Advances in predictive control and feedback equilibrium-seeking with applications to autonomous water resource recovery

Otacilio Bezerra Leite Neto

Aalto University

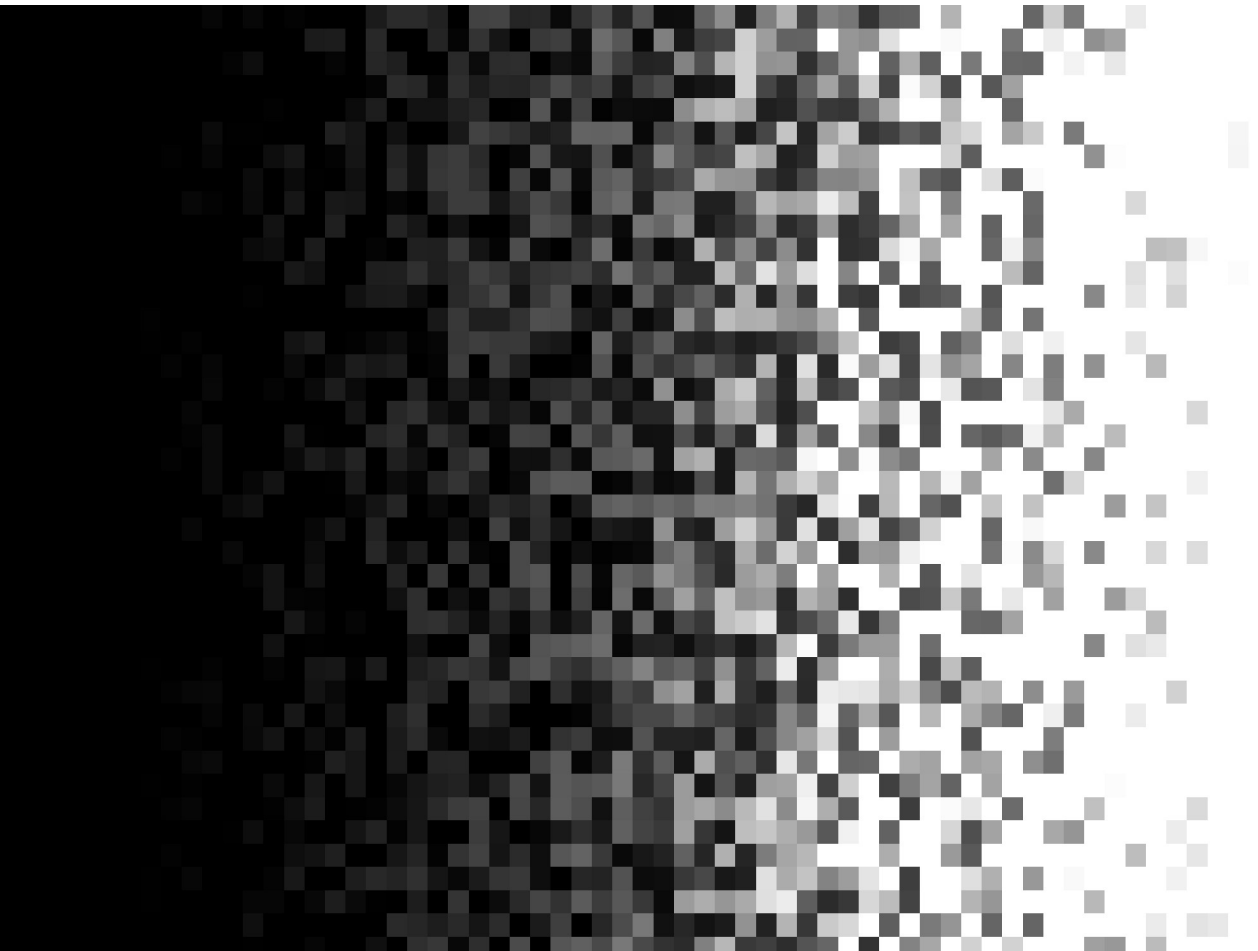

Aalto University publication series
Doctoral Theses 233/2025

Advances in predictive control and feedback equilibrium-seeking with applications to autonomous water resource recovery

Otacilio Bezerra Leite Neto

A doctoral thesis completed for the degree of Doctor of Science (Technology) to be defended, with the permission of the Aalto University School of Chemical Engineering, at a public examination held at the lecture hall Ke2 of the school on 05 December 2025 at 12:00.

Aalto University
School of Chemical Engineering
Department of Chemical and Metallurgical Engineering

Supervising professor
Professor Iiro Harjunkoski, Aalto University, Finland

Thesis advisors
Professor Michela Mulas, Federal University of Ceará, Brazil

Preliminary examiners
Professor Lorenz T. Biegler, Carnegie Mellon University, USA
Professor Carlos Ocampo-Martinez, Polytechnic University of Catalonia, Spain

Opponent
Professor Lorenz T. Biegler, Carnegie Mellon University, USA

Aalto University publication series
Doctoral Theses 233/2025

Image on the cover by Otacilio Bezerra Leite Neto

ISBN 978-952-64-2841-3 (paperback)
ISBN 978-952-64-2840-6 (PDF)
ISSN 1799-4934 (printed)
ISSN 1799-4942 (PDF)

https://urn.fi/URN:ISBN:978-952-64-2840-6

PunaMusta Oy
Helsinki 2025

Aalto University

ABSTRACT

Author
Otacilio Bezerra Leite Neto

Name of the doctoral thesis
Advances in predictive control and feedback equilibrium-seeking with applications to autonomous water resource recovery

Article-based thesis

Number of pages 138

Keywords
water resource recovery, feedback control, model predictive control, moving horizon estimation, multi-agent systems, Nash equilibrium, system level synthesis, fixed point iteration

A paradigm shift is underway: Wastewater, historically viewed as an environmental hazard, is now recognized as a sustainable source of clean water, energy, and nutrients. Under this context, wastewater treatment plants (WWTPs) are being rethought as water resource recovery facilities (WRRFs), plants that produce goods and energy using wastewater as a raw material. While current efforts focus on designing novel processes, this dissertation considers the role of automatic decision-making (that is, feedback control) in this transition instead. We propose two research directions: (i) The repurposing of already existing infrastructure to these emerging objectives, and (ii) The promotion of a wastewater-centered market through an interconnected network of WRRFs.

In the first direction, the thesis studies the predictive control of conventional biological WWTPs for water resource recovery tasks. The proposal consists of an output-feedback control framework comprised of an operating point optimizer, a model predictive controller, and a moving horizon estimator. The framework is designed to control a WWTP to recover nutrients directly into effluent streams, thus producing reused water of tailored quality, while ensuring that such an operation regime can be fully sustained by the energy recovered from processing sludge into biogas. Using well-established benchmark models, experimental results demonstrate the efficacy of this control strategy in operating a conventional WWTP as an energy-autonomous WRRF.

In the second direction, the thesis studies the control of (noncooperative) multi-agent systems. The underlying idea is to connect WRRFs into a supply chain exchanging wastewater and recovered resources. Here, the contribution is theoretical: it addresses the open problem of seeking game-theoretical equilibria (e.g., the Nash equilibrium) of output-feedback policies that incorporate both operational and informational constraints. The thesis proposes novel equilibrium-seeking algorithms in which noncooperative agents are able to converge to an equilibrium of feedback policies, while simultaneously operating their subsystem under the aforementioned constraints. We provide formal convergence certificates and demonstrate the algorithms in exemplary problems.

Our theoretical and experimental findings contribute a step towards the transition into zero-waste water resource recovery infrastructures. In turn, the contributions to multi-agent control should benefit applications in different domains, such as smart grids and supply chain management.

Preface

This doctoral thesis tells a story of unexpected changes. On its surface, this is a story of wastewater changing from an environmental hazard to a potential commodity. However, it is also the story of a guy moving from tropical Brazil to snowy Finland, from a computer engineering department to a chemical engineering department, to work on a thesis which changed from optimization-based control design to equilibrium-based control design. Changes can be uncomfortable, but they are often necessary—and I made an effort to welcome all of them. Since nothing can change on its own, here I express my gratitude to those who helped me in this journey.

First and foremost, I would like to express my deepest gratitude to my supervisor, Prof. Francesco Corona, who guided and supported me throughout this journey. I also extend my gratitude to Prof. Iiro Harjunkoski, who supervised the latter stages of this thesis, and to my thesis advisor, Prof. Michela Mulas. The many invaluable lessons that I learned from you helped me grow not only as an academic but also as a person.

During my time at Aalto University, I was fortunate to have very talented and caring people by my side. My sincere thanks to José Augusto Fontenele Magalhães (*aka*, Aukusti) and Jukka Kortela, my office mates, for their company and interesting discussions. A warm thank you to all my colleagues at work, especially to my friends Andreu Bernard Serra, Anna Zaykovskaya, Nicolaus Bijok, and Oliwer Śliczniuk. I will forever cherish our moments together, drinking coffee by the imaginary window at kitchen E301 while the koira watches over us. I also want to acknowledge the students whom I had the pleasure to teach in the “E7225 Advanced Process Control D” course. I hope that you learned from me as much as I learned from you.

This thesis would not have been possible without the love and support from my family. To my parents, Tereza Diana and Otacílio Filho, for always encouraging me to study and to chase my goals. To my siblings, Camilla Menezes and Matheus Menezes, for their companionship and for taking care of everything while I’m gone. Finally, most importantly, to my loving wife, Mariana Sátiro. Thank you for being here, for taking care of me, and for helping me walk through this arduous path. Te amo, lagartinha!

I acknowledge with gratitude the funding provided to me by the School of Chemical Engineering, Aalto University, which has financially supported this research project. I extend my gratitude to my advisory group, especially to Prof. Fabricio Oliveira and Sirje Liukko, for their feedback and support.

During these past years, I have spent the majority of my working hours wearing headphones/earbuds (as many would point out). Thus, I would also like to thank all the artists and bands who provided the soundtrack for this research. Finally, I am thankful to all of those not mentioned here who have directly or indirectly contributed to this thesis becoming a reality.

Helsinki, October 16, 2025,

Otacilio Bezerra Leite Neto

Contents

List of Publications

This thesis consists of an overview and of the following publications which are referred to in the text by their Roman numerals.

I Otacilio B. L. Neto, Michela Mulas, Francesco Corona. A model-based framework for controlling activated sludge plants. *Chemical Engineering Journal*, Volume 488, 150750, May 2024.

II Otacilio B. L. Neto, Michela Mulas, Iiro Harjunkoski, Francesco Corona. Predictive control of wastewater treatment plants as energy-autonomous water resource recovery facilities. Submitted to *Control Engineering Practice*, 2025.

III Otacilio B. L. Neto, Michela Mulas, Francesco Corona. SLS-BRD: A system-level approach to seeking generalized feedback Nash equilibria. Accepted for publication in *IEEE Transactions on Automatic Control*, 2025.

IV Otacilio B. L. Neto, Michela Mulas, Francesco Corona. A system level approach to generalized feedback Nash equilibrium seeking in partially-observed games. Submitted to *IEEE Transactions on Cybernetics*, 2025.

Author's contributions

Publication I: "A model-based framework for controlling activated sludge plants"

OBLN developed the proposed method, implemented the simulation model and supporting algorithms, analyzed the results, and wrote the initial draft of the manuscript. FC supervised the research. MM and FC critically reviewed the conclusions and supervised the writing.

Publication II: "Predictive control of wastewater treatment plants as energy-autonomous water resource recovery facilities"

OBLN developed the proposed method, implemented the simulation model and supporting algorithms, analyzed the results, and wrote the initial draft of the manuscript. FC supervised the research. MM, IH, and FC critically reviewed the conclusions and supervised the writing.

Publication III: "SLS-BRD: A system-level approach to seeking generalized feedback Nash equilibria"

OBLN developed the proposed method, formally analyzed its properties, implemented the experiments, and wrote the initial draft of the manuscript. FC supervised the research. MM and FC critically reviewed the conclusions and supervised the writing.

Publication IV: "A system level approach to generalized feedback Nash equilibrium seeking in partially-observed games"

OBLN developed the proposed method, formally analyzed its properties, implemented the experiments, and wrote the initial draft of the manuscript. FC supervised the research. MM and FC critically reviewed the conclusions and supervised the writing.

List of Abbreviations

ADMM	Alternating direction method of multipliers
ARE	Algebraic Riccati equation
BOD	Biochemical oxygen demand
BRD	Best-response dynamics
BSM	Benchmark simulation model
ECI	Energy cost index
FBS	Forward-backward splitting
FIR	Finite-impulse response
KPI	Key performance indicator
LQG	Linear quadratic Gaussian
MHE	Moving horizon estimation
MPC	Model predictive control
OPO	Operating point optimizer
SLS	System level synthesis
TN	Total nitrogen
TSS	Total suspended solids
WRRF	Water resource recovery facility
WWTP	Wastewater treatment plant
(v/G/F)NE	(variational/Generalized/Feedback) Nash equilibrium

List of Symbols

$\mathbb{R}$	field of real numbers
$\mathbb{N}$	set of all natural numbers, including zero
$x[n]$	nth sample of a discrete-time signal $x : \mathbb{N} \to \mathbb{R}^{N_x}$
$x(t)$	value at t of a continuous-time signal $x : \mathbb{R} \to \mathbb{R}^{N_x}$
$\hat{x}$	$\mathbb{Z}$-transform of a discrete-time signal (or mapping) x
$A * x[n]$	nth sample of the the convolution $A * x$
$\delta[n]$	nth sample of the unit impulse signal
$(x^1, \ldots, x^N)$	an ordered N-tuple
(x^p, x^{-p})	tuple with the pth element highlighted (not a reorder)
x^{-p}	all except the pth element of a tuple
$\frac{d}{dt}$	differentiation operator w.r.t time
E	expectation operator
Sp	sparsity, or support, operator (of a matrix)
$N_{\mathcal{S}}$	normal cone operator defined for a set $\mathcal{S}$
$\mathrm{proj}_{\mathcal{S}}$	projection operator defined for a set $\mathcal{S}$
$\mathbf{fix}(F)$	set of fixed-points of an operator F
$\mathbf{zer}(F)$	set of zeroes of an operator F
$[A]_{i,j}$	ith row and jth column element of a matrix A
I_N	$(N \times N)$ identity matrix or identity operator
e_i	ith column of an identity matrix
$\sigma_{\max}(A),\ \sigma_{\min}(A)$	largest and smallest singular value of matrix A

$\rho(A)$	spectral radius of matrix A
$\texttt{blkdiag}(\cdots)$	block-diagonal matrix with $(\cdots)$ in its main diagonal
$\ell_2^N(\mathbb{N})$	space of all square-summable N-dimensional discrete-time signals
$\ell_\infty^N(\mathbb{N})$	space of all bounded N-dimensional discrete-time signals
$\mathcal{L}(\ell_\infty^N, \ell_\infty^M)$	set of all bounded linear mappings from $\ell_\infty^N(\mathbb{N})$ to $\ell_\infty^M(\mathbb{N})$
$\frac{1}{z}\mathcal{RH}_\infty$	set of all real-rational strictly proper transfer functions
$\|\cdot\|_2$	Euclidean (vector) norm
$\|\cdot\|_F$	Frobenius (matrix) norm
$\|\cdot\|_{\ell_2}$	ℓ_2 norm, that is, the square-root of the square-sum of a sequence
$\|\cdot\|_{\ell_\infty}$	ℓ_∞ norm, that is, largest absolute value of a sequence

1. Introduction

A paradigm shift is underway: Wastewater, historically viewed as an environmental hazard, is being recognized as a sustainable source of clean water, energy, and nutrients [1–4]. The potential to commodify these resources, together with increasing water stresses and environmental concerns, is driving the wastewater sector to reorient wastewater treatment plants (WWTPs) to function as water resource recovery facilities (WRRFs), that is, plants that produce goods using wastewater as a raw material [5, 6]. In particular, the production of reused (or reclaimed) water of different qualities has attracted much attention. Water that is rich in nitrogen and phosphorus, nutrients which are abundant in domestic wastewater, is a valuable asset in agriculture [7]. Partially treated, non-potable water can support the demand from a variety of urban and industrial activities, such as toilet flushing, cleaning, and water cooling [8]. A treatment process capable of tailoring the quality of reused water can promote both economic development (the provision of a valuable commodity) and environmental sustainability (the conservation of freshwater resources), thus making it an essential tool in achieving circular economy and zero-waste goals. Finally, excess sludge from such a process can be potentially harvested for heat and electrical energy; thus ensuring an equalized water-energy nexus.

A technology that produces reused water of tailored quality with net-zero (or even net-negative) energy balance costs has been a central target of the research in transitioning from pollution abatement to resource recovery. While current research efforts focus on revamping the existing infrastructure with novel processes [9–15]—an undeniably costly endeavor [16]—, little has been investigated on how automatic decision-making (that is, feedback control [17–19]) can play a role in this transition. Under new operation policies, can the treatment systems already in place be rendered flexible enough to resource recovery tasks? How can these systems be repurposed and integrated to modern water reuse chains? This thesis aims to address these questions by considering two research directions in which process control could enable a circular, autonomous, water resource recovery.

Centralized water reuse chains

In a centralized water reuse chain, a WWTP is operated to produce reused water, energy, and raw materials, according to the resource demands from urban, industrial, and agricultural ecosystems (Figure 1.1). For this scenario, biological treatment of wastewater [20] offers the most promising platform: This infrastructure (i) is common to most urban centers, (ii) has a recognized potential for energy recovery through the production of biogas, and (iii) is centred on controlling the nitrogen and phosphorus content within wastewater—two major nutrients lost by conventional treatment. Moreover, biological WWTPs are increasingly instrumented and digitalized, thus facilitating the design and deployment of new operation policies [21].

Biological WWTPs are large-scale dynamical systems of notable complexity. Their real-time operation requires feedback control solutions capable of driving the plant using data from *in situ* sensors. Traditionally, this has been achieved by low-level regulators that control key process variables to match operational targets set during process design [22]. More recently, advanced methods such as model predictive control (MPC, [19]) have found noteworthy success in optimizing the operation of biological WWTPs: their performance was shown to improve considerably when using either linear [23–25] or nonlinear [26–29] MPC strategies, including when the process dynamics is represented by black-box models [30, 31]. While demonstrating the capabilities of model-based feedback control, these contributions have been limited to the classical goal of nitrogen removal; that is, they are not designed for resource recovery objectives. A unified control framework which is valid at a system-level and general enough to incorporate resource recovery operations to conventional WWTPs remains lacking.

