הטכניון
מכון טכנולוגי לישראל
TECHNION
Israel Institute of Technology

ספריות הטכניון
The Technion Libraries

בית הספר ללימודי מוסמכים ע"ש ארווין וג'ואן ג'ייקובס
Irwin and Joan Jacobs Graduate School

Exploiting the Interplay Between Spatial and Temporal Constraints in the Control of Multi-Agent Systems

Gal Barkai

Exploiting the Interplay Between Spatial and Temporal Constraints in the Control of Multi-Agent Systems

Research Thesis

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy

Gal Barkai

Submitted to the Senate
of the Technion — Israel Institute of Technology
Av 5785 Haifa August 2025

This research thesis was carried out under the supervision of Prof. Leonid Mirkin and Prof. Daniel Zelazo, in the Faculty of Mechanical Engineering.

Some results in this thesis have been published as articles by the author and research collaborators in conferences and journals during the course of the author's doctoral research period, the most up-to-date versions of which being:

G. Barkai, L. Mirkin, and D. Zelazo, "On sampled-data consensus: Divide and concur," *IEEE Control Systems Letters*, vol. 6, pp. 343–348, 2022.

G. Barkai, L. Mirkin, and D. Zelazo, "On the internal stability of diffusively coupled multi-agent systems and the dangers of cancel culture," *Automatica*, vol. 155, p. 111 158, 2023.

G. Barkai, L. Mirkin, and D. Zelazo, "An emulation approach to sampled-data synchronization," in *Proc. 62nd IEEE Conf. Decision and Control.*, 2023, pp. 6449–6454.

G. Barkai, L. Mirkin, and D. Zelazo, "An emulation approach to output-feedback sampled-data synchronization," in *Proc. 22nd European Control Conf.*, 2024.

G. Barkai, L. Mirkin, and D. Zelazo, "Asynchronous sampled-data synchronization with small communications delays," in *Proc. 63rd IEEE Conf. Decision and Control.*, 2024.

The research underlying this thesis was conducted honestly, in accordance with the ethical standards customary in academia. This applies, in particular, to the collection, processing and presentation of data, the description of and comparison with previous research work, etc. Also, the report on research activities and findings in this work is thorough and honest in accordance with the aforementioned standards.

Acknowledgements

First, I would like to thank both of my supervisors, Prof. Leonid Mirkin and Prof. Daniel Zelazo, for their guidance, encouragement, wisdom, and, most of all, patience. It is a rare thing to find even one supervisor so dedicated to a student's development as a researcher and to their well-being as a person, rather than to their own career. Somehow, when my dice were rolled, I got two such mentors. You shaped me as a researcher whether you like it or not, so I hope you are satisfied with the result.

My road to this point was not the typical one, and I would also like to thank several other people who, some unwittingly, paved my way here. To my high-school math teacher who told me "STEM is not for you, try something else" and put me in the lowest math level possible: you forced me to face the impostor syndrome far earlier than most. To Eyal Hershko, my high-school robotics teacher, who promptly ignored her judgment and ignited the passion for engineering in me. You showed me by example the difference between a teacher and a mentor, and how both are paramount for success. To Prof. Alon Wolf, who told me to stop wasting time on temporary jobs, took me by the hand, and brought me to the Technion pre-academic center. That simple conversation — which I doubt you even remember — made me stop stalling, face my fear of failure, and enroll. To my friends Leor Yamay and Yotam Kenneth, who are not only great friends but also took the time to help me weather the storm that is the first year at the Technion. To Ms. Daniela Avidan, whom I bothered incessantly with questions about missing proofs and generalizations after every ODE lecture — to the point she dragged me to the math department and suggested I enroll. I never did finish my degree in math, but your encouragement made me take more courses, learn new things, and, importantly, understand that the greatest hurdle in learning new things is being too afraid to try. To Prof. Rodolphe Sepulchre and Prof. Iman Shames, who invited me to their respective groups just to learn new things and perspectives.

And most of all, to my family. To my parents, who somehow supported and encouraged me no matter what, yet never pressured me when I was lost. I know you are proud of me, but I also hope you understand that your parenting — knowing when to give me space to make my own mistakes — was vital to my success. To my much, much smarter younger sister, Shahar, who talked to me as an equal even when I was learning high-school math and she was doing solid-state physics. I am also sorry for asking you to solve equations without real solutions when you were eight, without telling you. And of course, to my wonderful wife Maya. You are my rock, my comfort, my love. You brought our two amazing children into this world; you put your career on hold so I could pursue this dream. I have a truly marvelous proof of why you are perfect for me, but this page is too small to contain it. Everything was made possible because of you,

and nothing would have been worth it without you by my side.

Journey before destination.

The generous financial help of the Technion is gratefully acknowledged.

Contents

List of Figures

Abstract

This thesis investigates agreement and synchronization problems of continuous-time multi-agent systems (MASs) operating under sampled-data communication constraints. It challenges the prevailing reliance on diffusive consensus protocols, arguing that their structural simplicity can conceal fundamental performance limitations—particularly in the presence of external disturbances, measurement noise, and unstable dynamics.

The study begins with the consensus problem for integrator agents exchanging information intermittently at asynchronous sampling instants. A novel sampled-data protocol is proposed, emulating suitable global analog dynamics at each agent and transmitting sampled centroids of these emulators between neighbours. The closed-loop dynamics naturally decompose into centroid and disagreement components. The centroid subsystem evolves autonomously according to time-varying discrete consensus dynamics, independent of the sampling intervals. Under mild assumptions on persistency of connectivity and bounded sampling intervals, the system asymptotically converges to agreement. A simplified, scalable implementation is given for a specific choice of emulated topology.

The framework is then extended to general agreement of identical linear time-invariant (LTI) agents under the same connectivity assumptions, with both state-feedback and output-feedback designs accommodating practical sensing constraints. Conditions for exponential convergence are derived, and augmentations such as predictor-based updates for small delays and weighted update maps are proposed, preserving convergence guarantees while potentially offering improved performance.

A central theoretical contribution is the identification of a fundamental internal instability in diffusively coupled MASs when the agents satisfy certain constraints. For finite-dimensional agents, these manifest as a common unstable pole and direction canceled in the feedback loop, illustrating a unique multi-input multi-output (MIMO) phenomenon: pole cancellations due to deficient normal rank of the controller. This explains reported fragility to load disturbances and measurement noise, and proves that no purely diffusive feedback can stabilize such systems.

Motivated by this limitation, a 2 degrees-of-freedom (2DOF) distributed architecture is proposed, separating local stabilization from network-based coordination. This decouples plant dynamics from network filter design, enabling heterogeneous controllers, targeted disturbance rejection, and explicit noise shaping—capabilities absent

in standard consensus-like designs. An input–output analysis highlights these advantages and the architecture's ability to retain feedback in the consensus mode.

The analysis combines spectral graph theory, sampled-data control, and classical control tools, producing novel theoretical characterizations and practical design methods applicable for various LTI agents. The thesis concludes with open directions, including extensions of the emulation scheme to heterogeneous agents and the 2DOF architecture, systematic network filter design, and further studies of performance under communication and sensing constraints.

Notation and Abbreviations

$\mathbb{N}$	The set of natural numbers.
$\mathbb{N}_\nu$	The set $\{i \in \mathbb{Z} \mid 1 \leq i \leq \nu\}$.
$\mathbb{Z}$	The set of integers.
$\mathbb{Z}_+$	The set of positive integers.
$\mathbb{R}$	The field of real numbers.
$\mathrm{j}\mathbb{R}$	The set of pure imaginary numbers.
$\mathbb{C}$	The field of complex numbers.
$\mathbb{C}_\alpha$	The open complex right half-plane $\{s \in \mathbb{C} \mid \operatorname{Re} s > \alpha\}$.
$\bar{\mathbb{C}}_\alpha$	The closed complex right half-plane $\{s \in \mathbb{C} \mid \operatorname{Re} s \geq \alpha\}$.
$\mathfrak{R}(s)$	The real part of complex number s.
$\mathfrak{I}(s)$	The imaginary part of complex number s.
$\lvert s \rvert$	The modulus of complex number s.

