$ that maximizes the “distance” among the M posterior predictive distributions of the sampled mean and variance for experiment e , computed by using the m different models and the corresponding parameter posterior distributions obtained from previous experiments. Note that, in general, computing the distance among probability distributions is not a trivial task. However, as detailed in Section 7.2, the posterior predictive distributions of the sample moments can be approximated by Gaussian mixtures. One can then develop a computationally efficient experiment design procedure by using ad hoc definitions of distance for Gaussian mixtures. We finally note that the parameter posterior distributions can also be used to compute the Bayes factors among each pair of models, which can be used as test to terminate the iterative procedure.

10.2.2 Further research directions

Automatization and robustness to different environment conditions

Dynamic control of gene expression can have far-reaching implications for biotechnological applications and biological discovery. Nonetheless, it is at a very preliminary research-oriented implementation stage. The experiments described in Chapter 9, for example, were performed manually and required the constant presence of human supervision to take measurements, provide the correct pulses and, in general, guarantee a strictly regulated environment with constant temperature, density, etc. The interplay of experiment design, data collection and data inference was also performed in a custom made fashion. It is clear that to make the process of controlling gene expression technological appealing all these steps need to be automated and the proposed methods should be made robust to changes in the environment. A first step in this direction has been taken in [MARA⁺16], where an automated scheme to implement the control task described in Section 9.5 is suggested, based on a simple deterministic model. It would be of interest to devise a scheme where stochastic model calibration and re-identification of the parameters in case of environmental changes are included.

Internal vs external control

In this thesis we have focused on populations of stochastic biochemical reaction networks that can be controlled by using external signals, applied to the whole population. A different line of research has focused, on the other hand, on the design and integration of feedback loops inside the biochemical reaction network at the single cell level [BK12, BZK16, ZSRK16]. An important future direction is the comparison between these two approaches and possibly their integration.

Parameter inference and model selection in evolutionary games

We have briefly mentioned in Chapter 6 that, even though we focused on biochemical reaction networks, the tools developed in this thesis could be of interest also for other types of reaction networks. One important example is that of evolutionary games. Evolutionary game theory studies the learning process of a population of rational agents that can choose among a finite set of different strategies [Smi82, San10]. The basic assumptions are that (i) the strategy chosen by an agent results in a different payoff (i.e., reward) depending on the strategies chosen by the rest of the population and (ii) the agents iteratively change their strategies, in an attempt to maximize their payoffs. This “learning” process can be modeled by means of reactions that change the amount of players selecting a given strategy (i.e. the species) with a rate that depends on the payoff (i.e. the higher the payoff of the new strategy the oftener the reaction takes

place). Traditional learning dynamics rules, as best response, logit, replicator dynamics, etc. can be casted in this framework by choosing suitable reactions rates. The inference and model selection tools previously described could then be applied to identify and discriminate among these different rules based on population observations (e.g, surveys) and changing incentives (e.g., modifications of the payoff matrices).

Bibliography

- [AG00] F. Allen and D. Gale. Financial contagion. *Journal of Political Economy*, 108(1):1–33, 2000.
- [AKLP10] A. Abate, J.-P. Katoen, J. Lygeros, and M. Prandini. Approximate model checking of stochastic hybrid systems. *European Journal of Control*, 16(6):624–641, 2010.
- [AM92] R. M. Anderson and R. M. May. *Infectious diseases of humans: Dynamics and control*, volume 28. Wiley Online Library, 1992.
- [AO11] D. Acemoglu and A. Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49, 2011.
- [AOTS15] D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi. Networks, shocks, and systemic risk. Technical report, National Bureau of Economic Research, 2015.
- [AOY11] D. Acemoglu, A. Ozdaglar, and E. Yildiz. Diffusion of innovations in social networks. In *Proceedings of the IEEE Conference on Decision and Control and European Control Conference*, pages 2329–2334, Dec 2011.
- [ARM98] A. Arkin, J. Ross, and H. H. McAdams. Stochastic kinetic analysis of developmental pathway bifurcation in phage λ -infected *Escherichia coli* cells. *Genetics*, 149(4):1633–1648, 1998.
- [BB14] F. Bagagiolo and D. Bauso. Mean-field games and dynamic demand management in power grids. *Dynamic Games and Applications*, 4(2):155–176, 2014.
- [BC10] H. H. Bauschke and P. L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2010.
- [BCV15] F. Blanchini, P. Colaneri, and M. E. Valcher. Switched positive linear systems. *Foundations and Trends in Systems and Control*, 2(2):101–273, 2015.