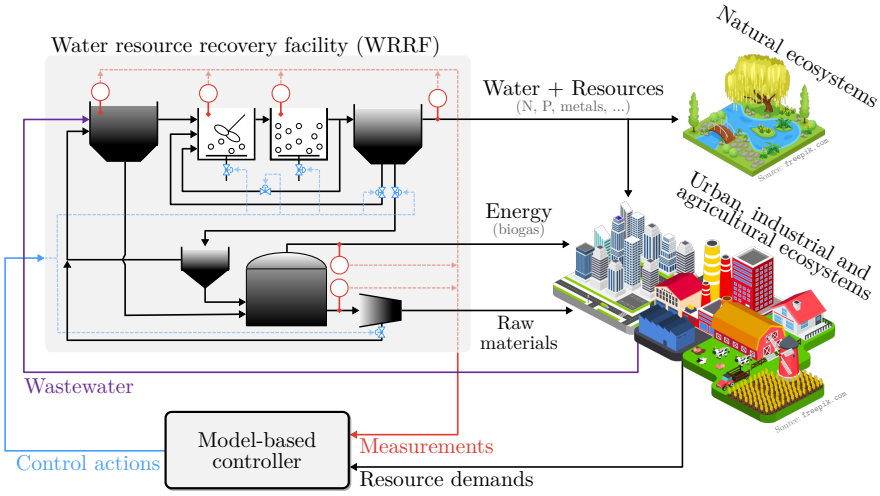

Figure 1.1. Diagram of an autonomous WRRF in a centralized water reuse chain. To satisfy the external resource demands, a (model-based) controller uses measurement data to determine the optimal actions for the plant's actuators.

Networked water reuse chains

In a networked water reuse chain, an interconnected grid of WRRFs is operated simultaneously while exchanging water and resources within themselves and their urban, industrial, and agricultural clients (Figure 1.2). The WRRFs are each controlled by a self-interested decision-making agent. Under this setup, a new market emerges with wastewater as the central commodity: the agents decide whether to consume or produce wastewater and its products based on the current demands. In addition to allowing more flexible and sophisticated operation strategies, such an architecture alleviates the shortcomings of centralized treatment systems, such as large maintenance costs and blending of heterogeneous sources [32, 33].

Being noncooperative multi-agent systems, this class of WRRF networks cannot be controlled using classical centralized methods. Dynamic game theory provides an alternative framework through the concept of noncooperative equilibria (e.g., the Nash equilibrium) [34, 35]. Under a dynamic game setting, agents design control policies to operate their subsystems while aware that their choice affects (and is affected by) the others' choices. Game-theoretical approaches to online decision-making have been proven successful in many applications, including communication networks [36, 37] and power systems [38–40]. The networked water reuse chain that we conceptualize would require equilibrium-seeking solutions to the synthesis of output-feedback policies which (i) are jointly stabilizing, (ii) incorporate operational constraints (e.g., actuator limits and restrictions on the effluent quality), and (iii) comply with the network's informational structure. Unfortunately, the existing methods cannot address these requirements and are mostly limited to state-feedback problems [41–50]. A systematic approach to equilibrium seeking for this class of dynamic games remains lacking.

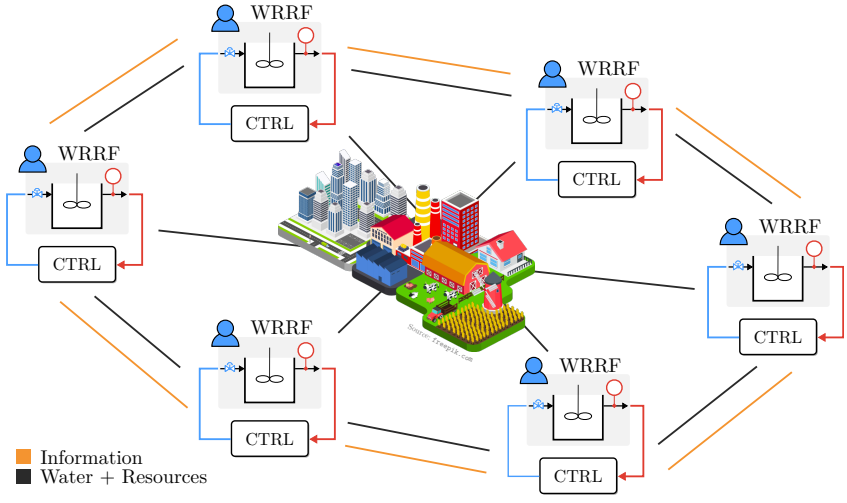

Figure 1.2. Diagram of autonomous WRRFs in a networked water reuse chain. The WRRFs are each managed by a decision-making agent, which deploys a feedback controller for their automatic operation based on self-interested objectives.

1.1 Summary of contributions

This thesis comprises four publications which address the literature gaps in the two previously discussed research directions. In the first direction, we propose a flexible control framework for the operation of conventional WWTPs as self-sufficient WRRFs (Publications I and II). In the second direction, our contribution is methodological: it addresses the open problem of seeking game-theoretical equilibria of control policies that incorporate both operational and informational constraints (Publications III and IV). The main contributions of each publication are as follows:

Publication I: We propose a control framework for rendering biological WWTPs flexible to both conventional and resource recovery objectives. The framework is centered on an MPC strategy for reference tracking, with the references assumed to reflect the environmental and economical targets set by the plant's management.

Publication II: We formalize the task of energy-autonomous production of reused water and design an MPC controller for operating a full-scale WWTP towards this objective. The results show that our control framework is capable of operating a WWTP as a self-sufficient WRRF, thus contributing a proof-of-concept on the feasibility of repurposing existing wastewater treatment infrastructures into centralized water reuse chains.

Publication III: An algorithm for seeking generalized feedback Nash equilibria (GFNE) is proposed. Leveraging a novel approach to control design, we construct a GFNE-seeking algorithm for solving dynamic games in which the players' policies are subjected to both operational and informational constraints. The algorithm is showcased in exemplary problems including price competition in oligopolistic markets.

Publication IV: An algorithm for GFNE-seeking in partially-observed games is proposed. The algorithm extends our previous work by considering the wider class of output-feedback problems, by taking a statistical approach to robust control, and by applying monotone operator theory to design a routine with proven convergence. The algorithm is showcased on the noncooperative control of a decentralized power-grid.

1.2 Outline of the thesis

This thesis consists of five chapters followed by Publications I–IV. The theoretical background for the proposed methods is covered in Chapter 2. A summary of our proposed control framework and its application to centralized water reuse chains is presented in Chapter 3. A summary of our proposed GFNE-seeking algorithms and their application to noncooperative dynamic games is presented in Chapter 4. Finally, Chapter 5 provides some concluding remarks and directions for future research.

2. Theoretical background

This chapter introduces main concepts from feedback control and game theory, which are fundamental for the contributions of this thesis. We first discuss model-based control problems, with focus on the model predictive control and system level synthesis approaches to designing optimal feedback policies. The chapter concludes with an overview of noncooperative dynamic games and their solution through equilibrium-seeking methods.

2.1 Model-based output-feedback control

A model-based approach to controller design is centred on a mathematical representation of the dynamics and observation process of the system to be controlled. In this thesis, the focus is on discrete-time state-space models

$$\Sigma : \begin{cases} x[n+1] = Ax[n] + B_w w[n] + B_u u[n]; & \text{(2.1a)} \\ y[n] = Cx[n] + D_w w[n], & \text{(2.1b)} \end{cases}$$

where $x[n] \in \mathbb{R}^{N_x}$ is the state of the system, $y[n] \in \mathbb{R}^{N_y}$ is its output (or measurement), $w[n] \in \mathbb{R}^{N_w}$ are external disturbances (to the state and the output), and $u[n] \in \mathbb{R}^{N_u}$ are the actionable inputs (or control actions), at time step $n \in \mathbb{N}$. A linear time-invariant system such as Σ often arises from the linearization of a nonlinear system around some desired steady-state (with a time-discretization performed if the model is in continuous-time).

A controller (Figure 2.1) is a digital device which determines the control signal u according to the (possibly nonlinear) output-feedback control policy

$$K : \ell_\infty^{N_y}(\mathbb{N}) \to \ell_\infty^{N_u}(\mathbb{N}), \quad u = Ky. \tag{2.2}$$

The control policy K is often designed to optimally achieve some technical objective (e.g., minimizing the distance from the state signal x to a given reference signal x^{ref}). For unconstrained systems, optimal policies can be obtained using the well-established $\mathcal{H}_2$/LQG and $\mathcal{H}_\infty$ control methods [17]. In general, however, analytical formulas for K are not available. A common

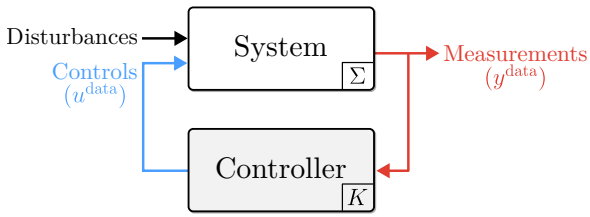

Figure 2.1. Diagram of an output-feedback controller. The superscript $(\cdot)^{\text{data}}$ distinguishes the signals actually exchanged with the system from those modelled by Σ.

approach is to describe the policy using computational routines: these include the popular *model predictive control* (MPC, [19]) methods. Recently, *system level synthesis* (SLS, [51]) emerged as a promising framework for the design of optimal policies which also incorporate structural constraints (typically, communication and actuation delays). Being central to this thesis, we overview the MPC and SLS methods in the following.

Model predictive control

In a model predictive controller (Figure 2.2), the control signal is defined by the causal recursive expression

$$u^{\text{data}}[n] := e_0\big(\text{MPC} \circ \text{MHE}\big)\big(u^{\text{data}}[n-N_e], \dots, u^{\text{data}}[n-1], \\ y^{\text{data}}[n-N_e], \dots, y^{\text{data}}[n]\big) \quad (2.3)$$

where e_0 is a selector operator[1], the mapping

$$\text{MPC} := (\hat{w}_0, \dots, \hat{w}_{N_e-1},\ \hat{x}_0, \dots, \hat{x}_{N_e}) \mapsto (u_0^\star, \dots, u_{N_c-1}^\star,\ x_0^\star, \dots, x_{N_c}^\star)$$

returns a solution to finite-horizon optimal control problems of the form

$$\underset{u[\cdot],\ x[\cdot]}{\text{minimize}} \quad \sum_{n=0}^{N_c-1} L_{c|n}(x[n], u[n]) + L_{c|N_c}(x[N_c]) \quad (2.4a)$$

$$\underset{n=0,\dots,N_c-1}{\text{subject to}} \quad \Sigma \text{ (Eq. 2.1) with } w[n] = \hat{w}_{N_e-1} \text{ and } x[0] = \hat{x}_{N_e}, \quad (2.4b)$$

$$x[n] \in \mathcal{X}_{c|n}, \quad u[n] \in \mathcal{U}_{c|n}, \quad (2.4c)$$

and the mapping

$$\text{MHE} := (u_0^{\text{data}}, \dots, u_{N_e-1}^{\text{data}},\ y_0^{\text{data}}, \dots, y_{N_e}^{\text{data}}) \mapsto (\hat{w}_0, \dots, \hat{w}_{N_e-1},\ \hat{x}_0, \dots, \hat{x}_{N_e})$$

returns a solution to finite-horizon optimal estimation problems of the form

$$\underset{w[\cdot],\ x[\cdot]}{\text{minimize}} \quad \sum_{n=0}^{N_e-1} L_{e|n}(x[n], w[n] \mid y_n^{\text{data}}) + L_{e|N_e}(x[N_e] \mid y_{N_e}^{\text{data}}) \quad (2.5a)$$

$$\underset{n=0,\dots,N_e-1}{\text{subject to}} \quad \Sigma \text{ (Eq. 2.1) with } u[n] = u_n^{\text{data}}, \quad (2.5b)$$

$$x[n] \in \mathcal{X}_{e|n}, \quad w[n] \in \mathcal{W}_{e|n}. \quad (2.5c)$$

[1] For an ordered $(N+1)$-tuple $x = (x_0, x_1, \dots, x_N)$, we have that $e_n x = x_n$.

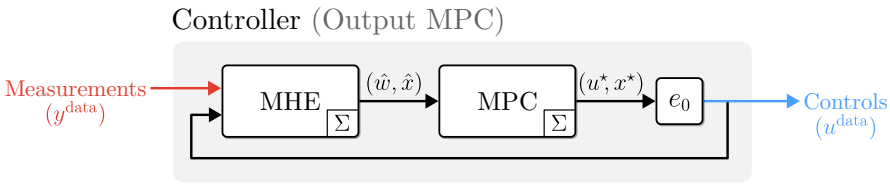

Figure 2.2. Diagram of an output-feedback model predictive control policy.

The Problems (2.4) and (2.5) are respectively known as the *model predictive control* (MPC) and *moving horizon estimation* (MHE) problems [19]. Explicitly, at each time step $n \in \mathbb{N}$, the policy in Eq. (2.3) consists of:

(i) collecting the past N_e control actions and N_e+1 measurements;

(ii) given this data, solving the MHE Problem (2.5) to obtain N_e disturbance estimates and N_e+1 state estimates;

(iii) given these estimates, solving the MPC Problem (2.4) to obtain N_c optimal controls and N_c+1 predicted state-responses;

(iv) deploying only the first optimal control action to the system.

Since it repeatedly solves dynamical optimization problems using data which shifts in time, this algorithm is known as a *receding-horizon control* or *output-feedback model predictive control* (Output MPC) policy.

The design of an Output MPC policy refers to the engineering task of choosing the horizon lengths $\{N_c, N_e\}$, cost functions $\{L_{c|(\cdot)}, L_{e|(\cdot)}\}$, and constraint sets $\{\mathcal{X}_{c|(\cdot)}, \mathcal{U}_{c|(\cdot)}, \mathcal{X}_{e|(\cdot)}, \mathcal{W}_{e|(\cdot)}\}$, such that the controller, when operating the physical system, achieves the desired objective while satisfying operational requirements. In special, the boundary costs $\{L_{c|N_c}, L_{e|0}\}$ and constraints $\{\mathcal{X}_{c|(N_c)}, \mathcal{U}_{c|N_c}, \mathcal{X}_{e|0}, \mathcal{W}_{e|0}\}$ require careful design to guarantee closed-loop stability [19]. The design should also ensure that Problems (2.4) and (2.5) can be solved with enough efficiency and reliability for the Output MPC policy to be implemented in practice. A common design choice (whenever possible) is to consider convex objective functions and constraint sets, allowing for these problems to be solved using interior-point methods [52–54] and other structure-exploiting algorithms [55–57]. Finally, we remark that the Output MPC architecture can be supplemented with auxiliary components, such as online model identifiers and disturbance predictors.

System level synthesis

Under a linear policy $u = Ky$, the $\mathbb{Z}$-transform of the system Σ (Eq. 2.1) allows us to express $(\hat{x}, \hat{u})$ in terms of the external disturbances,

$$\begin{bmatrix} \hat{x} \\ \hat{u} \end{bmatrix} = \begin{bmatrix} \hat{\Phi}_{xx} & \hat{\Phi}_{xx} B_u \hat{K} \\ \hat{K} C \hat{\Phi}_{xx} & \hat{K} + \hat{K} C \hat{\Phi}_{xx} B_u \hat{K} \end{bmatrix} \begin{bmatrix} \hat{\delta}_x \\ \hat{\delta}_y \end{bmatrix} := \begin{bmatrix} \hat{\Phi}_{xx} & \hat{\Phi}_{xy} \\ \hat{\Phi}_{ux} & \hat{\Phi}_{uy} \end{bmatrix} \begin{bmatrix} \hat{\delta}_x \\ \hat{\delta}_y \end{bmatrix}, \tag{2.6}$$

given matrix $\hat{\Phi}_{xx} = \left(zI - A - B_u \hat{K} C\right)^{-1}$, and signals $\hat{\delta}_x = B_w \hat{w}$ and $\hat{\delta}_y = D_w \hat{w}$. The transfer matrices $(\hat{\Phi}_{xx}, \hat{\Phi}_{xy}, \hat{\Phi}_{ux}, \hat{\Phi}_{uy})$ describe how the system reacts, in

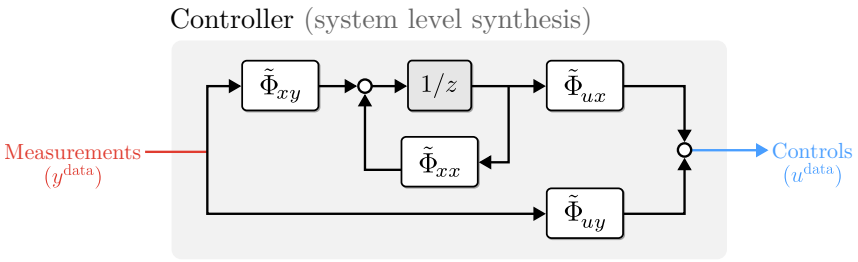

Figure 2.3. Diagram of an output-feedback controller under a system level parametrization. The policy $\hat{K} = \hat{\Phi}_{uy} + \hat{\Phi}_{ux}\hat{\Phi}_{xx}^{-1}\hat{\Phi}_{xy}$ is represented using the auxiliary transfer matrices $\tilde{\Phi}_x = z(I - z\hat{\Phi}_{xx})$, $\tilde{\Phi}_{ux} = z\hat{\Phi}_{ux}$, $\tilde{\Phi}_{xy} = -z\hat{\Phi}_{xy}$ and $\tilde{\Phi}_{uy} = \hat{\Phi}_{uy}$.

closed-loop, to the disturbances: They are denoted *system level responses*. Using this representation, the following results hold [51, Theorem 5.1]:

(a) The policy $\hat{K}$ is internally stabilizing if $(\hat{\Phi}_{xx}, \hat{\Phi}_{xy}, \hat{\Phi}_{ux}, \hat{\Phi}_{uy})$ satisfies

$$\hat{Z}_{AB}\begin{bmatrix}\hat{\Phi}_{xx} & \hat{\Phi}_{xy}\\ \hat{\Phi}_{ux} & \hat{\Phi}_{uy}\end{bmatrix} = \hat{Z}_{I0} \quad \text{and} \quad \begin{bmatrix}\hat{\Phi}_{xx} & \hat{\Phi}_{xy}\\ \hat{\Phi}_{ux} & \hat{\Phi}_{uy}\end{bmatrix}\hat{Z}_{AC} = \hat{Z}_{I0}^{\mathsf{T}} \tag{2.7}$$

with $\hat{Z}_{AB} = [zI - A \;\; B_u]$, $\hat{Z}_{AC} = [zI - A^{\mathsf{T}} \;\; C^{\mathsf{T}}]^{\mathsf{T}}$, and $\hat{Z}_{I0} = [I \;\; 0]$.