---

e_i	The ith standard basis vector in a field $\mathbb{F}^n$.
I_ν	The $\nu \times \nu$ identity matrix.
$\operatorname{diag}\{A_i\}$	A block-diagonal matrix with diagonal elements A_i.
A'	The complex-conjugate transpose of a matrix A.
$\operatorname{Im} A$	The image (range) space of A.
$\ker A$	The kernel (null) space of A.
$\operatorname{spec}(A)$	The set of all eigenvalues of the matrix A.
$\operatorname{tr}(A)$	The trace of a square matrix A.
$\underline{\sigma}(A)$	The minimal singular value of A.
$\lVert A \rVert_\bullet$	The $\bullet$ norm of matrix (vector) A.
$A \otimes B$	The Kronecker product of matrices A and B.
$\mathbb{1}$	The all-ones vector.
$P_{\mathbb{1}}$	The orthogonal projection on $\operatorname{Im} \mathbb{1}$.
$\mathcal{G}$	A mathematical graph with node set $\mathcal{V}$ and edge set $\mathcal{E}$.
$A_{\mathcal{G}}$	The adjacency matrix of graph $\mathcal{G}$.
$D_{\mathcal{G}}$	The degree matrix of graph $\mathcal{G}$.
E	The incidence matrix of graph $\mathcal{G}$.
$L_{\mathcal{G}}$	The Laplacian matrix of graph $\mathcal{G}$.

$A^{\star}_{\mathcal{G}}$	The normalized adjacency matrix of graph $\mathcal{G}$.

$G(s)$	The transfer function of an LTI system G.
(A, B, C, D)	A state-space realization of a finite-dimensional LTI system G.
H_∞	Hardy space of holomorphic and bounded functions.
L_2	Lebesgue space of square integrable functions.
$\|\cdot\|_\infty$	The H_∞ norm of a system or L_∞ norm of a signal.
$\mathrm{pdir_i}(G, p)$	The input direction of a pole p of $G(s)$.
$\mathrm{pdir_o}(G, p)$	The output direction of a pole p of $G(s)$.
$\mathrm{zdir_i}(G, z)$	The input direction of a zero z of $G(s)$.
$\mathrm{zdir_o}(G, z)$	The output direction of a zero z of $G(s)$.

2DOF	2 degrees-of-freedom
A/D	analog to digital
BIBO	bounded-input bounded-output
D/A	digital to analog
LTI	linear time-invariant (A system described by linear time-invariant differential (difference) equations.)
MAS	multi-agent system (A collection of independent dynamical units, interacting to achieve a common goal.)
MIMO	multi-input multi-output (A dynamical system with multiple inputs and multiple outputs.)
SISO	single-input single-output (A dynamical system with a single input and a single output.)
ZOH	zero-order hold (zero-order hold)

Chapter 1

Introduction

> It's a dangerous business, Frodo, going out your door. You step onto the road, and if you don't keep your feet, there's no knowing where you might be swept off to.
>
> J.R.R Tolkien, *The Fellowship of the Ring* [1]

Many engineering problems can be understood through the lens of large collections of simpler, interacting subsystems. When these subsystems are coupled—for example, when spatially separated sensors and actuators must cooperate, or when otherwise independent units share a common objective—the overall plant is typically called *large-scale*. The analysis and control of such systems has been an active research area for more than five decades [2], [3].

A large-scale plant can always be represented as a standard MIMO system and, in principle, controlled with standard synthesis tools. However, the resulting controllers are almost always *centralized*: each control input depends on measurements from *all* outputs. Guaranteeing reliable, low-latency communication between every sensor–actuator pair becomes technically challenging and prohibitively expensive as the dimension grows [4]. This has driven sustained interest in *decentralized* designs whose controller matrices are block-diagonal, so that every subsystem is driven solely by its local measurements. Unfortunately, such severe sparsity constraints introduce their own difficulties. For example, even for two interacting systems with a quadratic cost optimal design is NP-complete [5], [6], and certain plants cannot be stabilized by block-diagonal controllers [7].

A pragmatic compromise is *distributed* control, in which the controller matrix is only *sparse* rather than block-diagonal, allowing each subsystem to exploit limited information from carefully chosen neighbors. This can be visualized by pixel diagrams, where the ith row corresponds to the ith control signal, u_i, and each column represents an agent's measurement, y_j. A diagonal structure (Figure 1.1(a)) demands no communication since each control signal uses only local measurements, $u_i = k_i(y_i)$ for

some operator k_i. A cyclic structure (Figure 1.1(b)) requires two links for each agent, $u_i = k_i(y_{i-1}, y_i, y_{i+1})$, and still more exotic patterns are also possible (Figure 1.1(c)).

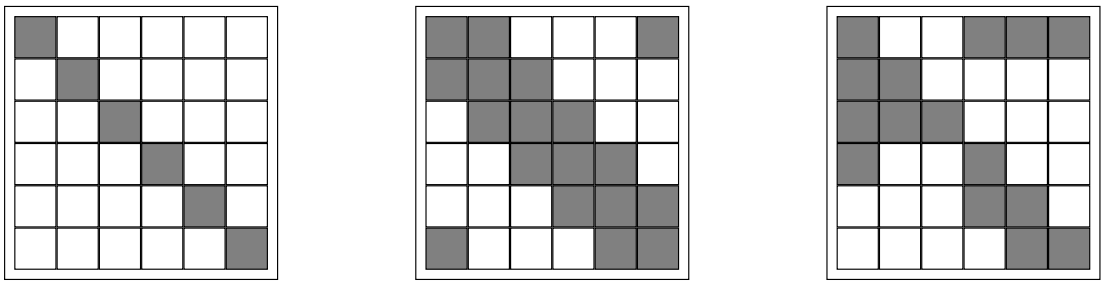

(a) Diagonal (decentralized) structure.

(b) Cycle structure.

(c) General sparse structure.

Figure 1.1: Illustration of distributed controller structures: a shaded ijth element represent the jth measurement is available for the ith control.

How these "information structures'' influence synthesis has been studied for decades [3]. For certain special cases—positive plants, for instance—optimal distributed controllers can even be computed in closed form [8]. In the generic LTI setting, however, convexity of the optimal synthesis problem is guaranteed only when plant and controller share a *quadratically invariant* sparsity pattern [9], [10]. Tractability deteriorates further when considering more realistic scenarios, for example when the information topology is time- or parameter- varying, or when the communication itself is intermittently sampled.

Among the many systems amenable to distributed control, MASs stand out as both a rich source of theoretical challenges and a key application domain. Here, a network of dynamical agents pursues a common, *cooperative* objective. Coordination tasks such as velocity agreement in vehicle platoons [11], formation keeping for multi-robot teams [12], [13], or distributed sensing and estimation [14], [15] all fit naturally in this framework. A fully decentralized controller is rarely adequate—cooperation demands some level of information sharing—yet a fully centralized architecture inherits the communication and implementation pitfalls described. The goal, therefore, is to devise distributed controllers that are as sparse as possible while still satisfying the cooperative specification and any temporal constraints.

1.1 Agreement problems in multi-agent systems

A prime example of cooperative objectives is reaching agreement between autonomous dynamical systems in the presence of communication constraints. It is a fundamental problem in numerous scientific disciplines such as opinion dynamics [16], distributed algorithms [17]–[19], physics [20], and control [21], [22].

$\boldsymbol{\mathcal{P}_1}$: Consider ν possibly heterogeneous agents, each with dynamics P_i, control inputs u_i, and outputs y_i. Design bounded control signals u_i such that the agents asymp-

totically reach *output agreement* in the sense that

$$\lim_{t\to\infty} \left|y_i(t) - y_j(t)\right| = 0 \qquad \forall\, i, j \in \mathbb{N}_\nu, \tag{1.1}$$

for all initial conditions.

If in addition $\lim_{t\to\infty} y_i(t) = \text{const}$, we say that the agents achieved *consensus*, and if the consensus value is 0 we say that the agreement is *trivial*, since this is akin to independent stabilization of the agents.

This would be an elementary problem if every agent had continuous access to the output of all other agents. However, this scenario is not practical in networked applications, where information exchange is costly. Hence, additional constraints are introduced. A natural assumption in large-scale networks is that each agent can interact only with a subset of the group often called its neighbors. Formally, this can be expressed as the constraint that the ith agent can communicate only with agents whose indices belong to a *neighborhood set* $\mathcal{N}_i \subset \mathbb{N}_\nu \setminus \{i\}$, which can be fixed or time-varying. We refer to such constraints as *spatial* constraints. Controllers satisfying prescribed spatial constraints are naturally *distributed*.

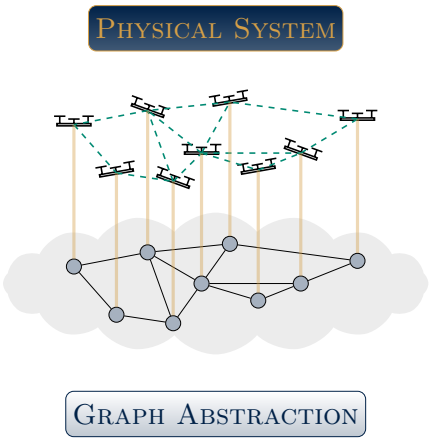

Figure 1.2: Modeling multi-agent systems via graphs: each agent is mapped to a node, and each communication channel to an edge.