- [BDJPG08] G. Batt, H. De Jong, M. Page, and J. Geiselman. Symbolic reachability analysis of genetic regulatory networks using discrete abstractions. *Automatica*, 44(4):982–989, 2008.
- [Ber07] V. Berinde. *Iterative approximation of fixed points*. Springer, 2007.
- [BFY13] A. Bensoussan, J. Frehse, and P. Yam. *Mean field games and mean field type control theory*. Springer, 2013.
- [BHHK03] C. Baier, B. Haverkort, H. Hermanns, and J.-P. Katoen. Model-checking algorithms for continuous-time Markov chains. *IEEE Transactions on Software Engineering*, 29(6):524–541, 2003.
- [BHT10] V. D. Blondel, J. M. Hendrickx, and J. N. Tsitsiklis. Continuous-time average-preserving opinion dynamics with opinion-dependent communications. *SIAM Journal on Control and Optimization*, 48(8):5214–5240, 2010.
- [BK12] C. Briat and M. Khammash. Computer control of gene expression: Robust setpoint tracking of protein mean and variance using integral feedback. In *Proceedings of the IEEE Conference on Decision and Control*, pages 3582–3588, 2012.
- [BK15] Y. Bramoullé and R. Kranton. Games played on networks. In Y. Bramoullé, A. Galeotti, and B. Rogers, editors, *The Oxford Handbook of the Economics of Networks*, chapter 5. Oxford University Press, Oxford, 2015.
- [BM99] A. Bemporad and M. Morari. Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35(3):407–427, 1999.
- [BO99] T. Başar and G. J. Olsder. *Dynamic noncooperative game theory*. Classics in Applied Mathematics. SIAM, 1999.
- [BP13] D. Bauso and R. Pesenti. Mean field linear quadratic games with set up costs. *Dynamic Games and Applications*, 3(1):89–104, 2013.
- [BT97] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and distributed computation: Numerical methods*. Athena Scientific, 1997.
- [BV04] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [BZK16] C. Briat, C. Zechner, and M. Khammash. Design of a synthetic integral feedback circuit: Dynamic analysis and DNA implementation. *ACS Synthetic Biology*, 2016.

- [CF03] N. Chabrier and F. Fages. Symbolic model checking of biochemical networks. In C. Priami, editor, *Computational Methods in Systems Biology*, pages 149–162. Springer, 2003.
- [CH12] R. Cornes and R. Hartley. Fully aggregative games. *Economics Letters*, 116(3):631–633, 2012.
- [Chi09] C. Chidume. *Geometric properties of Banach spaces and nonlinear iterations*. Lecture Notes in Mathematics. Springer, 2009.
- [CLLV14] H. Chen, Y. Li, R. H. Y. Louie, and B. Vucetic. Autonomous demand side management based on energy consumption scheduling and instantaneous load billing: An aggregative game approach. *IEEE Transactions on Smart Grid*, 5(4):1744–1754, 2014.
- [DeG74] M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, 1974.
- [DHZ06] P. Dubey, O. Haimanko, and A. Zapechelnyuk. Strategic complements and substitutes, and potential games. *Games and Economic Behavior*, 54(1):77–94, 2006.
- [DM07] A. Donzé and O. Maler. Systematic simulation using sensitivity analysis. In A. Bemporad, A. Bicchi, and G. Buttazzo, editors, *Hybrid Systems: Computation and Control*, pages 174–189. Springer, 2007.
- [DMDJ06] P. Del Moral, A. Doucet, and A. Jasra. Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(3):411–436, 2006.
- [DPK⁺15] T. Danino, A. Prindle, G. A. Kwong, M. Skalak, H. Li, K. Allen, J. Hasty, and S. N. Bhatia. Programmable probiotics for detection of cancer in urine. *Science translational medicine*, 7(289):289ra84, 2015.
- [EB15] S. R. Etesami and T. Başar. Game-Theoretic Analysis of the Hegselmann-Krause Model for Opinion Dynamics in Finite Dimensions. *IEEE Transactions on Automatic Control*, 60(7):1886–1897, 2015.
- [ESPP⁺06] H. El-Samad, S. Prajna, A. Papachristodoulou, J. Doyle, and M. Khammash. Advanced methods and algorithms for biological networks analysis. *Proceedings of the IEEE*, 94(4):832–853, 2006.
- [FJ99] N. E. Friedkin and E. C. Johnsen. Social influence networks and opinion change. *Advances in group processes*, 16(1):1–29, 1999.