(b) A response $(\hat{\Phi}_{xx}, \hat{\Phi}_{xy}, \hat{\Phi}_{ux}, \hat{\Phi}_{uy})$ satisfying Eq. (2.7) is achieved by the stabilizing policy $\hat{K} = \hat{\Phi}_{uy} - \hat{\Phi}_{ux}\hat{\Phi}_{xx}^{-1}\hat{\Phi}_{xy}$ (Figure 2.3).

The system level approach to control synthesis consists of determining a desired system level response and then reconstructing the policy which achieves it, as opposed to searching for an optimal control policy K directly. The class of system level synthesis problems of the form

$$\underset{\hat{\Phi}_{xx},\hat{\Phi}_{xy},\hat{\Phi}_{ux},\hat{\Phi}_{uy}}{\text{minimize}} \quad J(\hat{\Phi}_{xx}, \hat{\Phi}_{xy}, \hat{\Phi}_{ux}, \hat{\Phi}_{uy}) \tag{2.8a}$$

$$\underset{n=0,1,\ldots}{\text{subject to}} \quad \hat{Z}_{AB}\begin{bmatrix}\hat{\Phi}_{xx} & \hat{\Phi}_{xy}\\ \hat{\Phi}_{ux} & \hat{\Phi}_{uy}\end{bmatrix} = \hat{Z}_{I0}, \quad \begin{bmatrix}\hat{\Phi}_{xx} & \hat{\Phi}_{xy}\\ \hat{\Phi}_{ux} & \hat{\Phi}_{uy}\end{bmatrix}\hat{Z}_{AC} = \hat{Z}_{I0}^{\mathsf{T}} \tag{2.8b}$$

$$\begin{bmatrix}\hat{\Phi}_{xx} & \hat{\Phi}_{xy}\\ \hat{\Phi}_{ux} & \frac{1}{z}\hat{\Phi}_{uy}\end{bmatrix} \in \tfrac{1}{z}\mathcal{RH}_\infty \cap \mathcal{S}_\Phi, \tag{2.8c}$$

$$(\hat{\Phi}_{xx}\hat{\delta}_x + \hat{\Phi}_{xy}\hat{\delta}_y)[n] \in \mathcal{X}, \quad (\hat{\Phi}_{ux}\hat{\delta}_x + \hat{\Phi}_{uy}\hat{\delta}_y)[n] \in \mathcal{U} \tag{2.8d}$$

describes the task of finding an optimal system level response. The objective functionals $J(\cdot)$ is often obtained by using Eq. (2.6) to convert an objective originally expressed in terms of $(\hat{x}, \hat{u})$ into a system-level equivalent. The constraints are designed to impose both structural and operational restrictions: The *system level parametrization* (Eq. 2.8b) enforces the policy $\hat{K}$ to be internally stabilizing. The *structural constraints* (Eq. 2.8c) encode informational and structural requirements that the policy needs to satisfy;

for instance, time-delays imposed by the controller's hardware. Finally, the *operational constraints* (Eq. 2.8d) encode the restrictions on the state $x[\cdot]$ and control $u[\cdot]$ signals to be satisfied by K when operating the system.

In summary, a (linear) feedback policy can be synthesized by designing and solving the SLS Problem (2.8), offline, then programming the equivalent control policy into physical hardware. Numerically, an SLS problem can be solved by using $\hat{\Phi}(z) = \sum_{n=0}^{\infty} \frac{1}{z}\Phi[n]$ to express the objective and constraint functions in terms of the spectral components of the system level responses. Many relevant objective functionals $J(\cdot)$ and constraints sets $\{\mathcal{S}_\Phi, \mathcal{X}, \mathcal{U}\}$ lead to Problem (2.8) becoming a finite-dimensional convex optimization problem, and thus solvable even for large-scale decentralized systems–a major motivation for the development of the SLS framework [51, 58, 59]. Moreover, when compared to receding-horizon methods, SLS provides a more principled approach to describe structural constraints such as actuation and communication delays in a control network.

2.2 Multi-agent control systems and noncooperative games

In multi-agent systems (Figure 2.4), the centralized approach to control design (Section 2.1) becomes impractical: Decision-making agents are often self-interested, noncooperative, and lack coordination. If designing their policies without accounting for their rivals' actions, these agents risk failing their objectives and, in the worst case, destabilizing the whole system [60]. The framework of dynamic game theory thus emerges as an alternative for the design of efficient, yet strategically stable, control policies [34, 35].

An N_P-player noncooperative game, defined by the tuple

$$\mathcal{G} := (\mathcal{P}, \{\mathrm{U}^p\}_{p\in\mathcal{P}}, \{J^p\}_{p\in\mathcal{P}}), \tag{2.9}$$

describes the problem of players $p \in \mathcal{P} := \{1, 2, \ldots, N_P\}$ each deciding on a strategy $u^p \in \mathrm{U}^p(u^{-p})$ to optimize an objective $J^p : \mathcal{U}^1 \times \cdots \times \mathcal{U}^{N_P} \to \mathbb{R}$.

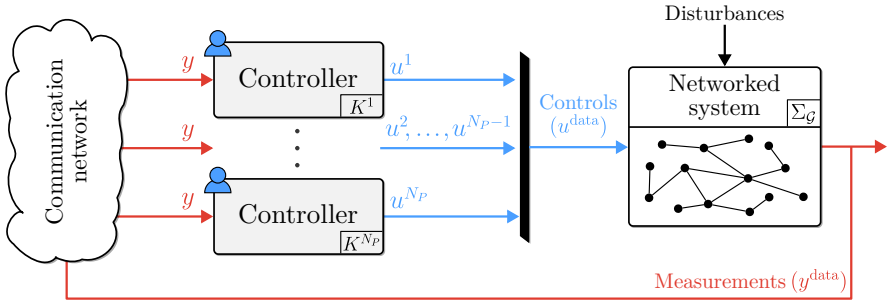

Figure 2.4. Diagram of a multi-agent control system. The measurements y^{data} are sent to a communication network and then broadcasted to the agents (possibly with different delays). The agents simultaneously employ the policies $(K^1, \ldots, K^{N_P})$ to compute their control actions, which are then deployed to the system.

The sets $\mathcal{U}^p$ ($\forall p \in \mathcal{P}$) define the actions available to the players, with $\mathrm{U}^p : \mathcal{U}^{-p} \rightrightarrows \mathcal{U}^p$ restricting this choice based on their rivals' strategies. As such, players are coupled through both their objectives and feasible actions.

A strategy profile $u = (u^1, \dots, u^{N_P})$ is a solution to the game $\mathcal{G}$ when it is feasible and agreeable to all players, if acting rationally. A widely accepted solution concept is that of a *generalized Nash equilibrium* (GNE, [61]). A profile $u^\star = (u^{1^\star}, \dots, u^{N_P^\star})$ is a GNE solution to $\mathcal{G}$ when

$$u^{p^\star} \in \mathrm{BR}^p(u^{-p^\star}) \quad (\forall p \in \mathcal{P}) \tag{2.10}$$

given the best-response mappings $\mathrm{BR}^p : \mathcal{U}^{-p} \rightrightarrows \mathcal{U}^p$,

$$\mathrm{BR}^p(u^{-p}) \coloneqq \underset{u^p}{\arg\min} \left\{ J^p(u^p, u^{-p}) \mid u^p \in \mathrm{U}^p(u^{-p}) \right\}. \tag{2.11}$$

Equivalently, $u^\star \in \mathrm{BR}(u^\star)$ (or, $u^\star \in \mathbf{fix}(\mathrm{BR})$) for the joint-best-response mapping $\mathrm{BR}(u) \coloneqq \mathrm{BR}^1(u^{-1}) \times \cdots \times \mathrm{BR}^{N_P}(u^{-N_P})$. Informally, a GNE is a profile from which no player can improve its outcome by deviating unilaterally.

The GNE solution concept effectively translates $\mathcal{G}$ into a fixed-point search problem, inspiring algorithmic approaches for its solution. We are interested in *equilibrium-seeking algorithms*, when players actively search for their equilibrium policies in a distributed manner. An important class of such algorithms is known as *best-response dynamics* (BRD, [62]). Broadly, they consist of players repeatedly choosing best-responses to their rivals' current strategies until simultaneously converging to a GNE. In addition to their simplicity and decentralized nature, mimicking how noncooperative players would behave in reality, these routines are ensured to converge when the best-response mappings $\{\mathrm{BR}^p\}_{p\in\mathcal{P}}$ are either nonexpansive or contractive [63]. Regrettably, assessing the (sub)continuity properties of such optimizer mappings is a complex task for all but the simplest problems [64], and thus BRD algorithms will often not provide convergence guarantees. Under the refined concept of a *variational generalized Nash equilibrium* (vGNE, [61, 65]), alternative GNE-seeking algorithms can be developed. In particular, this refinement allows the game $\mathcal{G}$ to be reformulated as a monotone inclusion problem and solved using operator-splitting algorithms and related methods [66, 67]. The monotone operator approach provides an unifying framework for the design of equilibrium-seeking routines which apply to a wide range of relevant noncooperative games. A disadvantage of such an approach is that the resulting routines often require a coordinator to perform some centralized computation; that is, the equilibrium-seeking cannot be fully decentralized. Despite their semi-decentralized nature, operator-splitting algorithms for vGNE-seeking are still the state-of-the-art due to their efficient numerics and well-understood convergence properties.

Dynamic noncooperative games

In a (linear and stationary) dynamic noncooperative game, the player's strategies are control signals designed to minimize functionals of the form

$$J^p(u^p, u^{-p}) \coloneqq \mathrm{E}\left[\sum_{n=0}^{\infty} L^p(x[n], u^p[n])\right] \tag{2.12}$$

given the state-space model

$$\Sigma_{\mathcal{G}} : \begin{cases} x[n+1] = Ax[n] + B_w w[n] + \sum_{p\in\mathcal{P}} B_u^p u^p[n]; & \text{(2.13a)} \\ \quad y[n] = Cx[n] + D_w w[n], & \text{(2.13b)} \end{cases}$$

describing the dynamics and measurement process of the underlying networked system. Here, w is a white noise process satisfying $\mathrm{E}w[n] = 0$ and $\mathrm{E}(w[n+k]w[n]^{\mathsf{T}}) = \delta[k]I_{N_w}$ for all $n, k \in \mathbb{N}$. In this case, any solution concept requires an assumption on the information available to each player, that is, the *information structure* of the dynamic game $\mathcal{G}$ under $\Sigma_{\mathcal{G}}$ [34]. We focus on closed-loop imperfect state information (CLIS) patterns, in which the pth player's control actions are represented by linear output-feedback policies

$$K^p : \ell_\infty^{N_y}(\mathbb{N}) \to \ell_\infty^{N_u^p}(\mathbb{N}), \quad u^p \coloneqq K^p y = \Phi_K^p * y, \tag{2.14}$$

given causal operators $K^p \in \mathrm{C}^p(K^{-p}) \subseteq \mathcal{L}(\ell_\infty^{N_y}, \ell_\infty^{N_u^p})$ defined by convolution kernels $\Phi_K^p = (\Phi_K^p[n])_{n\in\mathbb{N}}$. A plan of action with such an information structure is desirable over its open-loop counterpart: it allows players to detect noise disturbances and decision errors, then adapt their actions accordingly. The relations $\mathrm{C}^p : \ell_\infty^{N_u^{-p}} \rightrightarrows \ell_\infty^{N_u^p}$ $(p \in \mathcal{P})$ can be designed to restrict the players' choices to feedback policies that (i) are internally stabilizing and (ii) respect the communication structure of the networked system (imposing action, sensing, and communication delays).

As before, a policy profile $K^\star = (K^{1^\star}, \ldots, K^{N_P^\star})$ is a solution to the dynamic game $\mathcal{G}$ when $K^\star \in \mathbf{fix}(\mathrm{BR}_K)$ given the associated best-response mappings $\{\mathrm{BR}_K^p\}_{p\in\mathcal{P}}$. Since the strategies are feedback policies, this solution concept will be referred to as a *generalized feedback Nash equilibrium* (GFNE). Under this setup, the design of GFNE-seeking algorithms faces additional challenges when compared to the GNE-seeking problem: The best-response mappings $\{\mathrm{BR}_K^p\}_{p\in\mathcal{P}}$, corresponding to infinite-dimensional constrained optimization problems, cannot be analytically or numerically evaluated in general. As such, a GFNE-seeking via BRD routine becomes practically infeasible. Moreover, the policy profile $K \in \mathcal{L}(\ell_\infty^{N_y}, \ell_\infty^{N_u^p})$ lying in a Banach space [68] hinders the refining of the solution concept to a *variational generalized feedback Nash equilibrium* (vGFNE). Thus, vGFNE-seeking via operator-splitting also cannot be realized in practice. The aforementioned technical and fundamental limitations have limited the solution of dynamic games to specific problems and *ad-hoc* algorithms [41–50]. To the best of our knowledge, there are no principled approaches to GFNE-seeking in dynamic games with both operational and informational constraints.

3. Predictive control for self-sufficient resource recovery

This chapter overviews our contributions to the control of conventional WWTPs as self-sufficient WRRFs (Publications I and II). We first present the process and control objectives considered in our study, then overview our proposed Output MPC solution for its operation. The chapter is concluded with a summary of our main experimental findings.

3.1 Water resource recovery facilities

We consider biological wastewater treatment plants (Figure 3.1) comprised of two main processes [20]:

- The activated sludge process, in the *water line*, wherein microorganisms reduce ammonia present in the wastewater into nitrate, which is then reduced into nitrogen gas and released into the atmosphere;
- The anaerobic digestion process, in the *sludge line*, wherein organic material in the excess sludge produced by the plant is converted into biogas, which is then used as fuel for generating heat and electricity.

In general, a biological WWTP can be understood as a process that converts influent wastewater into three main products: treated water, biogas, and disposed sludge. The quality of the treated water is typically characterized by its concentrations of total suspended solids (TSS), organic matter (quantified as biochemical oxygen demand, BOD), and total nitrogen (TN, the sum of organic nitrogen, ammonia, and nitrates/nitrites). The quality of the biogas is characterized by its concentration of methane (CH_4). The quality of the disposed sludge is indirectly characterized by the concentrations of total suspended solids (TSS) in the outflow of the digester. During operation, those measurable quantities can be controlled by manipulating the recirculation of wastewater/sludge throughout the plant, and the addition of oxygen and supplementary carbon to the water line. Additional instrumentation (e.g., sensors for dissolved oxygen and temperature) can be deployed to further monitor the plant and assist in its operation.

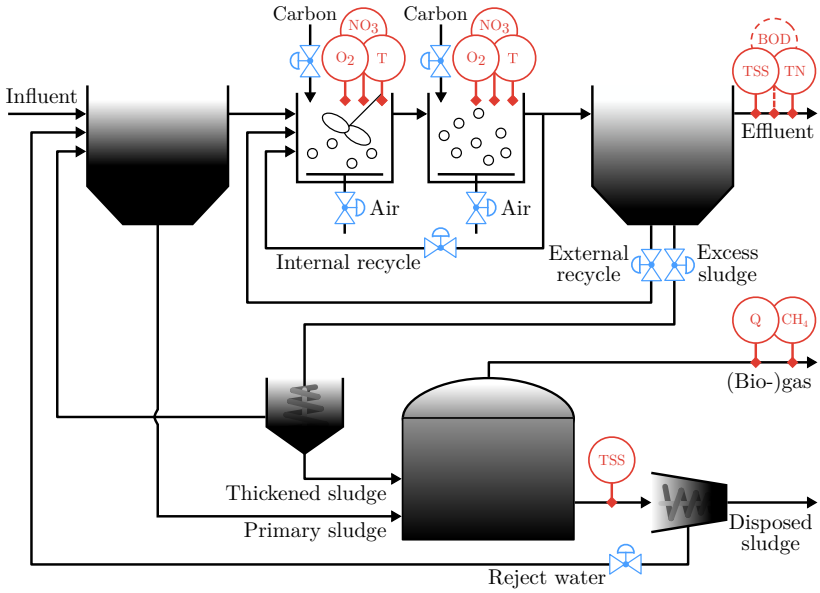

Figure 3.1. Process layout of a biological WWTP. The solid balloons refer to quantities measurable through online sensors, while the dashed balloons refer to quantities measurable only through laboratory analysis. (This figure is adapted from Publication II, © The authors).

Under the emerging water resource recovery paradigm, a WWTP should be operated to produce effluents of tailored quality while simultaneously performing energy recovery. In the following, we formalize a class of control objectives for WWTPs to be operated as a (energetically) self-sufficient WRRFs. Finally, the section is concluded with a discussion on the dynamical modelling of WWTPs and on their key performance indicators.

Water quality classes and energy cost index

The primary control objective for this class of WRRFs consists of targeting different qualities for its effluent water based on (time-varying) external demands. In this direction, we conceptualize three *water quality classes*:

Class A) Water suitable for sustaining natural water bodies;

Class B) Water suitable for urban and industrial activities;

Class C) Water suitable for agricultural irrigation and landscaping;

The classification is defined via upper limits on the effluent quality metrics (that is, concentrations of suspended solids, organic matter, and nitrogen) such that no harmful side-effects are expected if the treated water is used for the aforementioned purposes. We consider the limits reported in Table 3.1, based on existing regulations and well-established standards [69–72]. Intuitively, the energy required to treat the influent wastewater should be smaller the less stringent the quality limits. As such, it would be cost-efficient to produce Class A water only when the external demands cannot

Table 3.1. Water quality classes and their characterization in terms of total suspended solids (TSS), biochemical oxygen demand (BOD), and total nitrogen (TN). (This table is adapted from Publication II, © The authors).

Water class	Target application	Biochemical profile		
		TSS	BOD	TN
A	Environmental reuse	≤ 30 g/m^3	≤ 10 g/m^3	≤ 15 g/m^3
B	Urban/industrial reuse	≤ 30 g/m^3	≤ 15 g/m^3	≤ 30 g/m^3
C	Agricultural reuse	≤ 30 g/m^3	≤ 20 g/m^3	≤ 45 g/m^3

be satisfied with either Class B or Class C water. Similarly, producing Class B water should be avoided when Class C water suffices. Importantly, producing Class C water is beneficial not only because it is energy-inexpensive, but also due to nitrogen being a critical commodity for agriculture. The limit on the nitrogen content of this water is only meant to avoid issues that might be caused by the oversupply of nutrients [73].