These neighborhood sets and communication channels form a mathematical object called a graph. A simple graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ associates each agent with a node $v_i \in \mathcal{V}$ and each communication channel with an edge $e_{ij} = \{v_i, v_j\} \in \mathcal{E}$. When information exchange is bidirectional the graph is said to be undirected. The communication graph has substantial influence over the properties of the resulting closed-loop system, regardless of the particular controller used. This influence can be analyzed by exploiting the established a correspondence between graphs and certain matrix representations, in what is now called *spectral graph theory* [23], see a brief overview in Appendix A. This allows us to relate graph-theoretic properties such as connectivity, spanning trees, and cycles, to spectral properties of these matrices such as eigenvalues, positivity, and

invariant quantities. This is best illustrated in the following subsection, which considers the classical solution to the integrator consensus problem.

1.1.1 The consensus protocol

In its arguably simplest form, $\boldsymbol{\mathcal{P}_1}$ can be formulated in continuous time for a group of ν first-order autonomous agents described by

$$\dot{x}_i(t) = u_i(t), \quad x_i(0) = x_{0i}, \tag{1.2}$$

for all $i \in \mathbb{N}_\nu$, attempting to achieve consensus. One solution for the consensus problem for agents (1.2) which harnesses graph theory is the celebrated *consensus protocol* [21], [22], [24],

$$u_i(t) = \kappa \sum_{j \in \mathcal{N}_i} \Big(x_j(t) - x_i(t) \Big), \tag{1.3}$$

for some $\kappa > 0$. Note that only information within the ith neighborhood is used to generate this control, making it naturally distributed. The consensus protocol as shown in (1.3) exhibits the *relative sensing* nature of the protocol. Relative sensing appears naturally in MAS tasks, where absolute measurements are hard to obtain, such as space and aerial exploration and sensor localization, see [25]–[27] and the references therein. A different perspective of (1.3) is

$$u_i(t) = -\kappa \Big(|\mathcal{N}_i| x_i(t) - \sum_{j \in \mathcal{N}_i} x_j(t) \Big), \tag{1.4}$$

which emphasizes that agents exchange state-information by *communication*. This perspective is useful in problems where the agents are assumed to be communicating, and hence may exchange more complex, and possibly private, information.

Regardless of viewpoint, the analysis of the resulting closed-loop system is particularly elegant in the aggregate form. Assume that the spatial constraints describe a particular (undirected) graph. Then, using the notation from Appendix A, the aggregate version of the consensus protocol then reads

$$u(t) = -\kappa L_{\mathcal{G}} x(t), \text{ where } \left[\begin{array}{c:c} x(t) & u(t) \end{array}\right] \coloneqq \left[\begin{array}{c:c} x_1(t) & u_1(t) \\ \vdots & \vdots \\ x_\nu(t) & u_\nu(t) \end{array}\right],$$

and results in the collective closed-loop dynamics

$$\dot{x}(t) = -\kappa L_{\mathcal{G}} x(t), \quad x(0) = x_0, \tag{1.5}$$

with $L_{\mathcal{G}}$ being the *graph Laplacian*. Equation (1.5) is one of the main reasons for the renewed interest in graph theory within the control community. It clearly shows that, at least for integrators, the global dynamics are almost completely determined by the

underlying graph via $L_{\mathcal{G}}$. The benefit of this property is that we can tell quite a bit about the spectral properties of $L_{\mathcal{G}}$ without knowing the graph explicitly. In particular, we know that for any undirected graph, even disconnected, $\mathbb{1} \in \ker L_{\mathcal{G}}$. Since

$$x_i(t) = x_j(t), \quad \forall i, j \in \mathbb{N}_\nu \iff x(t) \in \operatorname{Im} \mathbb{1},$$

this implies that agreement is *always* a possible equilibrium of the consensus dynamics with undirected topologies. It turns out that for undirected graphs, dynamics (1.5) reach agreement if and only if $\mathcal{G}$ is connected. Moreover the convergence to agreement is exponential with rate determined by the second smallest eigenvalue of $\kappa L_{\mathcal{G}}$ and can be made arbitrarily fast by increasing the gain κ, see [22, Sec. 3.1]. Similar results, i.e., consensus iff connectivity, exist for both directed and time-varying graphs with suitable notions of connectivity, e.g. [21, Ch. 2].

1.1.2 General consensus-like protocols

Despite appearing in other fields for over 40 years, variations of (1.3) began appearing in the control literature only over the last two decades. In particular as part of several papers, published almost concurrently, which considered vehicle formation or heading problems [12], [28]–[31]. These papers which cite the simple discrete-time model from [20] as their inspiration, all considered variations of (1.3) acting on a group of identical agents attempting to drive their states into consensus without a common reference signal. In all cases, similar results, i.e., consensus iff connectivity, were derived. This successful marriage between graph and control theory led to a rapid surge of publications exploring the new subfield. For example, the Laplacian-like structure was shown to preserve passivity [32], and the graph and cycle structure were connected to the H_2 norm [33] of agreeing systems. In [34] a consensus model with an input was considered, and controllability properties were related to symmetries of the graph.

As the field evolved, variations of $\boldsymbol{\mathcal{P}_1}$ retained their significance. It turns out that agreement is a necessary building block for more sophisticated objectives, ranging from leaderless formation control [35] to coordinated wildfire monitoring [21]. Hence, we may consider a more general consensus-like protocol of the form

$$u_i = k_i \sum_{j \in \mathcal{N}_i} (y_j - y_i), \qquad i \in \mathbb{N}_\nu, \tag{1.6}$$

for some measured signals $y_i(t)$ and possibly dynamic controllers k_i. Commonly controllers k_i are chosen to be identical, since this facilitates the overall design. In particular, this reduces the affects of the network to a perturbation of the gain by the eigenvalues of the Laplacian, reducing agreement to a simultaneous stabilization problem [12], [36]. This property was also exploited to analyze the robustness of consensus-like protocols to uniform time-delays [37], [38] and sampling [29]. However, even when identical, the gains are not necessarily square, static, or time-invariant. For example in [39]

a group of second-order integrators was driven to consensus by using a variant of a PD controller, or equivalently full state-feedback. Others solved the consensus problem for general LTI agents with output measurements through a Riccati-based low-gain approach [36]. Time-varying controllers were also considered, mostly in the context of time-varying topologies [31] and noisy measurements [40], [41]. Over time, the assumption that the agents can communicate controller states became more widespread. This lead to several observer-based consensus-like controllers [42], [43], providing a methodical design process for homogeneous LTI agents.

Simultaneously, attention was directed toward two other important variations: i) agreement of heterogeneous agents, and ii) controlling the agreement trajectory. The former arises naturally in practical applications, since agents are often uncertain and thus not exactly homogeneous. The latter originated from the fact the aforementioned controllers could only drive agents to agreement trajectories generated by the homogeneous uncontrolled model, cf. [43, Thm. 2]. It can be formulated as solving $\boldsymbol{\mathcal{P}_1}$, while ensuring that

$$\lim_{t\to\infty} \|y_i(t) - r(t)\| = 0, \quad \forall i \in \mathbb{N}_\nu, \tag{1.7}$$

where $r(t) \in \mathcal{R}$ for some prescribed family of trajectories $\mathcal{R}$. This cumulated in the seminal result that a necessary condition for agreement is the existence of a common *internal model* between all the agents [44]. Combined with the prevalence of the communication viewpoint, this result paved the way to cooperative output regulation approaches. In these methods, each agent embeds a *generator* into their controller which can generate the required family of trajectories. For example, the set $\mathcal{R}$ is often generated by an LTI generator with uncontrolled dynamics

$$\begin{cases} \dot{r} = A_0 r \quad r(0) = r_0 \\ \eta = Rr \end{cases},$$

such that $\mathrm{spec}(A_0)$ are all in the closed left half-plane. The agents then locally track their generators, while cooperativly exchanging the generator's output and agreeing on them. This requires each agent to have a local controller based on the regulator equations [45, Ch. 1], which may result in the agents having different dynamic controllers, k_i [44, Thm. 5]. This result was later extended to non-linear agents, generators, and controllers [46].

Yet despite increasing generality, at the heart of all of these variants sits consensus protocol (1.6), or more generally controllers of the form

$$u_i = \sum_{j\in\mathcal{N}_i} k_{ij}(y_j - y_i), \qquad i \in \mathbb{N}_\nu, \tag{1.8}$$

for possibly dynamic controllers k_{ij}. Control laws which are based on a linear combination of some function of the relative outputs of all neighbors as in (1.8) are generally

called *diffusive*, and the closed-loop systems are called *diffusively coupled*. It seems that over the years this diffusive structure became almost an a priori assumption in control design of MASs, regardless of the actual problem. Even when no actual sensing is involved, for example [44], [47], the incoming information is processed via a consensus-like protocol using relative measurements.