- [FK07] F. Facchinei and C. Kanzow. Generalized Nash equilibrium problems. *4OR*, 5(3):173–210, 2007.
- [For05] S. Fortunato. Damage spreading and opinion dynamics on scale-free networks. *Physica A: Statistical Mechanics and its Applications*, 348:683–690, 2005.
- [FP03] F. Facchinei and J.-S. Pang. *Finite-dimensional variational inequalities and complementarity problems*. Operations Research. Springer, 2003.
- [FPR00] X. Freixas, B. Parigi, and J.-C. Rochet. Systemic risk, interbank relations, and liquidity provision by the central bank. *Journal of money, credit and banking*, 32(3):611–638, 2000.
- [FZ08] F. Fagnani and S. Zampieri. Randomized consensus algorithms over large scale networks. *IEEE Journal on Selected Areas in Communications*, 26(4):634–649, 2008.
- [Gal10] A. Galeotti. Network Games. In S. Goyal, M. O. Jackson, F. Vega-Redondo, and L. Yariv, editors, *The Review of Economic Studies*, volume 77, pages 218–244. Oxford University Press, 2010.
- [GGPL15] S. Grammatico, B. Gentile, F. Parise, and J. Lygeros. A mean field control approach for demand side management of large populations of thermostatically controlled loads. In *Proceedings of the IEEE European Control Conference*, pages 3548–3553, 2015.
- [Gil76] D. T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22(4):403–434, 1976.
- [Gil92] D. T. Gillespie. A rigorous derivation of the chemical master equation. *Physica A: Statistical Mechanics and its Applications*, 188(1–3):404–425, 1992.
- [Gil07] D. T. Gillespie. Stochastic simulation of chemical kinetics. *Annual Review of Physical Chemistry*, 58:35–55, 2007.
- [GJ13] J. Goutsias and G. Jenkinson. Markovian dynamics on complex reaction networks. *Physics Reports*, 529(2):199–264, 2013.
- [GK91] T. J. Graettinger and B. H. Krogh. Hyperplane method for reachable state estimation for linear time-invariant systems. *Journal of optimization theory and applications*, 69(3):555–588, 1991.

- [GM91] I. Gilboa and A. Matsui. Social stability and equilibrium. *Econometrica: Journal of the Econometric Society*, pages 859–867, 1991.
- [GPCL16] S. Grammatico*, F. Parise*, M. Colombino*, and J. Lygeros. Decentralized convergence to Nash equilibria in constrained deterministic mean field control. Accepted to *IEEE Transactions on Automatic Control*, to appear Nov. 2016.
- [GPL15] S. Grammatico, F. Parise, and J. Lygeros. Constrained linear quadratic deterministic mean field control: Decentralized convergence to Nash equilibria in large populations of heterogeneous agents. In *Proceedings of the IEEE Conference on Decision and Control*, pages 4412–4417, 2015.
- [GPP+17] B. Gentile*, F. Parise*, D. Paccagnan*, M. Kamgarpour, and J. Lygeros. Nash and Wardrop equilibria in aggregative games with coupling constraints. *ArXiv preprint:1702.08789*, 2017.
- [GR84] K. Goebel and S. Reich. *Uniform convexity, hyperbolic geometry, and nonexpansive mappings*. Taylor & Francis, 1984.
- [GS14] J. Ghaderi and R. Srikant. Opinion dynamics in social networks with stubborn agents: Equilibrium and convergence rate. *Automatica*, 50(12):3209–3215, 2014.
- [GTL13] L. Gan, U. Topcu, and S. H. Low. Optimal decentralized protocol for electric vehicle charging. *IEEE Transactions on Power Systems*, 28(2):940–951, 2013.
- [HBG+05] A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, and C. S. Woodward. SUNDIALS: Suite of Nonlinear and Differential/Algebraic Equation Solvers. *ACM Transactions on Mathematical Software*, 31(3):363–396, 2005.
- [HCM07] M. Huang, P. E. Caines, and R. P. Malhamé. Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ε -Nash equilibria. *IEEE Transactions on Automatic Control*, 52(9):1560–1571, 2007.
- [HCM12] M. Huang, P. E. Caines, and R. P. Malhamé. Social optima in mean field LQG control: Centralized and decentralized strategies. *IEEE Transactions on Automatic Control*, 57(7):1736–1751, 2012.
- [HCW98] T.-H. Ho, C. Camerer, and K. Weigelt. Iterated dominance and iterated best response in experimental “p-beauty contests”. *The American Economic Review*, 88(4):947–969, 1998.