An energy-autonomous WRRF should be capable of adjusting its biogas production to recover the energetic costs incurred by its operation. For this class of WWTPs, the net demand of electricity and heating is quantified by the energy cost index (ECI, in kWh/d) signal,

$$\mathrm{ECI} = \mathrm{AE} + \mathrm{PE} + \mathrm{ME} - \eta_E \mathrm{MP} + \max(0, \mathrm{HE} - \eta_H \mathrm{MP}), \tag{3.1}$$

given the aeration (AE, kWh/d), pumping (PE, kWh/d), mixing (ME, kWh/d), and heating (HE, kWh/d) energy required for operation. Energy generated from effluent biogas is expressed as methane production (MP, in kg CH_4/d) converted to electricity (with efficiency $\eta_E > 0$) and heat (with efficiency $\eta_H > 0$). A WWTP is characterized as energetically self-sufficient when its ECI, averaged over a specified period, is nonpositive.

State-space representation and key performance indicators

The complexity inherent to the bioprocesses in a WWTP leads to its dynamics being often represented by continuous-time nonlinear state-space models,

$$\Sigma_{\text{Plant}} : \begin{cases} \frac{d}{dt}x(t) = f(x(t), w(t), u(t)); & \text{(3.2a)} \\ y(t) = g(x(t), w(t), u(t)); & \text{(3.2b)} \\ z(t) = h(x(t), w(t), u(t)). & \text{(3.2c)} \end{cases}$$

The model Σ_{Plant} results from concatenating the state x, controls u (i.e., recycle, air, and external carbon flow-rates), disturbances w (i.e., influent wastewater and its composition), and measurable quantities y, of each process unit in the WWTP. In addition to the dynamics $f(\cdot)$ and measurement process $g(\cdot)$, this extended model also describes how a vector of key performance indicators (or KPIs) $z(t) \in \mathbb{R}^{N_z}$ is computed based on the state $x(t)$,

disturbances $w(t)$, and inputs $u(t)$, at time $t \in \mathbb{R}_{\geq 0}$. For the aforementioned energy-autonomous water resource recovery tasks, we consider

$$z = (\text{TSS}^{\text{eff}},\ \text{BOD}_5^{\text{eff}},\ \text{TN}^{\text{eff}},\ \text{ECI}), \tag{3.3}$$

comprising the biochemical profile of the effluent water (in terms of solids, organics, and nitrogen) and the plant's energy cost index, respectively. Hereafter, we assume no feedthrough of inputs such that $y(t) = g(x(t), w(t))$.

3.2 Output MPC: Overview and setup for resource recovery

We consider Output MPC (Section 2.1) as a systematic approach for designing controllers that operate WWTPs as energy-autonomous WRRFs. Our proposal consists of a *reference-tracking* control policy which regulates the plant's state to time-varying operating points whose associated KPIs reflect the desired objective: For the plant to maintain an energy-autonomous operation regime while producing reuse water of specified quality.

The controller (Figure 3.2) is designed to operate with a *control period* of $\Delta t = 1/96$ [days], each cycle starting at time $t_n = n\Delta t$ ($n = 1, 2, \ldots$). Due to the continuous-time and nonlinear nature of the plant's model, our formulation incorporates an *operating point optimizer* (OPO) which operates at a longer period $\Delta t_{\text{ref}} = 1$ [days] and is responsible for: (i) determining the operating conditions that, if enforced, lead to the currently desired profile of KPIs, and for (ii) producing a (discrete-time) linear approximation of the model Σ_{Plant} around this operating point. Under this setup, the MPC and MHE problems are jointly designed to perform output-feedback disturbance rejection such that the plant is regulated around these optimal operating points. In the following, we detail the main components of this controller. The reader is referred to Publications I and II for the technical details of its implementation and the explicit values of its tuning parameters.

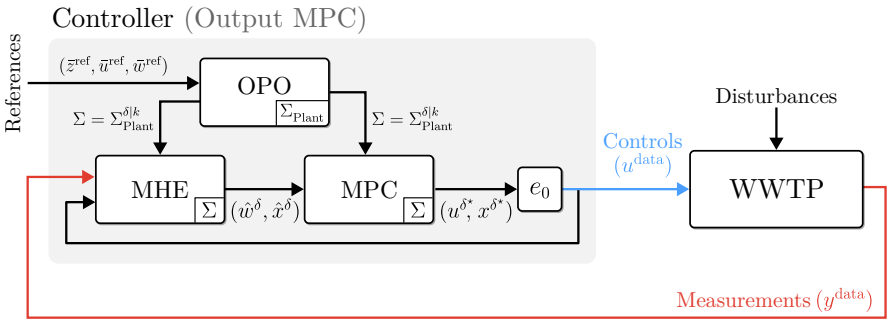

Figure 3.2. Diagram of the Output MPC policy including an operating point optimizer.

Operating point optimizer (OPO)

At each nth control cycle satisfying $n = (\Delta t_{\text{ref}}/\Delta t)k$ ($k = 1, 2, \ldots$), the OPO computes a pair $(x_k^{\text{ref}}, u_k^{\text{ref}})$ solving the nonconvex optimization problem

$$\underset{x_k,\, u_k}{\text{minimize}} \quad \left\| \begin{bmatrix} W_{z|\text{ref}} & \\ & W_{u|\text{ref}} \end{bmatrix} \begin{bmatrix} z_k - \bar{z}_k^{\text{ref}} \\ u_k - \bar{u}_k^{\text{ref}} \end{bmatrix} \right\|_2^2 \tag{3.4a}$$

$$\text{subject to} \quad 0 = f(x_k, \bar{w}_k^{\text{ref}}, u_k) \tag{3.4b}$$

$$z_k = h(x_k, \bar{w}_k^{\text{ref}}, u_k) \tag{3.4c}$$

$$z_k \in \mathcal{Z}_{\text{ref}}, \quad x_k \in \mathcal{X}_{\text{ref}}, \quad u_k \in \mathcal{U}_{\text{ref}}. \tag{3.4d}$$

The references $(\bar{z}_k^{\text{ref}}, \bar{w}_k^{\text{ref}}, \bar{u}_k^{\text{ref}})$ are, respectively, a desired KPI profile, a desired actuation setup, and the expected disturbances, at the kth OPO cycle. The weighting matrices $W_{z|\text{ref}}, W_{u|\text{ref}} \succeq 0$ and constraint sets $\mathcal{Z}_{\text{ref}}$, $\mathcal{X}_{\text{ref}}$, and $\mathcal{U}_{\text{ref}}$, (Eq. 3.4c) are user-defined and configured according to the desired objective. Specifically, to implement energy-autonomous resource recovery, we have $\bar{z}^{\text{ref}}$ represent a desired water quality (that is, we set reference values of TSS^{eff}, $\text{BOD}_5^{\text{eff}}$, and TN^{eff}, for each class in Table 3.1), while having $\mathcal{Z}^{\text{ref}}$ designed to restrict the ECI to be nonpositive.

Using the optimal operating point, the OPO produces the state-space

$$\Sigma_{\text{Plant}}^{\delta|k} : \begin{cases} x^\delta[n+1] = A_k x^\delta[n] + B_{w|k} w^\delta[n] + B_{u|k} u^\delta[n]; & \text{(3.5a)} \\ \quad y^\delta[n] = C_k x^\delta[n] + D_{w|k} w^\delta[n], & \text{(3.5b)} \end{cases}$$

obtained by linearizing the original model Σ_{Plant} around $p_k = (x_k^{\text{ref}}, \bar{w}_k^{\text{ref}}, u_k^{\text{ref}})$, then discretizing (in time) the resulting model using the period Δt. The signals $\{x^\delta, w^\delta, u^\delta\}$ describe deviations around p_k, with $y^\delta = y - g(x_k^{\text{ref}}, \bar{w}_k^{\text{ref}})$. We refer to Publications I and II for more details on these computations.

Model predictive control (MPC)

At each nth control cycle, the mapping $\text{MPC}(\cdot)$ solves the problem

$$\underset{u^\delta[\cdot],\, x^\delta[\cdot]}{\text{minimize}} \quad \sum_{n=0}^{N_c-1} \left\| \begin{bmatrix} W_{x|n} & \\ & W_{u|n} \end{bmatrix} \begin{bmatrix} x^\delta[n] \\ u^\delta[n] \end{bmatrix} \right\|_2^2 + \left\| W_{x|N_c} x^\delta[N_c] \right\|_2^2 \tag{3.6a}$$

$$\underset{n=0,\ldots,N_c-1}{\text{subject to}} \quad \Sigma_{\text{Plant}}^{\delta|k} \text{ (Eq. 3.5) with } w^\delta[n] = \hat{w}_{N_e-1}^\delta \text{ and } x^\delta[0] = \hat{x}_{N_e}^\delta, \tag{3.6b}$$

$$x_{\text{lb}}^\delta \preceq x^\delta[n] \preceq x_{\text{ub}}^\delta, \quad u_{\text{lb}}^\delta \preceq u^\delta[n] \preceq u_{\text{ub}}^\delta. \tag{3.6c}$$

The Problem (3.6) is designed to use the most recent linearized state-space model $\Sigma_{\text{Plant}}^{\delta|k}$ obtained by the OPO. Under this representation, the objective function (Eq. 3.6a) is chosen such that the controller drives the state x^δ and input u^δ deviations around the origin, i.e., it regulates the plant around the operating point p_k currently being tracked. The bound constraints (Eq.

3.6c) enforce technological restrictions related to the WWTP that must be satisfied by the controller (e.g., volume and actuation limits).

Due to the state constraints (Eq. 3.6c), the terminal weighting matrix $W_{x|N_c}$ requires careful tuning to ensure that this control policy is stabilizing. In this work, we obtain $W_{x|N_c}$ from the matrix solving the algebraic Riccati equation (ARE) associated with the unconstrained version of Problem (3.6). Under this setup, as the constraints sets $\{\mathcal{X}_{c|(\cdot)}, \mathcal{U}_{c|(\cdot)}\}$ are convex, compact, and contain the origin in their interior, the MPC policy is stabilizing [19].

Moving horizon estimation (MHE)

At each nth control cycle, the mapping $\mathrm{MHE}(\cdot)$ solves the problem

$$\underset{w^\delta[\cdot],\, x^\delta[\cdot]}{\text{minimize}} \quad \sum_{n=0}^{N_e-1} \left\| \begin{bmatrix} W_{y|n} & \\ & W_{w|n} \end{bmatrix} \begin{bmatrix} y^\delta[n] - y_n^{\text{data}|\delta} \\ w^\delta[n] - w_n^{\text{data}|\delta} \end{bmatrix} \right\|_2^2 + \left\| W_{y|N_e} \left(y^\delta[N_e] - y_{N_e}^{\text{data}|\delta} \right) \right\|_2^2 \tag{3.7a}$$

$$\text{subject to} \quad \Sigma_{\text{Plant}}^{\delta|k} \text{ (Eq. 3.5) with } u^\delta[n] = u_n^{\text{data}|\delta}, \tag{3.7b}$$

$$n=0,\ldots,N_e-1 \qquad x_{\text{lb}}^\delta \preceq x^\delta[n] \preceq x_{\text{ub}}^\delta, \quad w_{\text{lb}}^\delta \preceq w^\delta[n] \preceq w_{\text{ub}}^\delta. \tag{3.7c}$$

The Problem (3.7) considers the same model as in the MPC (Problem 3.6), that is, the most recent linearized model $\Sigma_{\text{Plant}}^{\delta|k}$ from the OPO. The objective function (Eq. 3.7a) is designed so that the estimator tracks the measurement data $y^{\text{data}|\delta} = y^{\text{data}} - g(x_k^{\text{ref}}, u_k^{\text{ref}})$ and some pre-specified disturbance reference $w^{\text{data}|\delta}$ (explicitly, we set $w^{\text{data}|\delta} = 0$). In practice, the estimation objective is to produce state estimates that closely match the measurement data while having the disturbance estimates match the expected profile $\bar{w}_k^{\text{ref}}$. Finally, the bound constraints (Eq. 3.7c) are designed to enforce physical constraints related to the WWTP to be satisfied by the estimates (e.g., the fact that influent flow-rates and composition are positive quantities).

The convergence of the MHE requires a careful tuning of the initial weighting matrix $W_{e|0}$. Again, $W_{e|0}$ is obtained from the matrix solving the ARE associated with the unconstrained version of Problem (3.7). Under the assumption that the estimation error is bounded, the proposed Output MPC (i.e., the combined MHE/MPC policy) is asymptotically stable [19].

3.3 Experimental results and discussion

The long-term operation of a conventional WWTP as an energy-autonomous WRRF (Section 3.1) under our Output MPC controller (Section 3.2) is validated in a simulation experiment (we refer to Publication II for more details). The Benchmark Simulation Model no. 2 (BSM2, [74]) is used to represent a large-scale plant subject to typical municipal influents. In this scenario, the targeted water quality alternates periodically between Class A and Class B, with Class C assumed to be the desired quality during cold months.

In this experiment, the OPO (Problem (3.4) and Eq. 3.5) is implemented using the software CasADi/IPOPT [54, 75]. The MPC (Problem 3.6) and the MHE (Problem 3.7) are both computed using an implementation of the structure-exploiting primal-barrier interior point method described in [55]. The computations are performed in a Dell Precision 3650 Tower Workstation (Intel Xeon W-1370P (16) @ 3.60GHz, 64GB of RAM).

Primary objective: Reuse water of tailored quality

The results (Figure 3.3) show that the Output MPC can operate the WWTP to provide water of the required quality class, consistently, despite influent disturbances. The biochemical profile of the treated water (in terms of TSS^{eff}, $\text{BOD}_5^{\text{eff}}$, and TN^{eff}) shows that the control performance is satisfactory except for sporadic rain and storm events. For the solids and organic matter:

- The effluent total suspended solids (TSS^{eff}) are kept under the 30 g/m^3 limit (complying with the requirements for all quality classes), aside from occasional storm events. We remark on the extreme events at the 1^{st} month ($t \in [0, 31)$ days) and 10^{th} month ($t \in [273, 304)$ days), when TSS^{eff} varies notably and stays uncorrected for a long period.
- The effluent 5-days measured biochemical oxygen demand ($\text{BOD}_5^{\text{eff}}$) is kept under the 10 g/m^3 (complying with the requirements for all quality classes) throughout the entire period. Again, the performance is less satisfactory only during extreme events.

The effluent total nitrogen (TN^{eff}), being the nutrient we are interested in recovering, changes periodically according to the targeted water quality:

- During odd months, the controller targets quality Class A by regulating TN^{eff} around 7.5 g/m^3. The performance is mostly consistent, with a noticeable but small decline in the 1^{st} month ($t \in [0, 31)$). In addition to the rain event, the task is more demanding at this period as the lower temperature conditions affect the nitrogen removal efficiency. In contrast, the nitrogen removal efficiency is noticeably higher in the latter months, when temperatures increase. The daily-average signal during odd months is always below the 15 g/m^3 limit, and thus complies with the requirements for Class A water.
- During even months, the controller targets quality Class B by regulating TN^{eff} around 22.5 g/m^3. In this case, performance is less consistent but still satisfactory. The controller struggles to track the reference during the 8^{th} month ($t \in [212, 243)$) and the 10^{th} month ($t \in [273, 304)$). Again, this is because higher temperature conditions amplify the nitrogen removal efficiency of this treatment process. The daily-average signal during even months is within the $(15, 30]$ g/m^3 limits and thus complies with the requirements for Class B water.
- During the cold season (months 3–5), the controller targets quality Class C by regulating TN^{eff} around 37.5 g/m^3. The reference is tracked

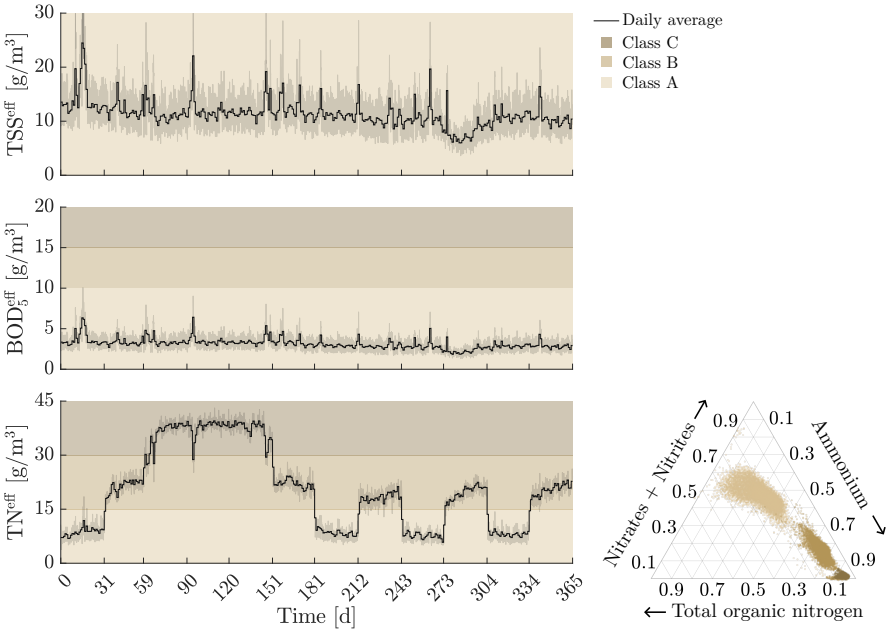

Figure 3.3. Simulation results: Effluent quality in terms of total suspended solids (TSS^{eff}), biochemical oxygen demand (BOD_5^{eff}) and total nitrogen (TN^{eff}). Shaded regions indicate the different water quality classes (A, B, and C). The ternary plot shows the composition of TN^{eff} at each time-step in terms of inorganic and organic nitrogen forms. (This figure is adapted from Publication II, © The authors).

consistently, except for a few isolated events. Specifically, the effluent concentrations of this nutrient decrease abruptly at days $t \approx 65$, $t \approx 95$, and $t \approx 145$. These periods align with the occurrence of rain events, indicating that this operational regime is more sensitive to extreme influent disturbances. Regardless, the daily-average signal during this period is mostly within the $(30, 45]$ g/m^3 limits and thus complies with the requirements for Class C water.