1.2 Structure and goals of this thesis

While such results highlight the elegance and utility of the consensus framework, practical deployments reveal important limitations under realistic conditions. Sampling or communication delays, measurement noise, and exogenous disturbances all interact intricately with the network structure. This interaction makes tuning rules opaque and reduces robustness margins. Even worse, robustness can be hard to define and quantify due to the coupling via the network structure. Consider, for example, a group of integrator agents attempting to reach consensus with sampled-data interaction. Although this is perhaps the simplest variant of $\boldsymbol{\mathcal{P}_1}$, adhering to the structure of protocol (1.6) requires much deeper analysis. Even synchronous periodic sampling requires the gains to be tuned with respect to both the sampling period h and the eigenvalues of the Laplacian [29]. If the sampling is intermittent, event-triggered, or asynchronous, more conservative gains are needed [48]. Similarly, uniform transmission delays – where each pair $y_i - y_j$ are delayed by the same amount – also constrain the allowable controller gains [29], [38], [49]. Because the graph, and thus Laplacian eigenvalues, is generally unknown, satisfying even basic constraints necessitates adopting a conservative design approach. This is a stark contrast to traditional control problems involving integrators, where such constraints often have analytical solutions with guaranteed stability margins.

Such limitation under non nominal conditions are not confined to information constraints like delays and sampling. It is well documented that diffusively-coupled systems behave poorly when affected by external signals. Measurement noise significantly degrades performance [33, §III-A], while disturbance and uncertainties can hardly be attenuated even by dynamic [43], [50] or non-linear [51] controllers. To cope with the difficulties, various relaxed assumptions are adopted. Some allow the controllers to be time-varying [40], [52] or non-linear [53], and in both cases some boundedness or convergence assumptions must be made on the external signals. Moreover, if these assumptions fail, the resulting trajectories exhibit certain common traits that can be associated with instability. It is notable that despite being well documented, these failings are rarely investigated for themselves. This alarming fragility and instability were not properly explained, and different architectures are hardly explored.

These oddities served as the main motivation for the research detailed in this dissertation, which advocates we take a more control-oriented perspective on MASs control problems as a whole, and the agreement problem in particular. Historically, the suc-

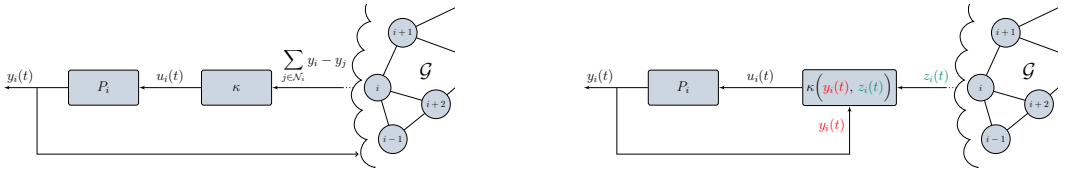

(a) Only diffusive measurement.

(b) Independent network (z_i) and local (y_i) measurements.

Figure 1.3: Distributed control paradigms for agents P_i and controller κ: a singe uniform measurement, or separation to local and networked measurements.

cess of control theory lies in match-making the right controllers to the right problems. This can be seen in matching 2DOF control to classical servo-regulation [54], [55], predictors to dead-time systems [56], [57], or generalized sample-and-hold functions for sampled-data systems [58]. A common theme in all of these examples is that they exploit something in the structure of the problem to either simplify the design or improve the performance. This is most evident in 2DOF controllers, where the design exploits the fact that there are two distinct measurement channels, one for the output and one for the reference, to expedite the design. In sampled-data control the key was to not assume a priori an ideal sampler and a zero-order hold (ZOH), and instead to consider both analog to digital (A/D) and digital to analog (D/A) devices as free design components.

The main focus of this dissertation is to bring similar insights into the realm of MASs and agreement problems. First, much like in 2DOF, we note that (1.6) can be rewritten in communication form much like the vanilla (1.4),

$$u_i = -|\mathcal{N}_i|\, k_i y_i + k_i \sum_{j\in\mathcal{N}_i} y_j \coloneqq -|\mathcal{N}_i|\, k_i y_i + k_i z_i, \qquad i \in \mathbb{N}_\nu.$$

This results in at least two different signals, the local output y_i and the network signal z_i. This is not a mere algebraic manipulation; often these signals are subject to different constraints. For example, it is reasonable to expect that network-sourced information will be sampled less frequently or will be more heavily affected by noise. This alone makes problems such as "agreement of sampled-data MASs" much more nuanced. Second, we may want to process these two measurements differently to try and exploit the interplay between their differing constraints. We argue that adhering rigidly to the diffusive consensus structure may inadvertently introduce conservatism and complexity in the design. In fact, as we shall show, when additional constraints appear, even modest departures from the conventional form can yield both performance gains and dramatic simplification of the design process. Unlocking these gains, however, demands a deeper look at the interplay between communication topology, sampling, delays, and plant dynamics. The difference between the two approaches is illustrated in Figure 1.3.

The research in this thesis focuses on explicitly considering control structures such

as Figure 1.3(b), rather than the diffusive structures in Figure 1.3(a), as well as highlighting the shortcomings of the latter. Chapter 2 begins with deeper inspection of the classic consensus protocol, and how naïvely discretizing (1.3), i.e., using an ideal sampler and a ZOH, requires conservative and small gains. The rest of this chapter is dedicated to an alternative architecture, published in [59], which eliminates this conservatism by incorporating generalized sampled-and-hold functions inspired by optimal sampled-data control. This is made possible by viewing the consensus protocol through the lens of Figure 1.3(b), and separating the intermittent network measurement z_i from the fast and reliable y_i. Notably, the proposed architecture guarantees consensus under asynchronous and intermittent communication, even for possibly time-varying directed graphs. The only requirement is that a weak connectivity assumption holds. Specifically, it requires only a necessary assumption for discrete-time integrators under switching graphs, and uniformly bounded sampling intervals.

Building on the sampled-data architecture from Chapter 2, Chapter 3 extends the discussion to general LTI agents that must *synchronise* to a time-varying reference. Motivated by structural insights from the integrator case, we propose a new state-feedback control law to emulate a different control law instead of the standard Laplacian-based one. We then show that when combined with the update scheme from the integrator case, synchronization is achieved under the same communication constraints. We then extend the results to controllers using output measurements only in two different ways: i) by directly emulating an observer-based control law, or ii) by including a simplified observer into the original scheme. Both for state and output measurements we provide a scalable and low-order implementation. The state-feedback case was presented (without some of the proofs) in [60], while the second output-feedback version was published in [61].

Chapter 4 concludes our investigation of sampled-data agreement problems, by deriving several smaller results most of which are unpublished. Specifically, we investigate the convergence rate of the emulation scheme, as well as provide graph-theoretical conditions to ensure exponential convergence. We also consider two augmented update maps. The first variation is a predictor capable of guaranteeing agreement under heterogeneous and time-varying small delays, a work published in [62]. The second demonstrates that adding weights to the original update scheme does not alter the convergence result. Such weights are commonly used in the literature in various ways to improve performance.

In Chapter 5 we deviate from the linear progression of increasingly complicated sampled-data agreement problems, and explicitly consider Figure 1.3(a) under diffusive controllers (1.8). This deviation is motivated by the odd, unstable looking, disturbance response of the consensus protocol as detailed in Section 5.1. This prompts a system-theoretic analysis of generic diffusive couplings, in an attempt to understand this behavior, as well as others exhibited in the literature. We prove that the diffusive structure *fundamentally* cannot internally stabilize certain classes of agents. This clari-

fies the disturbance amplification and instability reported in the literature, and further supports our focus on the architecture in Figure 1.3(b). For finite-dimensional agents this result has an intuitive interpretation as *cancellations* of common agent poles by the diffusive structure. This work was published in [63].

Motivated by the results of Chapter 5, Chapter 6 begins with revisiting the classical consensus protocol in an input-output framework. By deriving known results from this perspective it is possible to better understand the underlying structure and limitations of consensus-like protocols. By drawing parallels to servo-regulation problems, we develop an alternative architecture to diffusive coupling. Namely, we propose a 2DOF architecture that fuses classical servo-regulation ideas with distributed control requirements. It requires more careful tuning to ensure agreement, as opposed to classic consensus, but has demonstrably improved robustness and disturbance rejection.

We conclude our work in Chapter 7, where we summarize our findings and insights, and provide a few intriguing directions for future research. To keep the main text readable, three appendices are provided. Appendix A provides and introduction to graph theory, relevant notations, and some required standard results, while auxiliary results on the Kronecker product and matrix theory are collected in Appendix B. It is generally assumed that the reader is familiar with continuous-time LTI dynamical systems, yet some background and relevant results appear in Appendix C. The reader is referred to standard monographs for deeper coverage.