- [Hei06] T. Heikkinen. A potential game approach to distributed power control and scheduling. *Computer Networks*, 50(13):2295–2311, 2006.
- [Hes08] J. Hespanha. Moment closure for biochemical networks. In *Proceedings of the IEEE International Symposium on Communications, Control and Signal Processing*, pages 142–147, 2008.
- [HK02] R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence models, analysis, and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3):1–33, 2002.
- [HLW06] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.
- [HMW09] T. A. Henzinger, M. Mateescu, and V. Wolf. Sliding window abstraction for infinite Markov chains. In A. Bouajjani and O. Maler, editors, *Computer Aided Verification*, pages 337–352. Springer, 2009.
- [HWT13a] D. R. Hagen, J. K. White, and B. Tidor. Convergence in parameters and predictions using computational experimental design. *Interface Focus*, 3:20130008, 2013.
- [HWT13b] D. R. Hagen, J. K. White, and B. Tidor. Convergence in parameters and predictions using computational experimental design. *Interface Focus*, 3(4):20130008, 2013.
- [Jac96] J. A. Jacquez. *Compartmental analysis in biology and medicine, third edition*. Biomedware, 1996.
- [Jen10] M. K. Jensen. Aggregative games and best-reply potentials. *Economic theory*, 43(1):45–66, 2010.
- [JT04] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, 2004.
- [JT05] R. Johari and J. N. Tsitsiklis. A game theoretic view of efficiency loss in resource allocation. In E. H. Abed, editor, *Advances in Control, Communication Networks, and Transportation Systems*, pages 203–223. Springer, 2005.
- [JZ14] M. O. Jackson and Y. Zenou. Games on networks. In P. Young and S. Zamir, editors, *Handbook of game theory*, volume 4. 2014.

- [KC10] A. C. Kizilkale and P. E. Caines. Stochastic adaptive Nash certainty equivalence control: Population parameter distribution estimation. In *Proceedings of the IEEE Conference on Decision and Control*, pages 6169–6176, 2010.
- [KCRS11] M. Komorowski, M. J. Costa, D. A. Rand, and M. P. H. Stumpf. Sensitivity, robustness, and identifiability in stochastic chemical kinetics models. *Proceedings of the National Academy of Sciences of the USA*, 108(21):8645–8650, 2011.
- [KMC12] A. C. Kizilkale, S. Mannor, and P. E. Caines. Large scale real-time bidding in the smart grid: A mean field framework. In *Proceedings of the IEEE Conference on Decision and Control*, pages 3680–3687, 2012.
- [KNS12] J. Koshal, A. Nedić, and U. V. Shanbhag. A gossip algorithm for aggregative games on graphs. In *Proceedings of the IEEE Conference on Decision and Control*, pages 4840–4845, 2012.
- [KNS13] J. Koshal, A. Nedić, and U. V. Shanbhag. Regularized iterative stochastic approximation methods for stochastic variational inequality problems. *IEEE Transactions on Automatic Control*, 58(3):594–609, 2013.
- [KNS16] J. Koshal, A. Nedić, and U. V. Shanbhag. Distributed algorithms for aggregative games on graphs. *Operations Research*, 64(3):680–704, 2016.
- [Kuk04] N. S. Kukushkin. Best response dynamics in finite games with additive aggregation. *Games and Economic Behavior*, 48(1):94–10, 2004.
- [KV97] A. Kurzhanski and I. Valyi. *Ellipsoidal calculus for estimation and control*. Springer, 1997.
- [Lib11] D. Liberzon. *Calculus of variations and optimal control theory: A concise introduction*. Princeton University Press, 2011.
- [LKK09] C. H. Lee, K.-H. Kim, and P. Kim. A moment closure method for stochastic reaction networks. *The Journal of Chemical Physics*, 130(13):134107, 2009.
- [LL07] J.-M. Lasry and P.-L. Lions. Mean field games. *Japanese Journal of Mathematics*, 2:229–260, 2007.
- [LM13] N. Li and J. R. Marden. Designing games for distributed optimization. *IEEE Journal of Selected Topics in Signal Processing*, 7(2):230–242, 2013.