In Figure 3.3, we also visualize the composition of TN^{eff} at each time-step based on the achieved quality class. The results show that the stricter the quality limits, the higher the focus on recovering nitrogen in the form of ammonium. This is not surprising, as the influent nitrogen is mostly comprised of organic and ammonium nitrogen. While the former must be removed to avoid violating the quality limits on organic matter (BOD_5), the latter can be recovered by relaxing the nitrogen removal efficiency of the process. The fact that inorganic nitrogen is preferred for crop fertigation, together with the known toxicity of nitrite, makes this chemical profile ideal for clients requesting water of quality Class C (agricultural reuse). Interestingly, this behavior was not explicitly programmed into the controller, but emerged from its planning given the process model and the expected influent profile.

In summary, the results demonstrate the efficacy of this Output MPC strategy and highlight some facts on wastewater treatment and nitrogen

recovery. Firstly, it reinforces the well-known fact that nitrogen removal via biological treatment is reduced and amplified, respectively, by low and high temperature conditions. The conjecture is that conventional WWTPs operated as WRRFs should be scheduled to recover nutrients during cold seasons, when producing nitrogen-rich reuse water is shown to be a less demanding task. Similarly, the plant can be scheduled to focus on producing cleaner water during warm seasons. Finally, the results show that conventional WWTPs can still remove solids and organic matter while simultaneously manipulating the chemical profile of their water effluents.

Secondary objective: Nonpositive energy costs

The results (Figure 3.4) show that the operation implemented by the Output MPC is energy-autonomous: Except for two short periods (at days $t \approx 15$ and $t \approx 150$), the energy cost index (ECI) indicates that the plant can recover more heat and electrical energy daily than needed for its operation. The violations of the energy-autonomy constraints align with the occurrence of rain events, highlighting the effect of extreme influent disturbances. For this experiment, the WWTP produced a total of 83 GWh, with an average output of 2385 kWh/d of surplus energy. We conjecture that the plant can eliminate the need for buying external electricity (and even benefit from selling its surplus) by including energy-storage systems in its infrastructure. The energy recovery is also shown to depend on the targeted water quality:

- When targeting Class A, the output is 2022 kWh/d, on average. On a monthly-basis, this ranges from producing 1607 kWh/d (1^{st} month, $t \in [0, 31)$ days) to 2178 kWh/d (7^{th} month, $t \in [181, 212)$ days).
- When targeting Class B, the output is 2534 kWh/d, on average. On a monthly-basis, this ranges from producing 1877 kWh/d (6^{th} month, $t \in [151, 181)$ days) to 2933 kWh/d (2^{nd} month, $t \in [31, 59)$ days).
- When targeting Class C, the output is 2635 kWh/d, on average. This production is consistent throughout the cold season.

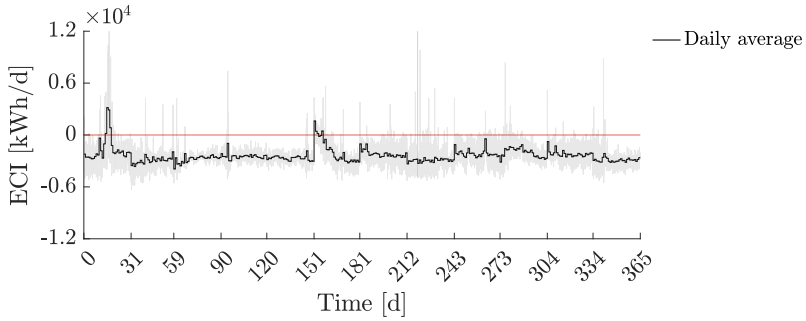

Figure 3.4. Simulation results: Energy cost index (ECI). Negative values imply that the WWTP is recovering more energy than necessary for its operation. (This figure is adapted from Publication II, © The authors).

The energy recovery performance concurs with the conclusions from the nutrient recovery task: It is more inexpensive to produce water of Class A during warm seasons, whereas producing water of Class B or C benefits from low-temperature conditions. In addition, the results show that set-point changes have a noticeable effect on the ECI. Specifically, switching the reference from Class B to Class A causes a short burst in the energy costs, whereas switching from Class A to Class B causes a sudden decrease.

The controller achieves this energy-autonomous operation not by boosting biogas production but instead by implementing cheaper aeration schemes while maintaining a regular methane production. Such a strategy reflects the fact that the controller was explicitly programmed not to minimize energy costs but only to ensure they are nonpositive. Under an alternative tuning, the controller would be able to prioritize energy production, likely at the expense of worsened nutrient recovery. Since our configuration achieves both the primary (resource recovery) and secondary (energy recovery) goals, it constitutes a solution to the problem of operating conventional WWTPs as energy-autonomous WRRFs.

4. Generalized feedback Nash equilibrium seeking

This chapter summarizes our contributions to solving generalized feedback Nash equilibrium problems (Publications III and IV). We demonstrate how the solutions to a relevant class of noncooperative dynamic games can be reformulated using SLS, then enabling the design of tractable GFNE-seeking algorithms. Here, the contribution is methodological and applicable to general networked systems, including, but not limited to, the networked water reuse chains conceptualized in Chapter 1.

4.1 Problem statement

The focus is on noncooperative dynamic games $\mathcal{G} \coloneqq (\mathcal{P}, \{\mathrm{U}^p\}_{p\in\mathcal{P}}, \{J^p\}_{p\in\mathcal{P}})$ defined by the best-response mappings $\mathrm{BR}_K^p : \ell_\infty^{N_u^{-p}} \rightrightarrows \ell_\infty^{N_u^p}$,

$$\mathrm{BR}_K^p(K^{-p}) \coloneqq \underset{u^p := K^p y}{\arg\min} \left\{ J^p(u^p, u^{-p}) \mid K^p \in \mathrm{C}^p(K^{-p}),\ u^p \in \mathrm{U}^p(u^{-p}) \right\}, \quad (4.1)$$

being to the solution set of the infinite-dimensional stochastic problem

$$\underset{u^p := K^p y}{\text{minimize}} \quad \mathrm{E}\left[\sum_{n=0}^{\infty} \left\| \begin{bmatrix} W_x^p & W_u^p \end{bmatrix} \begin{bmatrix} x[n] \\ u^p[n] \end{bmatrix} \right\|_2^2 \right] \quad (4.2a)$$

$$\text{subject to} \quad \Sigma_{\mathcal{G}} \text{ (Eq. 2.13) with } u^{\tilde{p}} = \Phi_K^{\tilde{p}} * y \ (\forall \tilde{p} \in \mathcal{P}), \quad (4.2b)$$

$$\underset{n=0,1,\ldots}{} \quad \Phi_K^p \in \mathcal{S}^p, \ \ (\Phi_K^p, \Phi_K^{-p}) \in \mathcal{S}^{\text{global}}, \quad (4.2c)$$

$$G_{u^p} u^p[n] \preceq 1, \ \ G_x x[n] + G_u u[n] \preceq 1, \quad (4.2d)$$

with the structural constraints $\mathrm{C}^p(\cdot)$ defined by the local $\mathcal{S}^p$ and global $\mathcal{S}^{\text{global}}$ constraint sets, and the operational constraints $\mathrm{U}^p(\cdot)$ defined as linear inequalities given the local G_{u^p} and global (G_x, G_u) matrices. We refer to Publications IV for all the details and assumptions on Problem (4.2). Under this formulation, the GFNE-seeking problem consists of players learning feedback policies that stabilize the underlying networked system $\Sigma_{\mathcal{G}}$ while competing for a limited shared resource (e.g., bandwidth in a network or

goods in a supply-chain) and subjected to a shared communication structure. Although specific, this class of linear-quadratic games covers a wide range of practically relevant control problems (*cf.* Chapter 3).

In this work, we propose to overcome both fundamental and practical limitations to the design of GFNE-seeking algorithms by leveraging the system level synthesis framework to policy design (Section 2.1). Specifically, our solution is centered on a specific, but equivalent, representation of feedback policies, $K = (K^1, \dots, K^{N_P})$, as system level responses to noise, $\Phi = (\Phi^1, \dots, \Phi^{N_P})$. This representation allows us to furnish "*system-level best-responses*" $\{\mathrm{BR}_\Phi^p\}_{p\in\mathcal{P}}$ which correspond to solving finite-dimensional optimal control problems incorporating both the operational and the structural restrictions in the form of convex constraints. Finally, we show that

$$K^\star \in \mathbf{fix}(\mathrm{BR}_K) \iff \Phi^\star \in \mathbf{fix}(\mathrm{BR}_\Phi).$$

The system-level approach to GFNE-seeking is presented in the following.

4.2 System-level best-response mappings

We start by demonstrating how dynamic games admit an equivalent system-level static game through a reformulation of their best-response mappings. In the frequency domain, each player's policy is associated with a transfer matrix $\hat{K}^p = \sum_{n=0}^{\infty} \frac{1}{z^n} \Phi_K^p[n] \in \mathcal{RH}_\infty$, defining the output-feedback $\hat{u}^p = \hat{K}^p \hat{y}$. Considering the $\mathbb{Z}$-transform of the state-space $\Sigma_\mathcal{G}$ (Eq. 2.13), the signals $(\hat{x}, \hat{u}^p)$ can be posed in terms of the noise $\hat{w}$ as

$$\begin{bmatrix} \hat{x} \\ \hat{u}^p \end{bmatrix} = \begin{bmatrix} \hat{\Phi}_{xx} & \sum_{\tilde{p}\in\mathcal{P}} \hat{\Phi}_{xx} B_u^{\tilde{p}} \hat{K}^{\tilde{p}}, \\ \hat{K}^p C \hat{\Phi}_{xx} & \hat{K}^p + \sum_{\tilde{p}\in\mathcal{P}} \hat{K}^p C \hat{\Phi}_{xx} B_u^{\tilde{p}} \hat{K}^{\tilde{p}} \end{bmatrix} \begin{bmatrix} \hat{\delta}_x \\ \hat{\delta}_y \end{bmatrix} \coloneqq \begin{bmatrix} \hat{\Phi}_{xx} & \hat{\Phi}_{xy} \\ \hat{\Phi}_{ux}^p & \hat{\Phi}_{uy}^p \end{bmatrix} \begin{bmatrix} \hat{\delta}_x \\ \hat{\delta}_y \end{bmatrix} \tag{4.3}$$

where $\Phi_{xx} = \left((zI - A) - \sum_{p\in\mathcal{P}} B_u^p \hat{K}^p C\right)^{-1}$, $\hat{\delta}_x = B_w \hat{w}$, and $\hat{\delta}_y = D_w \hat{w}$. As in Section 2.1, the *system level responses* $(\hat{\Phi}_{xx}, \hat{\Phi}_{xy}, \hat{\Phi}_{ux}^p, \hat{\Phi}_{uy}^p)$ describe how the state and the players' actions react to the noise under the output-feedback policy K^p. For ease of notation, let $\hat{\Phi}_{ux} = (\hat{\Phi}_{ux}^p)_{p\in\mathcal{P}}$ and $\hat{\Phi}_{uy} = (\hat{\Phi}_{uy}^p)_{p\in\mathcal{P}}$. In favour of a time-domain exposition, we will refer to these transfer matrices mostly through their spectral factors $\Phi_{(\cdot)} = (\Phi_{(\cdot)}[n])_{n\in\mathbb{N}}$. Under this setup, the following holds (see Publications III and IV for the full details):

(a) The policy profile $K = (K^1, \dots, K^{N_P})$ is internally stabilizing if the spectral factors of the profile $(\hat{\Phi}_{xx}, \hat{\Phi}_{xy}, \hat{\Phi}_{ux}, \hat{\Phi}_{uy})$ satisfy the system

$$\Sigma_{\Phi|\mathcal{G}} : \begin{cases} \Phi_{xx}[n+1] = A\Phi_{xx}[n] + \sum_{p\in\mathcal{P}} B_u^p \Phi_{ux}^p[n]; & \text{(4.4a)} \\ \Phi_{xy}[n+1] = A\Phi_{xy}[n] + \sum_{p\in\mathcal{P}} B_u^p \Phi_{uy}^p[n]; & \text{(4.4b)} \end{cases}$$

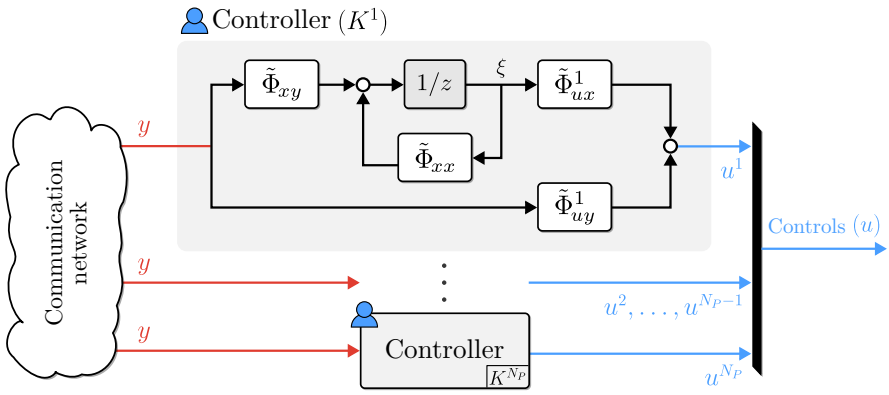

Figure 4.1. Feedback structure for the policies $\hat{K}^p = \hat{\Phi}^p_{uy} + \hat{\Phi}^p_{ux}\hat{\Phi}^{-1}_{xx}\hat{\Phi}_{xy}$ using the auxiliary transfer matrices $\tilde{\Phi}_x = z(I - z\hat{\Phi}_{xx})$, $\tilde{\Phi}^p_{ux} = z\hat{\Phi}^p_{ux}$, $\tilde{\Phi}_{xy} = -z\hat{\Phi}_{xy}$ and $\tilde{\Phi}^p_{uy} = \hat{\Phi}^p_{uy}$.

and its "dual"

$$\Sigma^*_{\Phi|\mathcal{G}} : \begin{cases} \Phi_{xx}[n+1] = \Phi_{xx}[n]A + \Phi_{xy}[n]C; & \text{(4.5a)} \\ \Phi^p_{ux}[n+1] = \Phi^p_{ux}[n]A + \Phi^p_{uy}[n]C, \quad (\forall p \in \mathcal{P}); & \text{(4.5b)} \end{cases}$$

with $\Phi_{xx}[0] = \Phi_{xy}[0] = \Phi^p_{ux}[0] = 0$ (for all $p \in \mathcal{P}$) and $\Phi_{xx}[1] = I_{N_x}$.

(b) A profile $(\hat{\Phi}_{xx}, \hat{\Phi}_{xy}, \hat{\Phi}_{ux}, \hat{\Phi}_{uy})$ satisfying $(\Sigma_{\Phi|\mathcal{G}}, \Sigma^*_{\Phi|\mathcal{G}})$ is achieved by the policies $\hat{K}^p = \hat{\Phi}^p_{uy} - \hat{\Phi}^p_{ux}\hat{\Phi}^{-1}_{xx}\hat{\Phi}_{xy}$ (Figure 4.1), implemented as

$$\Sigma_{K^p} : \begin{cases} \xi[n+1] = -\sum_{k=0}^{n} \tilde{\Phi}_{xx}[k]\xi[n-k] - \sum_{k=0}^{n} \tilde{\Phi}_{xy}[k]y[n-k]; & \text{(4.6a)} \\ u^p[n] = \quad \sum_{k=0}^{n} \tilde{\Phi}^p_{ux}[k]\xi[n-k] + \sum_{k=0}^{n} \Phi^p_{uy}[k]y[n-k]. & \text{(4.6b)} \end{cases}$$

with auxiliary kernels $\tilde{\Phi}_{xx} = (\Phi_{xx}[n+2])_{n\in\mathbb{N}}$, $\tilde{\Phi}_{xy} = (\Phi_{xy}[n+1])_{n\in\mathbb{N}}$, and $\tilde{\Phi}_{ux} = (\Phi_{ux}[n+1])_{k\in\mathbb{N}}$, and *internal state* $\xi = (\xi[n])_{n\in\mathbb{N}}$ with $\xi[0] = 0$.

The system level parametrization enables a methodology for policy synthesis consisting of searching the space of stabilizing policies directly through $(\Phi_{xx}, \Phi_{xy}, \Phi^p_{ux}, \Phi^p_{uy})$, for each $p \in \mathcal{P}$. In this direction, consider that the players plan their control actions by choosing a desired system level response $\Phi^p_u = (\Phi^p_{ux}, \Phi^p_{uy})$ to the noise, then implementing their feedback policy as the system Σ_{K^p} (Eq. 4.6). The original best-response mappings $\{\mathrm{BR}^p_K\}_{p\in\mathcal{P}}$ are equivalent to system-level best-response mappings $\{\mathrm{BR}^p_\Phi\}_{p\in\mathcal{P}}$,

$$\mathrm{BR}^p_\Phi(\Phi^{-p}_u) \coloneqq \arg\min_{\Phi^p_u} \left\{ J^p_\Phi(\Phi^p_u, \Phi^{-p}_u) \mid \Phi^p_u \in \mathrm{C}^p_\Phi(\Phi^{-p}_u), \quad \Phi^p_u \in \mathrm{U}^p_\Phi(\Phi^{-p}_u) \right\}, \tag{4.7}$$

corresponding to the solution set of the system level synthesis problem

$$\underset{\Phi^p_{ux},\Phi^p_{uy}}{\text{minimize}} \quad \sum_{n=0}^{\infty} \left\| \begin{bmatrix} W_x^p & W_u^p \end{bmatrix} \begin{bmatrix} \Phi_{xx}[n] & \Phi_{xy}[n] \\ \Phi^p_{ux}[n] & \Phi^p_{uy}[n] \end{bmatrix} \begin{bmatrix} B_w \\ D_w \end{bmatrix} \right\|_F^2 \tag{4.8a}$$

$$\underset{n=0,1,\dots}{\text{subject to}} \quad \Sigma_{\Phi|\mathcal{G}} \text{ (Eq. 4.4)}, \;\; \Sigma^*_{\Phi|\mathcal{G}} \text{ (Eq. 4.5)} , \tag{4.8b}$$

$$\begin{bmatrix} \Phi_{xx}[n] & \Phi_{xy}[n] \\ \Phi^p_{ux}[n] & \Phi^p_{uy}[n] \end{bmatrix} \in \mathcal{S}^p_\Phi[n], \tag{4.8c}$$

$$\begin{bmatrix} G_x & G_u \\ & \widetilde{G}_{u^p} \end{bmatrix} \begin{bmatrix} \Phi_{xx} & \Phi_{xy} \\ \Phi_{ux} & \Phi_{uy} \end{bmatrix} \begin{bmatrix} B_w \\ D_w \end{bmatrix} * w[n] \preceq 1, \tag{4.8d}$$

where we let $\widetilde{G}_{u^p} = e_p^\mathsf{T} \otimes G_{u^p}$ to simplify its statement. The translation from the structural constraints C^p to the equivalent C^p_Φ is done by (i) enforcing internal stability using the system level parametrization (Eq. 4.8b), and by (ii) enforcing any structural specification directly to the spectral factors of the closed-loop maps (Eq. 4.8c), consequently restricting the structure of the policy realization Σ_{K^p} (Eq. 4.6). In this formulation, we remark that the players' policies are only linked through the (Φ_{xx}, Φ_{xy}) responses. Finally, the translation from the operational constraints U^p to U^p_Φ (Eq. 4.8d) follows directly from the definition of $(\Phi_{xx}, \Phi_{xy}, \Phi^p_{ux}, \Phi^p_{uy})$ in Eq. (4.3).