Chapter 2

On Sampled-Data Consensus: Divide and Concur

> "Begin at the beginning," the King said, gravely, "and go on till you come to the end; then stop."
>
> Lewis Carrol, *Alice in Wonderland* [64]

In this chapter, we take our first steps into the world of constrained agreement problems. It begins with a quick recap of the consensus version of $\boldsymbol{\mathcal{P}_1}$ for integrators controlled by consensus protocol (1.3). We first formalize the results alluded to in Subsection 1.1.1, and then show the differences when naïvely discretizing (1.3). This motivates our alternative sampled-data architecture, which is the main topic of this chapter.

2.1 Between discretization and sampled-data

Consider the integrator consensus problem as defined in Section 1.1, with controller (1.3). That is, a group of ν first-order autonomous agents described by

$$\dot{x}_i(t) = u_i(t), \quad x_i(0) = x_{0i},$$

with the aggregate closed-loop dynamics

$$\dot{x}(t) = -\kappa L_{\mathcal{G}} x(t), \quad x(0) = x_0$$

as defined in (1.5), where $L_{\mathcal{G}}$ is the graph Laplaican of the underlying communication topology. The first property of these dynamics we wish to introduce is the invariance of the centroid, as stated below.

Proposition 2.1.1 (Proposition 3.13 in [22])**.** *Let $L^i_{\mathcal{G}}$ be the in-degree Laplacian of $\mathcal{G}$,*

with left eigenvector $q'L^i_{\mathcal{G}} = 0$. *The quantity* $q'x(t)$ *is an invariant of dynamics* (1.5). *In particular, the* centroid, $\bar{x}(t) = (1/\nu)\mathbb{1}'x(t)$, *is an invariant of dynamics* (1.5) *for all undirected graphs.*

As mentioned in Chapter 1, for undirected graphs, dynamics (1.5) reach agreement if and only if $\mathcal{G}$ is connected. This is a well known result even for directed graphs, with various proofs available in standard textbooks, e.g. [22, Thm. 3.12] or [21, Thm. 2.8]. Despite being well known, it may be instructive to understand the mechanism behind them. To this end, we state and prove below the simplest case in a slightly non-orthodox way to provide some intuition required for the rest of this dissertation.

Proposition 2.1.2. *Consider* ν *identical integrator agents controlled by* (1.3) *for some scalar* $\kappa > 0$ *and an undirected graph* $\mathcal{G}$. *The following statements are equivalent.*

1. $\mathcal{G}$ *is connected.*

2. *The agents reach average consensus for all initial conditions.*

Proof Consider a representation of the aggregate system in the frequency domain, where the closed-loop can be written as

$$x = \left(I_\nu + \frac{\kappa}{s}L_{\mathcal{G}}\right)^{-1} x_0$$

with x_0 is the aggregate initial condition response. Since $\mathcal{G}$ is undirected, $L_{\mathcal{G}}$ is positive semi-definite and there exists a unitary transformation U such that $UL_{\mathcal{G}}U' = \operatorname{diag}\{\lambda_i\}$, which allows us to rewrite the closed-loop dynamics as

$$x = U' \operatorname{diag}\{S_i\} U x_0$$

with

$$S_i(s) \coloneqq \frac{1}{1 + \frac{\kappa\lambda_i}{s}}.$$

Moreover, since $L_{\mathcal{G}}\mathbb{1} = 0$, we can pick a unitary U such that $U'e_1 = (1/\sqrt{\nu})\mathbb{1}$ and $e_1'U = (1/\sqrt{\nu})\mathbb{1}'$. Note that we can rewrite the above dynamics as

$$x = \sum_{i=2}^{\nu} S_i(s) U' e_i e_i' U x_0 + S_1(s)\frac{1}{\nu}\mathbb{1}\mathbb{1}'x_0$$

and that $S_1(s) \equiv 1$.

$\Longrightarrow$ Since $\mathcal{G}$ is undirected and connected, by Proposition A.0.1 $\lambda_1 = 0$ and $\lambda_i > 0$ for all $i > 1$. This implies that for all $i > 1$

$$S_i(s) = \frac{s}{s + \kappa\lambda_i}$$

is stable, thus by the final value theorem

$$\lim_{t\to\infty} x(t) = \frac{1}{\nu}\mathbb{1}\mathbb{1}'x_0 = \bar{x}_0\mathbb{1}.$$

$\Longleftarrow$ Since the agents reach average consensus for all initial conditions and the centroid is invariant, we have

$$\lim_{t\to\infty} (x(t) - \mathbb{1}\bar{x}(t)) = \lim_{t\to\infty} (I_\nu - P_{\mathbb{1}})\, x(t) = 0$$

where $P_{\mathbb{1}} \coloneqq (1/\nu)\mathbb{1}\mathbb{1}'$ is the orthogonal projection on $\operatorname{Im}\mathbb{1}$ [65, Thm. 7.5]. From the previous decomposition we have

$$(I_\nu - P_{\mathbb{1}})\, x = \sum_{i=2}^{\nu} S_i(s) U' e_i e_i' U x_0,$$

therefore $S_i(s)$ must be stable for all $i > 1$ and any $\kappa > 0$, which requires that $\lambda_i > 0$ for all $i > 1$. Thus $L_{\mathcal{G}}$ has a simple eigenvalue at 0 and the rest strictly positive, implying that $\mathcal{G}$ is connected. ■

Before moving on to the sampled-data case, there are two things that are important to notice for the sequel. First, for undirected graphs the centroid is an invariant quantity. It is unaffected by the dynamics and predetermines the final value regardless of the gain κ. Second, the convergence rate directly depends on κ and the graph structure. This is readily seen in the frequency domain, since the closed-loop poles are determined by

$$S_i(s) = \frac{s}{s + \kappa\lambda_i}$$

with poles $p_i = -\kappa\lambda_i$. Hence, the most dominant mode (excluding the agreement direction) is $-\kappa\lambda_2$, which can be made arbitrarily fast without risking the overall stability of the system. This may not be the case in discrete-time or considering more complex dynamics.

To this point we only discussed the problem in continuous time, where the spatial constraints limited the subset of neighbors each agent can communicate with. However, networked implementations may impose additional limitations on the information exchange between agents, related to the *time instances* at which communication is possible. Specifically, we assume that agents can convey information to neighbors only at time instances $t = s_k$, $k \in \mathbb{Z}_+$, for a strictly increasing monotone sequence of *sampling instances* $\{s_k\}_k$. If all agents transmit their states simultaneously, at each s_k, sampling is said to be *synchronous*. If the ith agent at each s_k receives information from only a subset of its neighbors, sampling is referred to as *asynchronous*. These concepts are illustrated in Figure 2.1, where each time axis denotes an agent, and each impulse a sampling instance. Note that even in the asynchronous case of Figure 2.1(c), there is only one sampling sequence that encompasses all agents. For example, even

thought agent 3 does not transmit at $t = s_1$, agent 1 does, hence $t = s_1 \in \{s_k\}_k$. This formalism can naturally accommodate changing interconnection topologies. There is

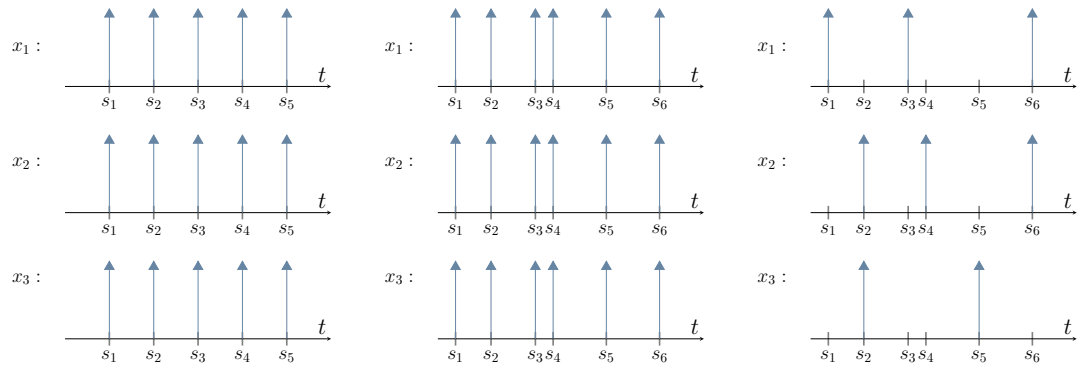

(a) Periodic and synchronous. (b) Aperiodic and synchronous. (c) Asynchronous.