- [Lor07] J. Lorenz. Continuous opinion dynamics of multidimensional allocation problems under bounded confidence: More dimensions lead to better chances for consensus. *European Journal of Economic and Social Systems*, 19(2):16, 2007.
- [LVP10] I. Lestas, G. Vinnicombe, and J. Paulsson. Fundamental limits on the suppression of molecular fluctuations. *Nature*, 467(7312):174–178, 2010.
- [MA79] R. M. May and R. M. Anderson. Population biology of infectious diseases: Part I. *Nature*, 280(5722):455–461, 1979.
- [MA13] A. Miliias-Argeitis. Computational methods for simulation, identification and model selection in systems biology. *Doctoral dissertation, ETH Zurich*, 2013.
- [MARA⁺16] A. Miliias-Argeitis, M. Rullan, S. K. Aoki, P. Buchmann, and M. Khammash. Automated optogenetic feedback control for precise and robust regulation of gene expression and cell growth. *Nature Communications*, 7:12546, 2016.
- [MAS09] J. R. Marden, G. Arslan, and J. S. Shamma. Cooperative control and potential games. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 39(6):1393–1407, 2009.
- [MASSO⁺11] A. Miliias-Argeitis, S. Summers, J. Stewart-Ornstein, I. Zuleta, D. Pincus, H. El-Samad, M. Khammash, and J. Lygeros. In silico feedback for in vivo regulation of a gene expression circuit. *Nature Biotechnology*, 29:1114–1116, 2011.
- [Mat92] A. Matsui. Best response dynamics and socially stable strategies. *Journal of Economic Theory*, 57(2):343–362, 1992.
- [MB12] A. Mirtabatabaei and F. Bullo. Opinion dynamics in heterogeneous networks: Convergence conjectures and theorems. *SIAM Journal on Control and Optimization*, 50(5):2763–2785, 2012.
- [MCH13] Z. Ma, D. S. Callaway, and I. A. Hiskens. Decentralized charging control of large populations of plug-in electric vehicles. *IEEE Transactions on Control Systems Technology*, 21(1):67–78, 2013.
- [MDBDB11] F. Menolascina, M. Di Bernardo, and D. Di Bernardo. Analysis, design and implementation of a novel scheme for in-vivo control of synthetic gene regulatory networks. *Automatica*, 47(6):1265–1270, 2011.