The mappings $\{\mathrm{BR}^p_\Phi\}_{p\in\mathcal{P}}$ are still intractable as Problem (4.8) is also infinite-dimensional and stochastic. In this case, however, both issues can be tackled by a special design of the structural constraints $\{\mathcal{S}^p_\Phi\}_{p\in\mathcal{P}}$ and by a robust reformulation of the operational constraints $\{\mathrm{U}^p_\Phi\}_{p\in\mathcal{P}}$.

Structural constraints: Finite-impulse responses and sparsity patterns

The problem in BR^p_Φ ($p \in \mathcal{P}$) can be made finite-dimensional by restricting the choice of K^p to policies with finite-impulse response (FIR) kernels. Using the constraint sets $\mathcal{S}^p_\Phi = (\mathcal{S}^p_\Phi[n])_{n\in\mathbb{N}}$, this can be enforced by defining

$$\mathcal{S}^p_\Phi[n] = \left\{ \begin{bmatrix} \Phi_{xx} & \Phi_{xy} \\ \Phi^p_{ux} & \Phi^p_{uy} \end{bmatrix} \;\middle|\; \begin{bmatrix} \Phi_{xx} & \Phi_{xy} \\ \Phi^p_{ux} & \Phi^p_{uy} \end{bmatrix} = 0 \right\}, \quad \text{for all } n \geq N, \tag{4.9}$$

given a FIR horizon of $N > 1$. In practice, this is equivalent to requiring $\Phi_{xx}[N] = \Phi_{xy}[N] = \Phi^p_{ux}[N] = \Phi^p_{uy}[N] = 0$ then restricting Problem (4.8) to the finite sequence of spectral factors $(\Phi_{xx}[n], \Phi_{xy}[n], \Phi^p_{ux}[n], \Phi^p_{uy}[n])_{n=1}^N$ and $\Phi^p_{uy}[0]$. The best-response mappings $\{\mathrm{BR}^p_\Phi\}_{p\in\mathcal{P}}$ are then equivalent to solving finite-dimensional optimization problems, with such solutions still being equivalent representations of the policies that solve the infinite-dimensional Problem (4.2). Additionally, the FIR kernels simplify the realization of the policies K^p by fixing the total number of summand terms in Σ_{K^p} (Eq. 4.6).

Remark 4.1. The terminal constraints induced by Eq. (4.9) cannot be satisfied when the system $\Sigma_\mathcal{G}$ is not controllable and observable. While this

constraint can be relaxed (e.g., to take the form $\|\Phi[N]\|_2 \leq \gamma$ for a $\gamma > 0$), the stability and sub-optimality properties of the resulting control policies are well-understood only for the state-feedback case [51, 76].

In this work, we design the remaining constraint sets $(\mathcal{S}_\Phi^p[n])_{n=0}^N$ to encode actuation and communication delays imposed by the structure of the networked system. Specifically, we use the sparsity constraints

$$\mathcal{S}_\Phi^p[n] = \left\{ \begin{bmatrix} \Phi_{xx} & \Phi_{xy} \\ \Phi_{ux}^p & \Phi_{uy}^p \end{bmatrix} \;\middle|\; \mathrm{Sp}\begin{bmatrix} \Phi_{xx} & \Phi_{xy} \\ \Phi_{ux}^p & \Phi_{uy}^p \end{bmatrix} = \mathrm{Sp}\begin{bmatrix} A^{\alpha_n} & A^{\alpha_n} C^\mathsf{T} \\ {B_u^p}^\mathsf{T} A^{\alpha_n} & {B_u^p}^\mathsf{T} A^{\alpha_{n+1}} C^\mathsf{T} \end{bmatrix} \right\}, \tag{4.10}$$

with $\alpha_n = \max(0, \lfloor (n - d_a)/d_c \rfloor)$ and $\mathrm{Sp}[\,\cdot\,]$ denoting the support of a matrix. Under this setup, the policy K^p ($p \in \mathcal{P}$) satisfy that the ith state component $(x_i[n])_{n\in\mathbb{N}}$ can only be affected by the ith action component $([B_u u^p[n]]_i)_{n\in\mathbb{N}}$ after an *actuation delay* of $d_a \geq 0$ time-steps. In addition, it specifies that measurements propagate in the communication network (assumed to have the same topology as the dynamical system) with a *communication delay* of $d_c > 0$. The policies satisfying this information pattern are referred to as (d_a, d_c)-*delayed feedback policies*. Intuitively, the delays (d_a, d_c) restrict the space of feasible control policies, thus affecting the performance of the closed-loop system (when compared to a policy with instantaneous actuation/communication). The ability to impose sparsity to feedback policies, a central feature of the SLS framework, enables the solution of GFNE problems with asymmetric information patterns; a major challenge in the field [77, 78]. Finally, we note that this choice is not restrictive and our proposed GFNE-seeking algorithms generalize to a broader class of informational constraints (such as those discussed in [58]).

Operational constraints: Robust and chance constraints

The problem in BR_Φ^p ($p \in \mathcal{P}$) can be made deterministic by using the available information about the noise process w to replace the operational constraints (Eq. 4.8d) with robust equivalents (see [53, Chapter 4.4]). If the filtered noise process is known to be bounded, $(B_w, D_w)w \in \mathcal{W} = \{P_w \zeta \mid \|\zeta\|_{\ell_\infty} \leq 1\}$, then Problem (4.8) can be rendered robust by considering

$$\begin{aligned} \sup \left\{ \begin{bmatrix} G_x & G_u \\ & \widetilde{G}_{u^p} \end{bmatrix}_{i,:} \begin{bmatrix} \Phi_{xx} & \Phi_{xy} \\ \Phi_{ux} & \Phi_{uy} \end{bmatrix} P_w * \zeta[n] \;\middle|\; \|\zeta\|_{\ell_\infty} \leq 1 \right\} \leq 1 \\ \iff \left\| \left(\begin{bmatrix} G_x & G_u \\ & \widetilde{G}_{u^p} \end{bmatrix}_{i,:} \begin{bmatrix} \Phi_{xx} & \Phi_{xy} \\ \Phi_{ux} & \Phi_{uy} \end{bmatrix} P_w \right)^\mathsf{T} \right\|_{\ell_1} \leq 1, \end{aligned} \tag{4.11}$$

for every ith row of the matrix $\mathrm{col}([G_x \; G_u], [0 \; \widetilde{G}_{u^p}])$. In practice, at the expense of conservativeness, this ensures that the synthesized policies will satisfy the operational constraints for all possible realizations of the noise.

Alternatively, when an explicit bound on the noise process is not available, the operational constraints can be rendered robust using a statistical approach. Under the assumption that w is an additive white Gaussian noise (AWGN) process, then Problem (4.8) can be rendered robust by considering

$$\mathbf{prob}\left(\begin{bmatrix} G_x & G_u \\ & \widetilde{G}_{u^p} \end{bmatrix}_{i,:} \begin{bmatrix} \Phi_{xx} & \Phi_{xy} \\ \Phi_{ux} & \Phi_{uy} \end{bmatrix} \begin{bmatrix} B_w \\ D_w \end{bmatrix} * w[n] \leq 1 \right) \geq \rho$$

$$\iff \left\| \left(\begin{bmatrix} G_x & G_u \\ & \widetilde{G}_{u^p} \end{bmatrix}_{i,:} \begin{bmatrix} \Phi_{xx} & \Phi_{xy} \\ \Phi_{ux} & \Phi_{uy} \end{bmatrix} \begin{bmatrix} B_w \\ D_w \end{bmatrix} \right)^{\mathsf{T}} \right\|_{\ell_2} \leq \frac{1}{Q(\rho)}, \quad (4.12)$$

given a probability $\rho \in (0.5, 1)$ and with $Q : [0, 1] \to \mathbb{R}$ being the quantile function of the standard Normal distribution. We remark that this approach cannot ensure that the synthesized policy K^p will satisfy the operational constraints for all actual realizations of the noise w (even without the Gaussian assumption). However, at the risk of conservativeness, K^p can still be designed to ensure with arbitrarily high probability $\rho > 0.5$ that such constraints will be satisfied during operation, if the noise is Gaussian.

4.3 System-level approach to equilibrium seeking

The equivalence between BR_K and BR_Φ implies that a policy profile $K^\star$ is a GFNE (i.e., $K^\star \in \mathbf{fix}(\mathrm{BR}_K)$) if its corresponding response $\Phi_u^\star$ is a GNE of the associated "system-level" game (i.e., $\Phi_u^\star \in \mathbf{fix}(\mathrm{BR}_\Phi)$). In turn, this implies that searching for an equilibrium policy $K \in \mathcal{L}(\ell_\infty^{N_y}, \ell_\infty^{N_u})$, in a Banach space, can be reformulated as finding an equivalent equilibrium response $\Phi \in \ell_2[0, N]$, in a Hilbert space isomorphic to the Euclidean space (due to being finite-dimensional); that is, the problem is equivalent to a static game. In this section, we leverage this fact to realize our system-level approach to GFNE-seeking in dynamic games. Specifically, two algorithmic solutions are presented: Best-response dynamics and operator-splitting.

4.3.1 Best-response dynamics for GFNE-seeking

The routine for GFNE-Seeking based on the system-level best-responses $\{\mathrm{BR}_\Phi\}_{p\in\mathcal{P}}$ is outlined in Algorithm 1. At each kth update, it consists of each pth player collecting the information about its rival's current system level responses ($\Phi_{u|k}^{-p}$) then computing a best-response ($\Phi_{u|k+1}^{p}$) which will be used to realize its output-feedback policy K_{k+1}^p. The operations can be performed in parallel, without any need for coordination other than having the players exchanging information before evaluating BR_Φ^p.

A distinct feature of this GFNE-Seeking routine is the policy update (line 4): In practice, it corresponds to each player plugging the responses $(\Phi_{xx|k}, \Phi_{xx|k}, \Phi_{ux|k}^p, \Phi_{uy|k}^p)$ on their policy's representation Σ_{K^p} (Eq. 4.6), with (Φ_{xx}, Φ_{xy}) computed according to the system $\Sigma_{\Phi|\mathcal{G}}$ (Eq. 4.4). Importantly,

Algorithm 1: GFNE-Seeking via BRD

Initialize $\Phi_{u|0} = (\Phi^1_{u|0}, \dots, \Phi^{N_P}_{u|0})$;
for $k = 0, 1, 2, \dots$ **do**
  for *each player* $p \in \mathcal{P}$ **do**
    $\hat{K}^p_k \coloneqq \hat{\Phi}^p_{uy|k} - \hat{\Phi}^p_{ux|k} \hat{\Phi}^{-1}_{xx|k} \hat{\Phi}_{xy|k}$;
    $\Phi^p_{u|k+1} \coloneqq (1-\eta)\Phi^p_{u|k} - \eta \mathrm{BR}^p_\Phi(\Phi^{-p}_{u|k})$;

this routine does not require knowledge of the state or input signals and thus players can update their policies while simultaneously operating the underlying networked system. This enables the GFNE-Seeking via BRD algorithm to be performed in games which cannot be interrupted for players to redesign their control policies (for instance, markets and smart grids).

Summary of the experimental results

We evaluate the performance of the GFNE-Seeking via BRD routine in two exemplary problems on the noncooperative control of networked systems (see Publication III for the full details). The systems are fully-observed and thus the players operate their subsystems via state-feedback policies. The results (Figure 4.2) demonstrate the efficacy of the routine in approaching an equilibrium policy profile $\hat{K}^\star = (\hat{\Phi}^{p^\star}_{uy} - \hat{\Phi}^{p^\star}_{ux}(\hat{\Phi}^\star_{xx})^{-1}\hat{\Phi}^\star_{xy})_{p\in\mathcal{P}}$. The convergence rates are shown to be linear and dependent on the structure of the problem: In the first problem, a 3-player game defined for an unstable chain networked system, a distance of order 10^{-6} to a (unique) GFNE solution is achieved under 400 iterations. In the second problem, a 4-player dynamic pricing game for an oligopolistic market, a (non-unique) GFNE solution is approached at a slower rate—a distance of order 10^{-6} is achieved only after 900 iterations. In both cases, players are restricted to (1,1)-delayed policies and must satisfy a set of operational constraints for all realizations of the bounded noise; but only for the second problem are the operational constraints coupled. The policies are shown to be stabilizing and to have satisfactory performance in operating the systems alongside the execution of the GFNE-seeking routine (see Publication III, Figures 5 and 8).

The simplicity and decentralized nature of the GFNE-Seeking via BRD algorithm make this routine appealing for practical implementation. However, despite the empirical evidence, its convergence properties are still difficult to establish *a priori*: The convergence to a GFNE requires BR_Φ to be L_{BR_Φ}-Lipschitz with $L_{\mathrm{BR}_\Phi} \leq 1$ [63]. Regrettably, aside from very specific cases (see Publication III), obtaining such a constant is demanding even for the linear-quadratic SLS Problem (4.8) due to analytical formulas for the mappings $\{\mathrm{BR}^p_\Phi\}_{p\in\mathcal{P}}$ not existing. Despite this limitation, this routine is still a valid approach to many control problems in which players must agree to an equilibrium policy with minimal coordination.

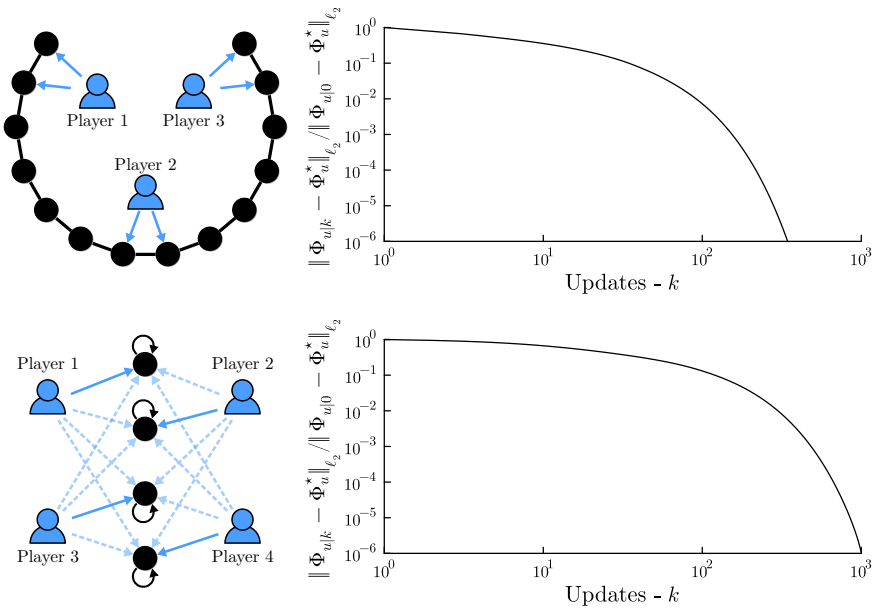

Figure 4.2. GFNE-Seeking via BRD: The network topology and convergence of the equilibrium-seeking routine for two exemplary dynamic games. In the first problem **(top)**, players operate distinct sections of a chain-networked system with the goal of stabilizing its state. In the second problem **(bottom)**, players apply price changes to their product (solid lines) to regulate their demands, indirectly affecting their competitors' products (dashed lines). (Adapted from Publication III, © The authors, reproduced from IEEE under CC BY 4.0).

4.3.2 Operator-splitting for vGFNE-seeking

In this section, we let $\Phi_x = [\Phi_{xx}\ \Phi_{xy}]$, $\Phi_u^p = [\Phi_{ux}^p\ \Phi_{uy}^p]$, and $W_w = (B_w, D_w)$. The concept of a *variational generalized feedback Nash equilibria* (vGFNE) can be defined indirectly through the system level parametrization: A policy profile $K^\star = (K^{1^\star}, \dots, K^{N_P^\star})$ with $\hat{K}^{p^\star} = \hat{\Phi}_{uy}^{p^\star} - \hat{\Phi}_{ux}^{p^\star}(\hat{\Phi}_{xx}^\star)^{-1}\hat{\Phi}_{xy}^\star$ is a vGFNE when $\Phi_u^\star = (\Phi_u^{1^\star}, \dots, \Phi_u^{N_P^\star})$ is a vGNE of the game defined by the objectives

$$J_\Phi^p(\Phi_u^p, \Phi_u^{-p}) = \sum_{n=0}^{N} \left\| \begin{bmatrix} W_x^p & W_u^p \end{bmatrix} \begin{bmatrix} \Phi_x[n] \\ \Phi_u^p[n] \end{bmatrix} W_w \right\|_F^2 \quad (\forall p \in \mathcal{P}),$$

and the global set of feasible strategies

$$\mathcal{U}_\Phi = \{\Phi_u \mid \Sigma_{\Phi|\mathcal{G}} \text{ (Eq. 4.4)},\ \ \Sigma_{\Phi|\mathcal{G}}^* \text{ (Eq. 4.5)},\ \ \Phi_x[N] = \Phi_u[N] = 0,$$

$$\mathrm{Sp} \begin{bmatrix} \Phi_x[n] \\ \Phi_u[n] \end{bmatrix} = \mathrm{Sp} \begin{bmatrix} A^{\alpha_n} & A^{\alpha_n} C^\mathsf{T} \\ B_u^\mathsf{T} A^{\alpha_n} & B_u^\mathsf{T} A^{\alpha_n + 1} C^\mathsf{T} \end{bmatrix}, \quad (\forall n),$$

$$\left\| \left(\begin{bmatrix} G_x & G_u \\ & \widetilde{G}_u \end{bmatrix}_{i,:} \begin{bmatrix} \Phi_x \\ \Phi_u \end{bmatrix} W_w \right)^\mathsf{T} \right\|_{\ell_2} \leq \frac{1}{Q(\rho)}, \quad (\forall i) \Bigg\},$$

Algorithm 2: vGFNE-Seeking via FB-Splitting

Initialize $\Phi_{u|0} = (\Phi^1_{u|0}, \dots, \Phi^{N_P}_{u|0})$;
for $k = 0, 1, 2, \dots$ **do**
  for *each player* $p \in \mathcal{P}$ **do**
    $\hat{K}^p_k \coloneqq \hat{\Phi}^p_{uy|k} - \hat{\Phi}^p_{ux|k} \hat{\Phi}^{-1}_{xx|k} \hat{\Phi}_{xy|k}$;
    $\Phi^p_{u|k+1/2} \coloneqq \Phi^p_{u|k} - \eta \nabla_{\Phi^p_u} J^p_\Phi(\Phi^p_{u|k}, \Phi^{-p}_{u|k})$;
  for *coordinator* **do**
    $\Phi_{u|k+1} \coloneqq \mathrm{proj}_{\mathcal{U}_\Phi}(\Phi_{u|k+1/2})$;

with $\widetilde{G}_u = \mathrm{blkdiag}(G_{u^1}, \dots, G_{u^{N_P}})$. Due to the FIR constraint, the feasible set $\mathcal{U}_\Phi$ is finite-dimensional and $\Phi_x = (\Phi_x[n])_{n=1}^N$ and $\Phi^p_u = (\Phi^p_u[n])_{n=0}^N$ can be represented as matrices by stacking their spectral factors. Using standard results from monotone operator theory and variational analysis [61], the vGFNE-seeking problem can be posed as the monotone inclusion problem

$$\text{Find } \Phi^\star_u \text{ such that } 0 \in F_\Phi(\Phi^\star_u) + N_{\mathcal{U}_\Phi}(\Phi^\star_u), \tag{4.13}$$

given the pseudo-gradient $F_\Phi(\Phi_u) = (\nabla_{\Phi^p_u} J^p_\Phi(\Phi^p_u, \Phi^{-p}_u))_{p \in \mathcal{P}}$ and with $N_{\mathcal{U}_\mathcal{G}}$ being the normal cone operator defined for the closed convex set $\mathcal{U}_\mathcal{G}$.