Figure 2.1: Different sampling patterns for $\nu = 3$ agents: each time axis corresponds to an agent, and an impulse on the ith axis indicates the corresponding agent transmits information.

rich literature on sampled-data consensus, studying problems under synchronous and asynchronous sampling, time- and event-triggered sampling mechanisms, see [48] and the references therein. Still, the vast majority of available approaches assumes that the control signal is piecewise constant, unchanged between updates from neighbors, i.e., uses the zero-order hold as the D/A converter. This assumption facilitates the reduction of the problem to a pure discrete agreement, at least under synchronous sampling. Indeed, in the latter case the discretized aggregate system is

$$x[k+1] = x[k] + h[k]u[k],$$

where $x[k] \coloneqq x(s_k)$ and $h[k] \coloneqq s_{k+1} - s_k$ (sampling interval). Substituting in the discrete version of the consensus protocol,

$$u[k] = -\kappa L_{\mathcal{G}} x[k],$$

leads to the closed-loop dynamics

$$x[k+1] = (I - \kappa h[k] L_{\mathcal{G}}) x[k], \tag{2.1}$$

which can be thought of as the Euler approximation of (1.5).

Like in the continuous-time case, system (2.1) still has its only equilibrium at an agreement set. Yet the stability required by Proposition 2.1.2 is now guaranteed only if $\kappa h[k]$ is sufficiently small for all possible sampling intervals $h[k]$ [24]. This compromises the convergence rate for networks with large variability in sampling intervals. There are other agreement-reaching protocols for discrete systems [24], but they all

require sufficiently small, and normally conservative, gains to guarantee stability. For asynchronous sampling, the situation is even more dire, as illustrated in the following example.

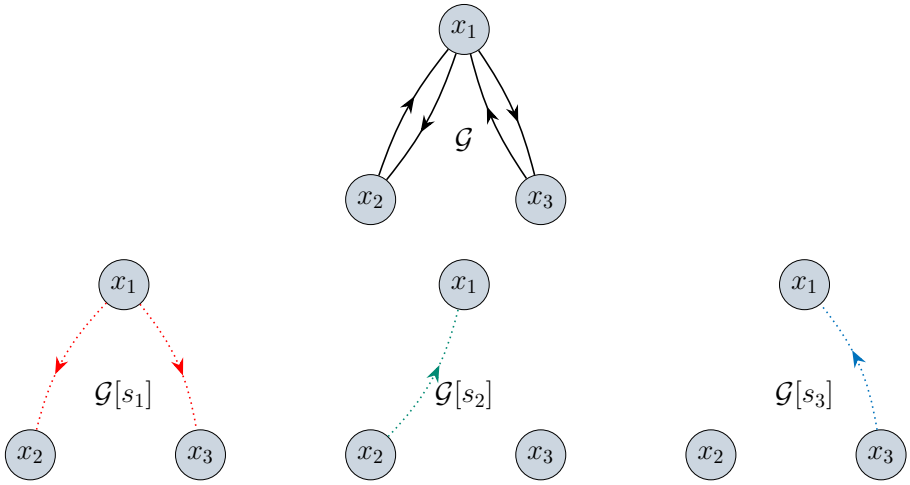

Figure 2.2: Top: original graph. Bottom: possible subgraphs induced by the asynchronous sampling sequence.

Example 2.1.3. Consider a simple case of $\nu = 3$ agents, interacting over a graph with edges $\mathcal{E} = \{(1,2),(2,1),(1,3),(3,1)\}$. Note that this graph is undirected, hence the in and out Laplacians are equal and given by

$$L_{\mathcal{G}} = \begin{bmatrix} 2 & -1 & -1 \\ -1 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix}.$$

Now assume that the agents are sampled in an asynchronous way, such that agent 2 transmits Δ_2 seconds after agent 1, and agent 3 transmits Δ_3 seconds after agent 2. Moreover, assume without loss of generality that $\Delta_3 > \Delta_2$, and that agent 1 does not transmit again before agent 3 does. Consequently, the first three sampling instances are

$$s_1 = s_1, \quad s_2 = s_1 + \Delta_2, \quad s_3 = s_1 + \Delta_3,$$

and each sampling instance induces a different *directed subgraph* of $\mathcal{G}$ based on the transmitting agent. This is illustrated in Figure 2.2, and results with the following in-degree Laplacians

$$L_{\mathcal{G}}[s_1] = \begin{bmatrix} 0 & 0 & 0 \\ -1 & 1 & 0 \\ -1 & 0 & 1 \end{bmatrix}, \quad L_{\mathcal{G}}[s_2] = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad \text{and} \quad L_{\mathcal{G}}[s_3] = \begin{bmatrix} 1 & 0 & -1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

all of which are directed. ▽

As shown in Example 2.1.3, the asynchrony naturally turns $L_{\mathcal{G}}$ into a time-varying matrix, which often requires even smaller gains and more complex technical machinery [66].

Remark (discrete vs. sampled-data) 2.1. The simple case considered above exemplifies well a key difference between pure discrete-time and sampled-data networked setups. The latter typically has shift-varying discretized models, as realistic network sampling is intermittent. Moreover, parameters of those models are uncertain, unless we somehow know the next sampling instance. ▽

We argue for parting with the use of the zero-order hold, synchronized with measurement updates, in sampled-data consensus protocols. This custom does not appear to be justified by implementation requirements nowadays and might be an *atavism*, survived from the early days of computer-controlled systems. We aim at exploiting opportunities offered by the use of more sophisticated hold mechanisms to reach agreement over agents (1.2) under both spatial (distributed) and temporal (sampling) constraints.

The design of control waveforms, or D/A converters, is not new to sampled-data control per se. In the centralized setting, optimal converters for periodic [58], [67], [68] and intermittent [69] sampling rates can be designed. Similar control waveform generators are also used in the networked control context in [70], [71], also without spatial constraints. Their common property is that they in effect *emulate* the control signal of the desired analog closed-loop control in an open-loop way. This property was observed in [68, Sec. 6] and conjectured as a guiding principle for situations where direct performance-justified design of the hold function is not available.

We follow that logic to put forward a sampled-data consensus protocol based on emulating (1.4). The challenge is that such an emulation has to be carried out in a distributed manner, locally at each agent. We work it out by implementing a model of the *whole world*, viz. (1.5), at each agent in continuous time. Each local controller uses then the states of these emulators to mimic the continuous-time consensus protocol (1.4). Information about real states of neighboring agents is used to adjust local emulators each time this information becomes available at sampling instances. We propose sample and update protocols, based on the *centroids* of the corresponding emulators of each agent. This idea is illustrated in Figure 2.3, where each agent transmits the local centroid at $t = s_k$ to its corresponding neighbors.

The choice of transmitting the centroids is not unique, but provides certain distinct advantages. Namely, with such a choice, the resulting closed-loop dynamics can be divided into those of the decoupled centroids, which behave as an autonomous system, and those of local disagreements, which evolves according to the analog consensus dynamics driven by the emulator centroids. Remarkably, the dynamics of the centroids are also independent of the sampling intervals, i.e., are certain, unlike (2.1). As such, the analysis is simplified and global asymptotic convergence to agreement is proved under mild connectivity assumptions on the sampled-data topology and with no need

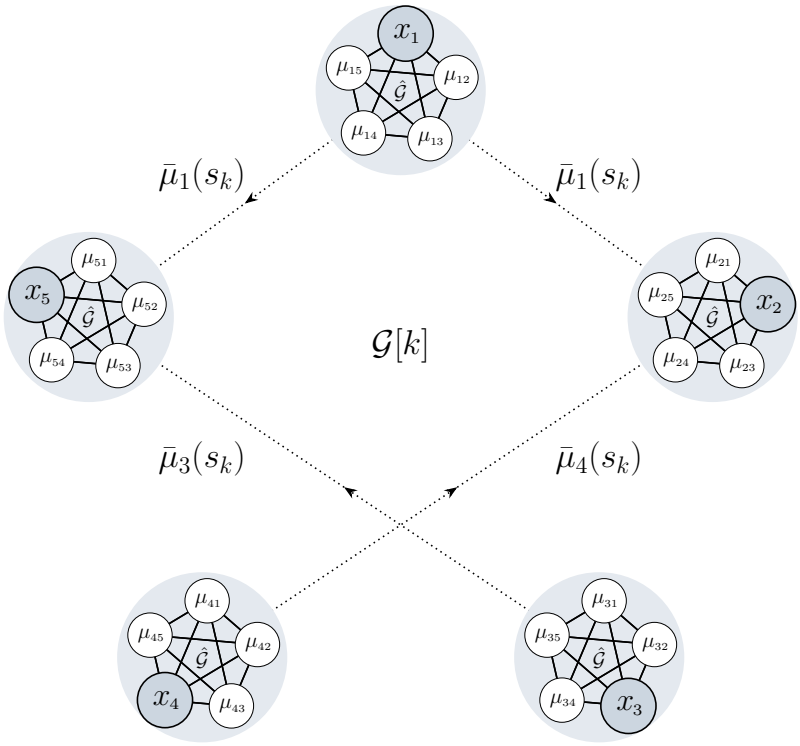

Figure 2.3: An illustration of the emulation architecture: each agent locally emulates all of the other continuously, and transmits their centroid when possible.

for *a priori* knowledge of sampling instances and specific bounds on sampling intervals. Another noteworthy property of the proposed architecture is that the spatial topology emulated locally need not match that of the actual network. This is demonstrated in Figure 2.3, where the emulated topology is the complete graph, while $\mathcal{G}[k]$ is not. This property can be further exploited to substantially reduce computational complexity, rendering it independent of the number of agents.