- [MGP⁺16] L. Möller, B. Gentile, F. Parise, S. Grammatico, and J. Lygeros. Constrained deterministic leader-follower mean field control. In *Proceedings of the IEEE American Control Conference*, pages 4687–4692, 2016.
- [MHS14] K. Ma, G. Hu, and C. J. Spanos. Distributed energy consumption control via real-time pricing feedback in smart grid. *IEEE Transactions on Control Systems Technology*, 22(5):1907–1914, 2014.
- [MJB14] A. Mirtabatabaei, P. Jia, and F. Bullo. Eulerian opinion dynamics with bounded confidence and exogenous inputs. *SIAM Journal on Applied Dynamical Systems*, 13(1):425–446, 2014.
- [MK06] B. Munsky and M. Khammash. The finite state projection algorithm for the solution of the chemical master equation. *The Journal of chemical physics*, 124(4):044104, 2006.
- [MK07] B. Munsky and M. Khammash. A multiple time interval finite state projection algorithm for the solution to the chemical master equation. *Journal of Computational Physics*, 226(1):818–835, 2007.
- [Mos69] U. Mosco. Convergence of variational of convex sets and of solutions inequalities. *Advances in Mathematics*, 3(4):510–585, 1969.
- [MS96] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
- [MS10] A. Montanari and A. Saberi. The spread of innovations in social networks. *Proceedings of the National Academy of Sciences of the USA*, 107(47):20196–20201, 2010.
- [MS14] J. R. Marden and Shamma. Game theory and distributed control. In P. Young and S. Zamir, editors, *Handbook of game theory*, pages 861–899. 2014.
- [MT05] I. M. Mitchell and J. A. Templeton. A toolbox of Hamilton-Jacobi solvers for analysis of nondeterministic continuous and hybrid systems. In M. Morari and L. Thiele, editors, *Hybrid systems: Computation and control*, pages 480–494. Springer, 2005.
- [NCMH12] M. Nourian, P. E. Caines, R. P. Malhamé, and M. Huang. Mean field LQG control in leader-follower stochastic multi-agent systems: Likelihood ratio based adaptation. *IEEE Transactions on Automatic Control*, 57(11):2801–2816, 2012.

- [NPP16] C. Nowzari, V. M. Preciado, and G. J. Pappas. Analysis and control of epidemics: A survey of spreading processes on complex networks. *IEEE Control Systems*, 36(1):26–46, 2016.
- [OHL⁺14] E. J. Olson, L. A. Hartsough, B. P. Landry, R. Shroff, and J. J. Tabor. Characterizing bacterial gene circuit dynamics with optically programmed gene expression signals. *Nature Methods*, 11:449–455, 2014.
- [OSAFM07] R. Olfati-Saber, J. Alex Fax, and R. M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.
- [OT09] A. Olshevsky and J. N. Tsitsiklis. Convergence speed in distributed consensus and averaging. *SIAM Journal on Control and Optimization*, 48(1):33–55, 2009.
- [OT14] E. J. Olson and J. J. Tabor. Optogenetic characterization methods overcome key challenges in synthetic and systems biology. *Nature Chemical Biology*, 10(7):502–511, 2014.
- [Pav07] L. Pavel. An extension of duality to a game-theoretic framework. *Automatica*, 43(2):226–237, 2007.
- [PCGL14] F. Parise, M. Colombino, S. Grammatico, and J. Lygeros. Mean field constrained charging policy for large populations of plug-in electric vehicle. In *IEEE Conference on Decision and Control*, pages 5101–5106, 2014.
- [PGCL15] F. Parise, S. Grammatico, M. Colombino, and J. Lygeros. On constrained mean field control for large populations of heterogeneous agents: Decentralized convergence to Nash equilibria. In *Proceedings of the IEEE European Control Conference*, pages 3316–3321, 2015.
- [PGGL15a] F. Parise, B. Gentile, S. Grammatico, and J. Lygeros. Network aggregative games: Distributed convergence to Nash equilibria. In *Proceedings of the IEEE Conference on Decision and Control*, pages 2295–2300, 2015.
- [PGGL15b] F. Parise, S. Grammatico, B. Gentile, and J. Lygeros. Network aggregative games and distributed mean field control via consensus theory. *arXiv preprint*. Available online at <http://arxiv.org/pdf/1506.07719.pdf>, 2015.
- [PKL16] D. Paccagnan, M. Kamgarpour, and J. Lygeros. On aggregative and mean field games with applications to electricity markets. In *European Control Conference*, 2016.