Problem (4.13) can be approached using operator splitting methods [79]. In general, these require the evaluation of the operators F_Φ, $N_{\mathcal{U}_\Phi}$, and/or their resolvents. A fundamental instance of such algorithms is the *forward-backward splitting* (FB-Splitting), based on the relationship

$$0 \in F_\Phi(\Phi^\star_u) + N_{\mathcal{U}_\Phi}(\Phi^\star_u) \iff \Phi^\star_u \in (I + \eta N_{\mathcal{U}_\Phi})^{-1}(I - \eta F_\Phi)(\Phi^\star_u),$$

for any given $\eta > 0$; that is, a vGNE must satisfy $\Phi^\star_u \in \mathbf{fix}(T) \subseteq \mathbf{fix}(\mathrm{BR}_\Phi)$ for the operator $T = (I + \eta N_{\mathcal{U}_\Phi})^{-1} \circ (I - \eta F_\Phi)$. The operator T induces a fixed-point iteration $\Phi_{u|k+1} = T(\Phi_{u|k})$ which converges to a unique vGNE (thus, a vGFNE of $\mathcal{G}$) whenever T is a contraction. Due to its composition form, this iteration results in the (semi)decentralized routine for vGFNE-Seeking via FB-Splitting in Algorithm 2. At each $k > 0$, it consists of two main steps:

(1) The players propose new policies to improve their objectives, simultaneously, given the current profile and disregarding the constraints.
(2) A coordinator collects these proposals and then computes (and broadcasts to the players) the closest admissible policy profile.

Computation

Numerically, the gradient $\nabla_{\Phi^p_u} J^p_\Phi(\cdot)$ can be evaluated through automatic differentiation and the projection operator $\mathrm{proj}_{\mathcal{U}_\Phi}(\cdot)$ by directly solving the associated optimization problem. However, the structure of the game $\mathcal{G}$ allows for more efficient numerics. The gradient can be computed as

$$\nabla_{\Phi^p_u} J^p(\cdot) = \Big(2\big((W^{p\mathsf{T}}_u W^p_u)\Phi^p_{u|k}[n] + (B^p_u)^\mathsf{T} \Delta^p_{x|k}[n]\big)\big(W_w W_w^\mathsf{T}\big)\Big)_{n=0}^N, \tag{4.14}$$

with sensitivities $\Delta^p_{x|k} = (\Delta^p_{x|k}[n])_{n=0}^N$ obtained by first forward-propagating $\Phi_{x|k}$ according with the system $\Sigma_{\Phi|\mathcal{G}}$ (Eq. 4.4), then backward-propagating

$$\Delta^p_{x|k}[n-1] = A^\mathsf{T}\Delta^p_{x|k}[n] + (W_x^{p\mathsf{T}}W_x^p)\Phi_{x|k}[n], \quad \Delta^p_{x|k}[N] = 0.$$

The projection map

$$\mathrm{proj}_{\mathcal{U}_\Phi}(\Phi_{u|k+1/2}) = \arg\min_{\Phi_u} \left\{ \|\Phi_{u|k+1/2} - \Phi_u\|_{\ell_2} \;\middle|\; \Phi_u \in \mathcal{U}_\Phi \right\} \tag{4.15}$$

is a sparse convex program that can be solved efficiently. However, it still has $(N{+}1)N_u(N_x{+}N_y)$ decision variables and $NN_x(N_x{+}N_y)$ equality constraints and thus scales poorly with the state-space dimensions (N_x, N_u, N_y) and FIR horizon N. The computational burden can be alleviated by exploiting the partially-separable structure of Problem (4.15) and solving it using the Alternating Direction Method of Multipliers (ADMM, [79]). In particular, Problem (4.15) can be decomposed into subproblems which are often row-wise and column-wise separable and can be further reduced into smaller problems to be solved in parallel. This leads to a substantial performance improvement, making Algorithm 2 scalable to very large-scale problems. We refer to [59] for more details on partially separable SLS problems.

Convergence properties

If the learning rate satisfies $\eta \in (0, 2M_{F_\Phi}/L^2_{F_\Phi})$, with M_{F_Φ} and L_{F_Φ} being the strong-monotonicity and Lipschitz constants of F_Φ, then Algorithm 2 converges to an unique vGFNE $\hat{K}^\star = (\hat{\Phi}^{p\star}_{uy} - \hat{\Phi}^{p\star}_{ux}\hat{\Phi}^{\star^{-1}}_{xx}\hat{\Phi}^\star_{xy})_{p\in\mathcal{P}}$, parametrised by $\Phi^\star_u \in \mathbf{fix}(\mathrm{BR}_\Phi)$. Moreover, the convergence is linear with rate

$$\lim_{k\to\infty} \frac{\|\Phi_{u|k+1} - \Phi^\star_u\|_{\ell_2}}{\|\Phi_{u|k} - \Phi^\star_u\|_{\ell_2}} = \sqrt{1 - \eta(2M_{F_\Phi} - \eta L^2_{F_\Phi})}. \tag{4.16}$$

which also corresponds to the Lipschitz constant L_T for the forward-backward operator $T = (I + \eta N_{\mathcal{U}_\Phi})^{-1} \circ (I - \eta F_\Phi)$. In this particular problem,

$$M_{F_\Phi} = \lambda_{\min}(H_\Phi + H_\Phi^\mathsf{T})\sigma^2_{\min}(W_w), \qquad L_{F_\Phi} = 2\sigma_{\max}(H_\Phi)\sigma^2_{\max}(W_w), \tag{4.17}$$

where $H_\Phi = [(H_\Phi^{pp})^\mathsf{T}(H_\Phi^{p\tilde{p}})]_{p,\tilde{p}\in\mathcal{P}}$ is formed by the block-triangular matrices,

$$\left[H_\Phi^{p\tilde{p}}\right]_{0\le n,n'\le N}, \quad \left[H_\Phi^{p\tilde{p}}\right]_{n,n'} = \begin{cases} W_u^p & \text{if } n = n' \text{ and } p = \tilde{p}, \\ W_x^p A^{(n-1)-n'} B_u^{\tilde{p}} & \text{if } n > n', \\ 0 & \text{otherwise.} \end{cases} \tag{4.18}$$

We refer to Publication IV for more details. In practice, this result implies that there is always a small enough learning rate $\eta > 0$ that can ensure the convergence of Algorithm 2. However, such a convergence can still be slow if F_Φ is ill-conditioned. Interestingly, note that the pseudo-gradient F_Φ is ill-conditioned if so is the noise-filtering matrix W_w, and that convergence rates are not affected by the noise if $W_w = I$. Moreover, it suggests that open-loop unstable games (when $\rho(A) \ge 1$) might require a careful tuning of W_x^p $(\forall p)$ to avoid the terms $W_x^p A^n B_u^{\tilde{p}}$ $(n = 1, \ldots, N{-}1)$ in Eq. (4.18) from exploding, then affecting the condition number of F_Φ.

Summary of the experimental results

We evaluate the performance of the vGFNE-Seeking via FB-Splitting routine in an exemplary problem on the noncooperative control of a decentralized power-grid (see Publication IV for the full details). The system is marginally stable, partially-observed, and players operate their individual subsystems via output-feedback policies. The players are restricted to (1,1)-delayed policies and must satisfy some operational constraints; in particular, they must ensure (with probability $\eta = 0.975$) that pairwise exchanges in the network are within a specified safety limit. The results (Figure 4.3) demonstrate the efficacy of the routine in approaching an equilibrium policy profile $\hat{K}^\star = (\hat{\Phi}_{uy}^{p^\star} - \hat{\Phi}_{ux}^{p^\star}(\hat{\Phi}_{xx}^\star)^{-1}\hat{\Phi}_{xy}^\star)_{p\in\mathcal{P}}$. Convergence is shown to be linear, with a distance of order 10^{-12} to the (unique) vGFNE solution being achieved under 10^4 iterations. Interestingly, despite slow, it is still roughly 100 times faster than that predicted using the Lipschitz constant $L_T \approx 0.99995$ (Eq. 4.16). The policies are also shown to be jointly stabilizing and to have satisfactory performance in operating the power grid alongside the execution of the vGFNE-Seeking routine (see Publication IV, Figure 4).

When compared to the GFNE-Seeking via BRD routine (Algorithm 1), the vGFNE-Seeking via FB-Splitting routine has the disadvantage of requiring an unbiased coordinator. However, it enjoys better computational properties and its convergence can be easily established even in the presence of the generalized constraints. Moreover, the coordinator can promote an additional layer of security and privacy: If each pth player is provided with $\Phi_x = [\Phi_{xx} \;\; \Phi_{xy}]$ (i.e., it is computed and broadcasted by the coordinator), then they do not need to have information about the other players to update or implement their policies. Regardless, the routine can still be made less centralized by having the players determine these quantities by sharing local information with their neighbors [66, 67].

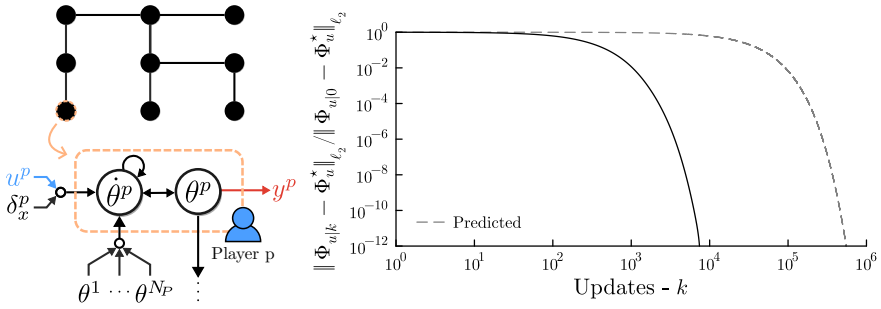

Figure 4.3. vGFNE-Seeking via FB-Splitting: The network topology and convergence of the equilibrium-seeking routine for an exemplary power-grid system. Each pth subsystem, controlled by the pth player, is observed by measuring the phase deviation θ^p and manipulated by applying a load (u^p) directly on its frequency deviation $\dot{\theta}^p$, which is affected by the phase of neighboring subsystems. (Adapted from Publication IV, © The authors).

5. Concluding remarks

Automatic decision-making is an important tool for the transition into circular water resource recovery infrastructures. Predictive control can render the conventional biological WWTP flexible to resource recovery objectives; thus enabling their repurposing for centralized water reuse chains. Equilibrium-seeking allows noncooperative agents to learn, simultaneously, optimal control policies to operate their subsystems in a decentralized network; thus paving the way for networked water reuse chains. Furthermore, predictive control and equilibrium-seeking are methods that can potentially complement the design of future water resource recovery systems.

The thesis presents novel results on the predictive control of WWTPs for resource recovery, using state-of-the-art methods and widely accepted benchmark systems. In Publication I, we conceptualized our Output MPC control framework and analyzed its components: the operating-point optimizer (OPO), the model predictive control (MPC), and the moving-horizon estimator (MHE). Publication II then employs this framework and obtains a proof-of-concept on the operation of a full-scale WWTP for producing reuse water of tailored quality given external, time-varying, demands. The control policy is also shown to enforce an energy-positive operation regime by regulating biogas production, implying that biological WWTPs have potential as renewable power plants. Finally, being highly customizable, our controller can be extended to different objectives (e.g., phosphorus recovery), influent conditions (e.g., non-domestic wastewater), and treatment systems (e.g., membrane bioreactors) than those considered in this study.

The thesis also presents important contributions to the field of dynamic games. In Publications III and IV, we propose algorithms which address the previously unsolved problem of GFNE-seeking of feedback policies which are stabilizing and subject to both operational and structural constraints. In general, their main contribution is the enabling of noncooperative agents to synthesize optimal control policies while operating large-scale, constrained, and decentralized, networked systems. Our approach is principled (based on the SLS methodology) and valid for a broad class of relevant problems, including, but not limited to, the class of constrained linear-quadratic dy-

namic games. The routines are shown to converge linearly to a GFNE solution, obtaining near-equilibrium solutions within a reasonable number of policy updates. Due to their generality, these methods apply to multi-agent problems in many different domains, such as microgrids and autonomous driving. Nevertheless, they would still cover the control of our conceptualized networked water reuse chains–considering, for instance, the *linearize-then-control* approach taken in Publications I and II.

5.1 Future research directions

Looking forward, the contributions of this thesis open several opportunities for further developments. Here, we discuss some relevant directions.

The predictive controller we propose deals with nonlinear dynamics by incorporating a steady-state optimization problem which is nonconvex, thus affecting its resilience. Addressing this issue, while maintaining the overall control architecture, a possible direction for future work would be to develop convex alternatives to the operating point optimizer. The online identification of linear dynamical models [76, 80] is a promising approach which also solves the need for an accurate first-principles model of the plant. Alternatively, learning convex steady-state models using data-based approaches [81] and designing Lyapunov functions using sum-of-squares [82] are methods that could enable an efficient search of operational points.

The system-level approach to GFNE-seeking can be refined in many directions. The algorithms we propose are based on well-established, but classical, fixed-point iteration methods: Adapting them to modern methods, such as the *preconditioned operator splitting* [67], could enhance their performance and practical applicability. For systems with nonlinear dynamics, future work should also consider adapting SLS to nonlinear systems or having the players collectively construct a linearized model of the networked system. Finally, we note that our solution is based on the existence of a global network model, shared with all agents. An interesting research problem concerns the design of GFNE-seeking algorithms for problems in which the players can only access information pertaining to their subsystems.

The paradigm of water resource recovery is still maturing, and there are many exciting challenges to be solved. A (perhaps next) task to be addressed is the investigation of different water reuse network topologies and the development of related benchmark problems. Such a research problem draws challenges not only from a technological perspective, but also from the urban planning required for devising a realistic concept. Finally, it remains to extend our control solutions to novel wastewater treatment technologies and, possibly, to applications in other domains.

References

[1] W. Verstraete, , and S. E. Vlaeminck, "ZeroWasteWater: short-cycling of wastewater resources for sustainable cities of the future," *Int. J. of Sustain. Dev. & World Ecol.*, vol. 18, no. 3, pp. 253–264, 2011.

[2] United Nations World Water Assessment Programme, *Wastewater: The untapped resource*, ser. The United Nations World Water Development Report 2017. UNESCO, 2017.

[3] S. Ricart and A. M. Rico, "Assessing technical and social driving factors of water reuse in agriculture: A review on risks, regulation and the yuck factor," *Agric. Water Manag.*, vol. 217, pp. 426–439, 2019.

[4] S. Kundu, B. K. Pramanik, P. Halder, S. Patel, M. Ramezani, M. Khairul, M. H. Marzbali, J. Paz-Ferreiro, S. Crosher, G. Short, A. Surapaneni, D. Bergmann, and K. Shah, "Source and central level recovery of nutrients from urine and wastewater: A state-of-art on nutrients mapping and potential technological solutions," *J. of Environ. Chem. Eng.*, vol. 10, no. 2, p. 107146, 2022.

[5] G. Sarpong and V. G. Gude, "Near future energy self-sufficient wastewater treatment schemes," *Int. J. of Environ. Res.*, vol. 14, pp. 479–488, 2020.

[6] P. Kehrein, M. van Loosdrecht, P. Osseweijer, M. Garfí, J. Dewulf, and J. Posada, "A critical review of resource recovery from municipal wastewater treatment plants - market supply potentials, technologies and bottlenecks," *Environ. Sci.: Water Res. & Technol.*, vol. 6, pp. 877–910, 2020.

[7] M. Mainardis, D. Cecconet, A. Moretti, A. Callegari, D. Goi, S. Freguia, and A. G. Capodaglio, "Wastewater fertigation in agriculture: Issues and opportunities for improved water management and circular economy," *Env. Pollution*, vol. 296, p. 118755, 2022.

[8] A. F. Meese, D. J. Kim, X. Wu, L. Le, C. Napier, M. T. Hernandez, N. Laroco, K. G. Linden, J. Cox, P. Kurup, J. McCall, D. Greene, M. Talmadge, Z. Huang, J. Macknick, K. A. Sitterley, A. Miara, A. Evans, K. Thirumaran, M. Malhotra, S. G. Gonzalez, P. Rao, J. Stokes-Draut, and J.-H. Kim, "Opportunities and Challenges for Industrial Water Treatment and Reuse," *ACS EST Eng.*, vol. 2, no. 3, 2022.