2.2 An emulation approach

As discussed, we aim at reaching agreement for the group of ν agents described by (1.2) under given spatial and temporal communication constraints. These constraints are determined by a strict monotonically increasing sequence of sampling instances $\{s_k\}_k$ and associated sets $\mathcal{N}_i[k] \subset \mathbb{N}_\nu \backslash \{i\}$ of neighbors of the ith agent that convey information about their states at each s_k. These information updates are intrinsically directional and not necessarily symmetric, i.e., $j \in \mathcal{N}_i[k]$ might not imply that $i \in \mathcal{N}_j[k]$. Each set $\mathcal{N}_i[k]$ can thus be associated with a directed graph, say $\mathcal{G}[k]$, having the Laplacian $L_{\mathcal{G}}[k] \in \mathbb{R}^{\nu \times \nu}$.

To ensure persistent connectivity in the whole scheme, we assume hereafter that

$\boldsymbol{\mathcal{A}_1}$: there is a strictly increasing sub-sequence of sampling indices $\{k_n\}_n$ such that for all $n \in \mathbb{Z}_+$

(i) the intervals $s_{k_{n+1}} - s_{k_n}$ are uniformly bounded, and

(ii) the union graph $\bigcup_{k=k_n}^{k_{n+1}-1} \mathcal{G}[k]$ contains a directed rooted tree.

Assumption $\boldsymbol{\mathcal{A}_1}$ is commonly employed in works related to coordination protocols over switching or time-varying graphs [21], [28], [31]. This assumption ensures that information propagates throughout the entire network persistently across bounded sampling intervals, leaving no nodes forever detached from the rest of the network. This is illustrated in Figure 2.4, where neither of the three graphs contain a directed rooted tree, yet their union does.

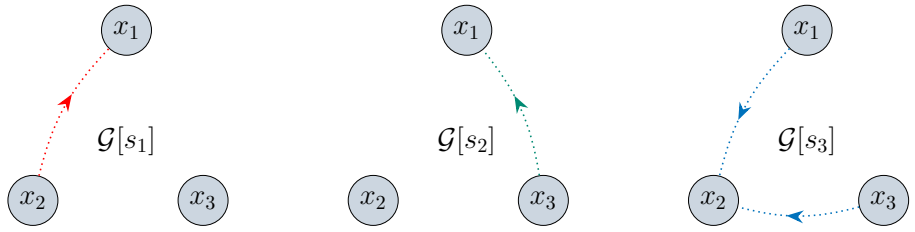

Figure 2.4: A sequence of graphs satisfying $\boldsymbol{\mathcal{A}_1}$: neither graph contains a directed rooted tree, yet their union does.

Remark (variants of $\boldsymbol{\mathcal{A}_1}$) 2.2. Some readers might be more familiar with variants of $\boldsymbol{\mathcal{A}_1}$ requiring a globally reachable node rather then a directed rooted tree. This is because of the inherent ambiguity when defining directed graph Laplacians. This goes beyond the obvious confusion regarding in or out degree Laplacians. For example, in some popular books (e.g. [22], [72]) $A_{\mathcal{G}}$ is defined as the transpose of (A.1), which changes the left and right kernels of the directed Laplacians. Unless otherwise specified, when talking about the directed graph Laplacian in this thesis, we refer to the *in-degree* Laplacian as defined in Appendix A. This choice allows us to naturally consider an arrow "going in to" agent i as that agent *receiving* information. For example, in Figure 2.4 in $\mathcal{G}[s_1]$ agent 2 transmitted information to agent 1. This imposes no loss of generality since the same results would hold, under slightly modified assumptions, for both in and out Laplacians defined elsewhere in the literature. ▽

Following the discussion in Section 2.1, the proposed architecture is based on emulating the "ideal" analog consensus behavior à la (1.5) at each agent. To this end, associate with each agent an emulated spatial communication topology, which is represented by a graph $\hat{\mathcal{G}}$ with Laplacian $L_{\hat{\mathcal{G}}} \in \mathbb{R}^{\nu \times \nu}$, and the corresponding set of neighbors $\hat{\mathcal{N}}_i \subset \mathbb{N}_\nu \setminus \{i\}$ of the ith agent. We emphasize that this $\hat{\mathcal{G}}$ need not match the actual spatio-temporal topology represented by $\mathcal{G}[k]$. All we require is that

$\boldsymbol{\mathcal{A}_2}$: the graph $\hat{\mathcal{G}}$ is undirected and connected,

which implies that $L_{\hat{\mathcal{G}}} = L'_{\hat{\mathcal{G}}}$ and its eigenvalue at the origin is single.

Remark (undirected vs. directed emulated graphs) 2.3. In fact, the emulator topology may also be directed. We require primarily the property that $L_{\hat{\mathcal{G}}}$ has a single eigenvalue at the origin with corresponding left eigenvector $\mathbb{1}$. This can be also achieved, for example, by weakly connected and balanced directed graphs. ▽

By the emulator at the ith agent we then understand the function $\mu_i(t) \in \mathbb{R}^\nu$, whose elements $\mu_{ij} \coloneqq [\mu_i]_j$ satisfy

$$\begin{aligned} \dot{\mu}_{ij}(t) &= -\kappa \sum_{l \in \hat{\mathcal{N}}_j} \big(\mu_{ij}(t) - \mu_{il}(t)\big), \quad \forall j \in \mathbb{N}_\nu \setminus \{i\} \\ \mu_{ii}(t) &= x_i(t), \end{aligned} \tag{2.2}$$

for some given initial conditions $\mu_{ij}(0)$. It is readily seen that (2.2) at each $j \neq i$ matches the jth row of (1.5). The purpose of (2.2) is to emulate the analog consensus protocol (1.4) at the ith agent by replacing the actual (remote) neighboring states x_j by their local clones μ_{ij}, i.e., as

$$u_i(t) = -\kappa\Big(|\hat{\mathcal{N}}_i|\mu_{ii}(t) - \sum_{j \in \hat{\mathcal{N}}_i} \mu_{ij}(t)\Big). \tag{2.3}$$

Taking into account that $\dot{\mu}_{ii} = u$, by (1.2), we end up with

$$\dot{\mu}_i(t) = -\kappa L_{\hat{\mathcal{G}}} \mu_i \tag{2.4}$$

as the collective dynamics of the ith agent.

It should be clear that the control law (2.2)–(2.3) is incomplete, for the resulted dynamics (2.4) are autonomous, not synchronized with other agents. We thus need to complement it by a synchronization mechanism satisfying given spatio-temporal constraints.

To this end, two aspects are to be decided:

(i) what information agents should broadcast about their own states, and

(ii) how emulators should utilize the conveyed information.

One can think of a number of possible approaches here, even if each agent may broadcast only a scalar signal. For instance, agents may broadcast their own sampled states, $x_j(s_k)$, or a function of the state of their complete emulator, $\phi(\mu_j(s_k))$. The receiving-side emulators may then update only components corresponding to the received updates, e.g., as $\mu_{ij}(s_k^+) = x_j(s_k)$, or all their states simultaneously.

Our choice is to broadcast the *centroid* of the corresponding emulator,

$$\bar{\mu}_j(t) \coloneqq \frac{1}{\nu}\mathbb{1}'\mu_j(t) = \frac{1}{\nu}\sum_{j=1}^{\nu} \mu_{ij}(t),$$

and is motivated by two insights. First, by Proposition 2.1.1 we know that the centroid is an invariant quantity of Laplacian dynamics, meaning it is not affected by the emulator "flow" dynamics (2.4). For the second insight, consider the following example.