- [PLR15] F. Parise, J. Lygeros, and J. Ruess. Bayesian inference for stochastic individual-based models of ecological systems: A pest control simulation study. *Frontiers in Environmental Science*, 3:42, 2015.
- [PP13] F. Parise and G. Picci. System identification for tide prediction in the Venetian lagoon. In *Proceedings of the IEEE European Control Conference*, pages 2994–2999, 2013.
- [PP14] F. Parise and G. Picci. Identification of high tide models in the Venetian lagoon: variable selection and G-LASSO. *IFAC Proceedings Volumes*, 47(3):10385–10390, 2014.
- [PPG⁺16] D. Paccagnan, F. Parise, B. Gentile, M. Kamgarpour, and J. Lygeros. Distributed computation of generalized Nash equilibria in aggregative games with coupling constraints. In *Proceedings of the IEEE Conference on Decision and Control*, to appear, 2016.
- [PPTF15] S. E. Parsegov, A. V. Proskurnikov, R. Tempo, and N. E. Friedkin. A novel multidimensional model of opinion dynamics in social networks. *arXiv preprint*. Available online at <http://arxiv.org/abs/1505.04920>, 2015.
- [PRL14] F. Parise, J. Ruess, and J. Lygeros. Grey-box techniques for the identification of a controlled gene expression model. In *Proceedings of the IEEE European Control Conference*, pages 1498–1503, 2014.
- [PVL14] F. Parise, M. E. Valcher, and J. Lygeros. On the reachable set of the controlled gene expression system. In *Proceedings of the IEEE Conference on Decision and Control*, pages 4597–4604, 2014.
- [PVL15] F. Parise, M. E. Valcher, and J. Lygeros. On the use of hyperplane methods to compute the reachable set of controlled stochastic biochemical reaction networks. In *Proceedings of the IEEE Conference on Decision and Control*, pages 1259–1264, 2015.
- [PVL16] F. Parise, M. E. Valcher, and J. Lygeros. Reachability analysis for switched affine autonomous systems and its application to controlled stochastic biochemical reaction networks. In *Proceedings of the IEEE Conference on Decision and Control*, to appear, 2016.
- [RL13] J. Ruess and J. Lygeros. Identifying stochastic biochemical networks from single-cell population experiments: A comparison of approaches based on the Fisher information. In *Proceedings of the IEEE Conference on Decision and Control*, pages 2703–2708, 2013.

- [RMAL13] J. Ruess, A. Miliias-Argeitis, and J. Lygeros. Designing experiments to understand the variability in biochemical reaction networks. *Journal of the Royal Society Interface*, 10(88):20130588, 2013.
- [RMASL11] J. Ruess, A. Miliias-Argeitis, S. Summers, and J. Lygeros. Moment estimation for chemically reacting systems by extended Kalman filtering. *The Journal of Chemical Physics*, 135:165102, 2011.
- [Roc70] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [Ros73] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [Rou05] T. Roughgarden. *Selfish routing and the price of anarchy*. MIT press Cambridge, 2005.
- [RPMA⁺15] J. Ruess*, F. Parise*, A. Miliias-Argeitis, M. Khammash, and J. Lygeros. Iterative experiment design guides the characterization of a light-inducible gene expression circuit. *Proceedings of the National Academy of Sciences of the USA*, 112(26):8148–8153, 2015.
- [Rue14] J. Ruess. Moment-based methods for the analysis and identification of stochastic models of biochemical reaction networks. *Doctoral dissertation, ETH Zurich*, 2014.
- [San10] W. H. Sandholm. *Population games and evolutionary dynamics*. The MIT press, 2010.
- [SCGB06] C. Seatzu, D. Corona, A. Giua, and A. Bemporad. Optimal control of continuous-time switched affine systems. *IEEE Transactions on Automatic Control*, 51(5):726–741, 2006.
- [Sin92] A. Sinclair. Improved bounds for mixing rates of Markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1(04):351–370, 1992.
- [Sma74] D. R. Smart. *Fixed point theorems*. Cambridge University Press Archive, 1974.
- [Smi82] J. M. Smith. *Evolution and the Theory of Games*. Cambridge university press, 1982.
- [Son98] E. D. Sontag. *Mathematical control theory: Deterministic finite dimensional systems*. Springer, 1998.