[9] B. Valverde-Pérez, E. Ramin, B. F. Smets, and B. G. Plósz, "EBP2R - an innovative enhanced biological nutrient recovery activated sludge system to produce growth medium for green microalgae cultivation," *Water Res.*, vol. 68, pp. 821–830, 2015.

[10] J. Wan, J. Gu, Q. Zhao, and Y. Liu, "COD capture: a feasible option towards energy self-sufficient domestic wastewater treatment," *Sci. Rep.*, vol. 6, no. 1, p. 25054, 2016.

[11] M. Xie, H. K. Shon, S. R. Gray, and M. Elimelech, "Membrane-based processes for wastewater nutrient recovery: Technology, challenges, and future direction," *Water Res.*, vol. 89, pp. 210–221, 2016.

[12] L. Peng, H. Dai, Y. Wu, Y. Peng, and X. Lu, "A comprehensive review of phosphorus recovery from wastewater by crystallization processes," *Chemosphere*, vol. 197, pp. 768–781, 2018.

[13] B. Wu, "Membrane-based technology in greywater reclamation: A review," *Sci. of the Total Env.*, vol. 656, pp. 184–200, 2019.

[14] K. Li, Q. Liu, F. Fang, R. Luo, Q. Lu, W. Zhou, S. Huo, P. Cheng, J. Liu, M. Addy, P. Chen, D. Chen, and R. Ruan, "Microalgae-based wastewater treatment for nutrients recovery: A review," *Bioresour. Tech.*, vol. 291, p. 121934, 2019.

[15] K. Solon, M. Jia, and E. I. P. Volcke, "Process schemes for future energy-positive water resource recovery facilities," *Water Sci. and Tech.*, vol. 79, no. 9, pp. 1808–1820, 2019.

[16] M. Faragò, A. Damgaard, J. A. Madsen, J. K. Andersen, D. Thornberg, M. H. Andersen, and M. Rygaard, "From wastewater treatment to water resource recovery: Environmental and economic impacts of full-scale implementation," *Water Res.*, vol. 204, p. 117554, 2021.

[17] K. Zhou and J. C. Doyle, *Essentials of robust control*. Prentice-Hall, 1998.

[18] D. E. Seborg, T. F. Edgar, D. A. Mellichamp, and F. J. Doyle III, *Process Dynamics and Control*, 4th ed. John Wiley & Sons, 2016.

[19] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl, *Model Predictive Control: Theory, Computation and Design*, 2nd ed. Nob Hill Publ., LLC., 2020.

[20] M. Henze, M. C. van Loosdrecht, G. A. Ekama, and D. Brdjanovic, *Biological wastewater treatment*. IWA publishing, 2008.

[21] G. Olsson, B. Carlsson, J. Comas, J. Copp, K. Gernaey, P. Ingildsen, U. Jeppsson, C. Kim, L. Rieger, I. Rodrìguez-Roda, J. Steyer, I. Takács, P. Vanrolleghem, A. Vargas, Z. Yuan, and L. Åmand, "Instrumentation, control and automation in wastewater - from London 1973 to Narbonne 2013," *Water Sci. & Technol.*, vol. 69, no. 7, pp. 1372–1385, 2014.

[22] L. Åmand, G. Olsson, and B. Carlsson, "Aeration control - A review," *Water Sci. and Technol.*, vol. 67, no. 1, pp. 2374–2398, 2013.

[23] B. Holenda, E. Dokomosa, A. Rèdey, and J. Fazakas, "Dissolved oxygen control of the activated sludge wastewater treatment process using model predictive control," *Comput. & Chem. Eng.*, vol. 33, no. 6, pp. 1270–1278, 2008.

[24] M. Mulas, S. Tronci, F. Corona, H. Haimi, P. Lindell, M. Heinonen, R. Vahala, and R. Baratti, "Predictive control of an activated sludge process: An application to the Viikinmäki wastewater treatment plant," *J. of Process Control*, vol. 35, pp. 89–100, 2015.

[25] I. Santín, C. Pedret, R. Vilanova, and M. Meneses, "Advanced decision control system for effluent violations removal in wastewater treatment plants," *Control Eng. Pract.*, vol. 49, pp. 60–75, 2016.

[26] W. Shen, X. Chen, and J. Corriou, "Application of model predictive control to the BSM1 benchmark of wastewater treatment process," *Comput. and Chem. Eng.*, vol. 22, no. 12, pp. 2849–2856, 2008.

[27] H.-G. Han, H.-H. Qian, and J.-F. Qiao, "Nonlinear multiobjective model-predictive control scheme for wastewater treatment process," *J. of Process Control*, vol. 24, no. 3, pp. 47–59, 2014.

[28] P. Vega, S. Revollar, M. Francisco, and J. M. Martín, "Integration of set point optimization techniques into nonlinear MPC for improving the operation of WWTPs," *Comput. and Chem. Eng.*, vol. 68, pp. 78–95, 2014.

[29] M. Sadeghassadi, C. J. B. Macnab, B. Gopaluni, and D. Westwick, "Application of neural networks for optimal-setpoint design and MPC control in biological wastewater treatment," *Comput. and Chem. Eng.*, vol. 115, pp. 150–160, 2018.

[30] M. Ekman, "Bilinear black-box identification and MPC of the activated sludge process," *J. of Process Control*, vol. 18, no. 7–8, pp. 643–653, 2008.

[31] C. Foscoliano, S. Del Vigo, M. Mulas, and S. Tronci, "Predictive control of an activated sludge process for long term operation," *Chem. Eng. J.*, vol. 304, pp. 1031–1044, 2016.

[32] M. A. Massoud, A. Tarhini, and J. A. Nasr, "Decentralized approaches to wastewater treatment and management: Applicability in developing countries," *J. of Env. Management*, vol. 90, no. 1, pp. 652–659, 2009.

[33] N. K. Singh, A. A. Kazmi, and M. Starkl, "A review on full-scale decentralized wastewater treatment systems: techno-economical approach," *Water Sci. & Technol*, vol. 71, no. 4, pp. 468–478, 2014.

[34] T. Başar and G. J. Olsder, *Dynamic noncooperative game theory*. SIAM, 1998.

[35] J. P. Hespanha, *Noncooperative game theory: An introduction for engineers and computer scientists*. Princeton University Press, 2017.

[36] L. Pavel, *Game Theory for Control of Optical Networks*. Birkhäuser, 2012.

[37] S. Maghsudi and S. Stańczak, "Channel selection for network-assisted D2D communication via no-regret bandit learning with calibrated forecasting," *IEEE Transactions on Wireless Communications*, vol. 14, no. 3, pp. 1309–1322, 2015.

[38] L. Chen, N. Li, S. H. Low, and J. C. Doyle, "Two market models for demand response in power networks," in *2010 First IEEE International Conference on Smart Grid Communications*, 2010, pp. 397–402.

[39] A.-H. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid," *IEEE Transactions on Smart Grid*, vol. 1, no. 3, pp. 320–331, 2010.

[40] W. Saad, Z. Han, H. V. Poor, and T. Basar, "Game-theoretic methods for the smart grid: An overview of microgrid systems, demand-side management, and smart grid communications," *IEEE Signal Processing Magazine*, vol. 29, no. 5, pp. 86–105, 2012.

[41] J. Engwerda, W. Van Den Broek, and J. M. Schumacher, "Feedback Nash equilibria in uncertain infinite time horizon differential games," in *Proceedings of the 14th International Symposium of Mathematical Theory of Networks and Systems, MTNS 2000*, 2000, pp. 1–6.

[42] D. Vrabie and F. Lewis, "Integral reinforcement learning for online computation of feedback Nash strategies of nonzero-sum differential games," in *49th IEEE Conference on Decision and Control*. IEEE, 2010, pp. 3066–3071.

[43] G. Kossioris, M. Plexousakis, A. Xepapadeas, A. de Zeeuw, and K.-G. Mäler, "Feedback Nash equilibria for non-linear differential games in pollution control," *Journal of Economic Dynamics and Control*, vol. 32, no. 4, pp. 1312–1331, 2008.

[44] R. Kamalapurkar, J. R. Klotz, and W. E. Dixon, "Concurrent learning-based approximate feedback-Nash equilibrium solution of N-player nonzero-sum differential games," *IEEE/CAA Journal of Automatica Sinica*, vol. 1, no. 3, pp. 239–247, 2014.

[45] P. V. Reddy and G. Zaccour, "Feedback Nash equilibria in linear-quadratic difference games with constraints," *IEEE Transactions on Automatic Control*, vol. 62, no. 2, pp. 590–604, 2017.

[46] ——, "Open-loop and feedback Nash equilibria in constrained linear–quadratic dynamic games played over event trees," *Automatica*, vol. 107, pp. 162–174, 2019.

[47] A. Tanwani and Q. Zhu, "Feedback Nash equilibrium for randomly switching differential–algebraic games," *IEEE Transactions on Automatic Control*, vol. 65, no. 8, pp. 3286–3301, 2019.

[48] D. Fridovich-Keil, E. Ratner, L. Peters, A. D. Dragan, and C. J. Tomlin, "Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games," in *2020 IEEE International Conference on Robotics and Automation*. IEEE, 2020, pp. 1475–1481.

[49] D. Fridovich-Keil, V. Rubies-Royo, and C. J. Tomlin, "An iterative quadratic method for general-sum differential games with feedback linearizable dynamics," in *2020 IEEE International Conference on Robotics and Automation*. IEEE, 2020, pp. 2216–2222.

[50] F. Laine, D. Fridovich-Keil, C.-Y. Chiu, and C. Tomlin, "The computation of approximate generalized feedback Nash equilibria," *SIAM Journal on Optimization*, vol. 33, no. 1, pp. 294–318, 2023.

[51] J. Anderson, J. C. Doyle, S. H. Low, and N. Matni, "System level synthesis," *Annual Reviews in Control*, vol. 47, pp. 364–393, 2019.

[52] C. V. Rao, S. J. Wright, and J. B. Rawlings, "Application of interior-point methods to model predictive control," *Journal of Pptimization Theory and Applications*, vol. 99, no. 3, pp. 723–757, 1998.

[53] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.

[54] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, 2006.

[55] Y. Wang and S. Boyd, "Fast model predictive control using online optimization," *IEEE Transactions on Control Systems Technology*, vol. 18, no. 2, pp. 267–278, 2009.

[56] S. Richter, S. Mariethoz, and M. Morari, "High-speed online MPC based on a fast gradient method applied to power converter control," in *Proceedings of the 2010 American Control Conference*. IEEE, 2010, pp. 4737–4743.

[57] G. Frison and J. B. Jørgensen, "Efficient implementation of the Riccati recursion for solving linear-quadratic control problems," in *2013 IEEE International Conference on Control Applications (CCA)*. IEEE, 2013, pp. 1117–1122.

[58] Y.-S. Wang, "A system level approach to optimal controller design for large-scale distributed systems," Dissertation (Ph.D.), California Institute of Technology, 2017.

[59] Y.-S. Wang, N. Matni, and J. C. Doyle, "Separable and localized system-level synthesis for large-scale systems," *IEEE Transactions on Automatic Control*, vol. 63, no. 12, pp. 4234–4249, 2018.

[60] S.-H. Wang and E. Davison, "On the stabilization of decentralized control systems," *IEEE Transactions on Automatic Control*, vol. 18, no. 5, pp. 473–478, 1973.

[61] F. Facchinei and C. Kanzow, "Generalized Nash equilibrium problems," *Annals of Operations Res.*, vol. 175, no. 1, pp. 177–211, 2010.

[62] B. Swenson, R. Murray, and S. Kar, "On best-response dynamics in potential games," *SIAM Journal on Control and Optimization*, vol. 56, no. 4, pp. 2734–2767, 2018.

[63] H. H. Bauschke and P. L. Combettes, *Convex analysis and monotone operator theory in Hilbert spaces*, 2nd ed. Springer, 2017.

[64] R. T. Rockafellar and R. J. Wets, *Variational analysis*. Springer, 1998.

[65] A. A. Kulkarni and U. V. Shanbhag, "On the variational equilibrium as a refinement of the generalized Nash equilibrium," *Automatica*, vol. 48, no. 1, pp. 45–55, 2012.

[66] P. Yi and L. Pavel, "An operator splitting approach for distributed generalized Nash equilibria computation," *Automatica*, vol. 102, pp. 111–121, 2019.

[67] G. Belgioioso, P. Yi, S. Grammatico, and L. Pavel, "Distributed generalized Nash equilibrium seeking: An operator-theoretic perspective," *IEEE Control Systems Magazine*, vol. 42, no. 4, pp. 87–102, 2022.

[68] J. B. Conway, *A course in functional analysis*, 2nd ed. Springer, 2007.

[69] U.S. EPA, *Guidelines for water reuse*, 2012.

[70] B. M. G. L. Alcalde-Sanz, "Minimum quality requirements for water reuse in agricultural irrigation and aquifer recharge - towards a water reuse regulatory instrument at EU level," EUR 28962 EN, Publ. Office of the EU, Tech. Rep., 2017.

[71] Council of the European Union, "Regulation (EU) 2020/741 of the European Parliament and of the Council," 2020, https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32020R0741&from=EN.

[72] F. Shoushtarian and M. Negahban-Azar, "Worldwide regulations and guidelines for agricultural water reuse: a critical review," *Water*, vol. 12, no. 4, p. 971, 2020.

[73] F. Albornoz, "Crop responses to nitrogen overfertilization: A review," *Scientia Horticulturae*, vol. 205, pp. 79–83, 2016.

[74] K. Gernaey, U. Jeppsson, P. Vanrolleghem, and J. Copp, *Benchmarking of Control Strategies for Wastewater Treatment Plants*, ser. Scientific and Technical Report Series No. 23. IWA Publ., 2014.

[75] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.

[76] S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu, "On the sample complexity of the linear quadratic regulator," *Foundations of Computational Mathematics*, vol. 20, no. 4, pp. 633–679, 2020.

[77] A. Nayyar and T. Başar, "Dynamic stochastic games with asymmetric information," in *51st IEEE Conference on Decision and Control*. IEEE, 2012, pp. 7145–7150.

[78] T. Başar, “Stochastic differential games and intricacy of information structures,” in *Dynamic Games in Economics*. Springer, 2014, pp. 23–49.

[79] E. K. Ryu and W. Yin, *Large-scale convex optimization: algorithms & analyses via monotone operators*. Cambridge University Press, 2022.

[80] L. Ljung, *System Identification: Theory for the User*. Prentice-Hall, 1999.

[81] C. M. Bishop and N. M. Nasrabadi, *Pattern recognition and machine learning*. Springer, 2006, vol. 4, no. 4.

[82] A. Papachristodoulou and S. Prajna, “On the construction of Lyapunov functions using the sum of squares decomposition,” in *Proceedings of the 41st IEEE Conference on Decision and Control, 2002.*, vol. 3, 2002, pp. 3482–3487.

Appendix 1. Key terms of the doctoral thesis

Key term: Water resource recovery facility
Definition: A plant which treats wastewater to produce clean water, energy, nutrients, and other resources.
Explanation: A water resource recovery facility is a plant implementing a wastewater treatment process (either biological or chemical) with the purpose of recovering resources contained within the wastewater. The plant might also include a sludge treatment line to process excess sludge from the wastewater treatment. The most common recovered resources include (treated) water, biofuels, and nutrients such as nitrogen and phosphorus.

Key term: Model predictive control
Definition: A model-based method of optimally controlling dynamical systems subjected to operational constraints.
Explanation: Model predictive control is a class of digital feedback controllers in which the control actions (or actionable inputs) are determined in real-time by solving a mathematical optimization problem. The optimization is based on a dynamical model of the system to be controlled, used to predict its response to the planned sequence of actions.

Key term: Equilibrium seeking algorithm
Definition: A class of distributed algorithms for computing an equilibrium solution to a noncooperative game.
Explanation: In non-cooperative game theory, equilibrium-seeking algorithms are computational routines in which players (decision makers) actively seek an equilibrium solution to their problem. Specifically, these are distributed algorithms in which players exchange information and perform the calculations/operations necessary to solve the game—as opposed to a solution being obtained and implemented by some centralized coordinator.

Keskeinen käsite: Vesivarojen talteenottolaitos
Määritelmä: Jätevedenpuhdistuslaitos, joka tuottaa puhdasta vettä, energiaa, ravinteita ja muita resursseja.
Selite: Vesivarojen talteenottolaitos on laitos, jossa toteutetaan jäteveden käsittelyprosessi (joko biologinen tai kemiallinen) jätevedessä olevien resurssien talteenottamiseksi. Laitoksessa voi olla myös lietteen käsittelylinja, jossa käsitellään jäteveden käsittelystä syntyvä ylimääräinen liete. Yleisimpiä talteenotettuja resursseja ovat (käsitelty) vesi, biopolttoaineet ja ravinteet, kuten typpi ja fosfori.

Keskeinen käsite: Mallipohjainen ennakoiva säätö
Määritelmä: Mallipohjainen menetelmä dynaamisten järjestelmien optimaaliseen säätöön, kun niihin kohdistuu operatiivisia rajoituksia.
Selite: Mallipohjainen ennakoiva säätö on digitaalisen takaisinkytkentäohjauksen luokka, jossa säätötoiminnot (tai toimivat ohjaustulot) määritetään reaaliajassa ratkaisemalla matemaattinen optimointiongelma. Optimointi perustuu ohjattavan järjestelmän dynaamiseen malliin, jota käytetään ennustamaan sen vaste suunnitellun toimintasekvenssin perusteella.

Keskeinen käsite: Tasapainohakuinen algoritmi
Määritelmä: Luokka hajautettuja algoritmeja tasapainoisten ratkaisujen laskemiseksi epäyhteistoiminnallisessa pelissä.
Selite: Epäyhteistoiminnallisessa peliteoriassa tasapainoa etsivät algoritmit ovat laskennallisia rutiineja, joissa pelaajat (päätöksentekijät) etsivät aktiivisesti tasapainoisia ratkaisuja ongelmaansa. Tarkemmin sanottuna nämä ovat hajautettuja algoritmeja, joissa pelaajat vaihtavat tietoja ja suorittavat pelin ratkaisemiseen tarvittavat laskelmat/toiminnot—toisin kuin ratkaisu, jonka joku keskitetty koordinaattori hankkii ja panee täytäntöön.

Business, Economy

Art, Design, Architecture

Science, Technology

Crossover

Doctoral Theses

Aalto DT 233/2025

ISBN 978-952-64-2841-3
ISBN 978-952-64-2840-6 (pdf)

Aalto University
School of Chemical Engineering
Department of Chemical and Metallurgical Engineering
aalto.fi