- [SPFP10] G. Scutari, D. P. Palomar, F. Facchinei, and J.-S. Pang. Convex optimization, game theory, and variational inequality theory. *IEEE Signal Processing Magazine*, 27(3):35–49, 2010.
- [SS08] V. Shahrezaei and P. S. Swain. The stochastic nature of biochemical networks. *Current opinion in biotechnology*, 19(4):369–374, 2008.
- [SSHTQ02] S. Shimizu-Sato, E. Huq, J. M. Tepperman, and P. H. Quail. A light-switchable gene promoter system. *Nature Biotechnology*, 20(10):1041–1044, 2002.
- [Sus83] H. J. Sussmann. Lie Brackets, Real Analyticity and Geometric Control. In R. W. Brockett, R. S. Millman, and H. J. Sussmann, editors, *Differential Geometric Control Theory*, volume 27, pages 1–116. Birkhauser Boston, 1983.
- [TCW04] G. Taguchi, S. Chowdhury, and Y. Wu. *Taguchi’s quality engineering handbook*. John Wiley & Sons Hoboken, 2004.
- [TX14] J. N. Tsitsiklis and Y. Xu. Efficiency loss in a Cournot oligopoly with convex market demand. *Journal of Mathematical Economics*, 53:46–58, 2014.
- [UMD⁺12] J. Uhlenhof, A. Miermont, T. Delaveau, G. Charvin, F. Fages, S. Bot-tani, G. Batt, and P. Hersen. Long-term model predictive control of gene expression at the population and single-cell levels. *Proceedings of the National Academy of Sciences of the USA*, 109(35):14271–14276, 2012.
- [VK07] N. G. Van Kampen. *Stochastic processes in physics and chemistry, third edition*. Elsevier, 2007.
- [Wac13] G. Wachsmuth. On LICQ and the uniqueness of Lagrange multipliers. *Operations Research Letters*, 41(1):78–80, 2013.
- [WB05] C. M. Waters and B. L. Bassler. Quorum sensing: Cell-to-cell communication in bacteria. *Annual Review of Cell and Developmental Biology*, 21:319–346, 2005.
- [Wil11] D. J. Wilkinson. *Stochastic modeling for systems biology*. CRC press, 2011.
- [YNS16a] F. Yousefian, A. Nedić, and U. V. Shanbhag. On smoothing, regularization and averaging in stochastic approximation methods for stochastic variational inequalities. *arXiv preprint*. Available online at <http://arxiv.org/pdf/1411.0209v2.pdf>, 2016.

- [YNS16b] F. Yousefian, A. Nedić, and U. V. Shanbhag. Self-tuned stochastic approximation schemes for non-Lipschitzian stochastic multi-user optimization and Nash games. *IEEE Transactions on Automatic Control*, 61(7):1753–1766, 2016.
- [ZC10] M. Zhong and C. G. Cassandras. Asynchronous distributed optimization with event-driven communication. *IEEE Transactions on Automatic Control*, 55(12):2735–2750, 2010.
- [ZF16] M. Zhua and E. Frazzoli. Distributed robust adaptive equilibrium computation for generalized convex games. *Automatica*, 63:82–91, 2016.
- [ZNUK12] C. Zechner, P. Nandy, M. Unger, and H. Koepl. Optimal variational perturbations for the inference of stochastic reaction dynamics. In *Proceedings of the IEEE Conference on Decision and Control*, pages 5336–5341, 2012.
- [ZRK⁺12] C. Zechner, J. Ruess, P. Krenn, S. Pelet, M. Peter, J. Lygeros, and H. Koepl. Moment-based inference predicts bimodality in transient gene expression. *Proceedings of the National Academy of Sciences of the USA*, 109(21):8340–8345, 2012.
- [ZSRK16] C. Zechner, G. Seelig, M. Rullan, and M. Khammash. Molecular circuits for dynamic noise filtering. *Proceedings of the National Academy of Sciences of the USA*, 113(17):4729–4734, 2016.

Curriculum vitae

Francesca Parise

born October 10th, 1988 in Italy.

- 2012-2016 PhD, in Control Engineering
ETH Zürich, Automatic Control Laboratory, Switzerland
- 2007-2012 Galilean School of Higher Education (100/100 cum Laude)
University of Padova, Italy
- 2010-2012 Master's Degree in Control Engineering (110/110 cum Laude)
University of Padova, Italy
- 2007-2010 Bachelor's Degree in Information Engineering (110/110 cum Laude)
University of Padova, Italy
- 2002-2007 Scientific high school diploma (100/100)
Liceo Scientifico G.Galilei, Dolo, Venezia, Italy