Example 2.2.1. Consider dynamics (1.4), and assume that $\mathcal{G}$ is *the complete graph*. For the complete graph each agent is connected to all of the others, hence

$$L_{\mathcal{G}} = \nu(I_\nu - P_{\mathbb{1}}) \quad \text{and} \quad \mathcal{N}_i = \mathbb{N}_\nu \setminus \{i\} \quad \forall i \in \mathbb{N}_\nu,$$

where $P_{\mathbb{1}} = (1/\nu)\mathbb{1}\mathbb{1}'$ is the orthogonal projection on $\operatorname{Im}\mathbb{1}$. For this scenario the consensus protocol can be rewritten as

$$\begin{aligned} u_i(t) &= -\kappa\Big((\nu-1)x_i(t) - \sum_{j\neq i} x_j(t)\Big) \\ &= -\kappa\nu\Big(x_i(t) - \frac{1}{\nu}\sum_{i}^{\nu} x_j(t)\Big) \end{aligned} .$$

Thus, in the "best" scenario the consensus protocol is essentially a proportional gain acting on the error between the local state and the global centroid. Moreover, since it is an invariant quantity in this case the agents essentially solve ν individual tracking problems using a particular proportional controller. ▽

When combined, these two insights show that the global centroid is somewhat akin to a reference that the agents attempt to track. Hence, it seems that the emulated centroid somehow encompasses all the required information. We then update components of μ_i at the receiver end as

$$\mu_{ij}(s_k^+) = \mu_{ij}(s_k) - \alpha_{ij}\sum_{l\in\mathcal{N}_i[k]}\big(\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)\big) \tag{2.5}$$

for all $j \neq i$ and some gains $\alpha_{ij} \in \mathbb{R}$. Summing up these updates and dividing them by ν, we have the following update algorithm for the centroid of the ith emulator:

$$\bar{\mu}_i(s_k^+) = \bar{\mu}_i(s_k) - \sum_{j\neq i}\frac{\alpha_{ij}}{\nu}\sum_{l\in\mathcal{N}_i[k]}\big(\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)\big).$$

This algorithm bears a resemblance to a discrete consensus protocol.

Having determined the control architecture, we are now in the position to analyze the closed-loop system and its capability to reach agreement. The first step in that direction is, naturally, to characterize the closed-loop dynamics.

Lemma 2.2.2. *Consider the set of agents described by* (1.2) *and controlled by* (2.3)*, where components of the emulators satisfy* (2.2) *and* (2.5)*. The cumulative state of the*

resulted closed-loop system $\mu(t) \in \mathbb{R}^{\nu^2}$ *satisfies the hybrid equations,*

$$\begin{cases} \dot{\mu}(t) = -\kappa(I_\nu \otimes L_{\hat{\mathcal{G}}})\mu(t), \quad \mu(0) = \mu_0 \\ \mu(s_k^+) = A_{\mathrm{jmp}}[k]\mu(s_k) \end{cases}, \tag{2.6}$$

for initial conditions μ_0 *comprised those of all* μ_{ij}*, where*

$$A_{\mathrm{jmp}}[k] \coloneqq I_{\nu^2} - \frac{1}{\nu}\sum_{i=1}^{\nu}\sum_{j\in\mathcal{N}_i[k]} \big(e_i(e_i - e_j)'\big) \otimes (\alpha_i \mathbb{1}'), \tag{2.7}$$

and $\alpha_i = \begin{bmatrix} \alpha_{i1} & \cdots & \alpha_{i\nu} \end{bmatrix}'$ *satisfy* $e_i'\alpha_i = 0$ *for all* $i \in \mathbb{N}_\nu$*.*

Proof The "flow" part of (2.6) is the cumulative block-diagonal version of (2.4). The "jump" part follows from (2.5) by the relations $\mu = \sum_i \sum_j e_i \otimes e_j \mu_{ij}$, $\bar{\mu}_i = (e_i' \otimes \mathbb{1}')\mu/\nu$, and with a little help of the mixed-product property of the Kronecker product in Proposition B.1.2. ∎

Equation (2.6) represents hybrid dynamics, with the continuous flow part, whose evolution is shaped by the block-diagonal matrix $-\kappa(I_\nu \otimes L_{\hat{\mathcal{G}}})$, and the discontinuous jump part, whose evolution is shaped by $A_{\mathrm{jmp}}[k]$. The former matrix is actually the Laplacian of the union of ν disconnected clones of $\hat{\mathcal{G}}$.

2.3 Divide and concur

As mentioned, the flow dynamics are comprised of ν disconnected copies of $\hat{\mathcal{G}}$. Importantly, these clones are then connected at jump stages which are dictated by $A_{\mathrm{jmp}}[k]$. The result below plays a key role to understand properties of this connectivity mechanism.

Lemma 2.3.1. *Every* $A_{\mathrm{jmp}}[k]$ *defined by* (2.7) *satisfies*

$$A_{\mathrm{jmp}}[k](I_\nu \otimes (I_\nu - P_{\mathbb{1}})) = I_\nu \otimes (I_\nu - P_{\mathbb{1}}), \tag{2.8a}$$

where $I_\nu - P_{\mathbb{1}}$ *is the orthogonal projection onto* $\ker \mathbb{1}'$*. Also,*

$$(I_\nu \otimes \mathbb{1}')A_{\mathrm{jmp}}[k] = \Big(I_\nu - \frac{1}{\nu}L_{\mathcal{G}}[k]\Big) \otimes \mathbb{1}', \tag{2.8b}$$

whenever $\mathbb{1}'\alpha_i = 1$ *for all* $i \in \mathbb{N}_\nu$*.*

Proof By the mixed-product property of the Kronecker product, $M_{\mathrm{r}}[k] \coloneqq (I_{\nu^2} - A_{\mathrm{jmp}}[k])(I \otimes (I - P_{\mathbb{1}}))$ satisfies

$$M_{\mathrm{r}}[k] = \frac{1}{\nu}\sum_{i=1}^{\nu}\sum_{j\in\mathcal{N}_i[k]} \big(e_i(e_i - e_j)'\big) \otimes (\alpha_i \mathbb{1}'(I - P_{\mathbb{1}})) = 0,$$

which proves (2.8a). Likewise,

$$M_{\mathbb{1}}[k] \coloneqq (I_\nu \otimes \mathbb{1}')(I_{\nu^2} - A_{\mathrm{jmp}}[k]) = \frac{1}{\nu}\sum_{i=1}^{\nu}\sum_{j\in\mathcal{N}_i[k]}\big(e_i(e_i-e_j)'\big)\otimes(\mathbb{1}'\alpha_i\mathbb{1}')$$
$$= \frac{1}{\nu}\sum_{i=1}^{\nu}\sum_{j\in\mathcal{N}_i[k]}\big(e_i(e_i-e_j)'\big)\otimes\mathbb{1}' = \frac{1}{\nu}L_{\mathcal{G}}[k]\otimes\mathbb{1}',$$

which yields (2.8b). ∎

It follows from (2.8b) that every $A_{\mathrm{jmp}}[k]$ has at least $\nu(\nu-1)$ eigenvalues at 1 and their right eigenspace, $\mathrm{Im}[I\otimes(I-P_{\mathbb{1}})]$, is independent of k. Hence, jumps alone *cannot* result in an agreement between agents either. But, as we show below, alternating flow and jump propagation actions does lead to agreement in the system.

To show that, we need two signals,

$$\bar{\mu}(t) \coloneqq \frac{1}{\nu}(I_\nu\otimes\mathbb{1}')\mu(t) \tag{2.9a}$$

which is the ν-dimensional vector comprised of the centroids $\bar{\mu}_i(t)$, and

$$\delta(t) \coloneqq (I_\nu\otimes(I_\nu - P_{\mathbb{1}}))\mu(t) \tag{2.9b}$$

which is the ν^2-dimensional vector of local disagreements. It is readily seen that

$$\mu(t) = (I\otimes\mathbb{1})\bar{\mu}(t) + \delta(t),$$

and $\|\mu(t)\|^2 = \nu\,\|\bar{\mu}(t)\|^2 + \|\delta(t)\|^2$ owing to the orthogonality of the centroid and disagreement. Hence, the boundedness of both $\bar{\mu}$ and δ implies that of μ itself. Moreover, because

$$x(t) = \sum_{i=1}^{\nu} e_i(e_i'\otimes e_i')\mu(t) = \bar{\mu}(t) + \sum_{i=1}^{\nu} e_i(e_i'\otimes e_i')\delta(t),$$

the states of agents x_i agree whenever the centroids agree, i.e., $\bar{\mu}(t)\in\mathrm{Im}\,\mathbb{1}$, and local disagreements vanish, i.e., $\delta(t)\to 0$. In this direction, we now show that for an appropriate choice of the weight vector α_i, the emulator centroid dynamics evolve according to a discrete consensus protocol over a switching graph.

Lemma 2.3.2. *If $\mathbb{1}'\alpha_i = 1$ for all $i\in\mathbb{N}_\nu$, then*

$$\bar{\mu}(s_{k+1}) = \Big(I - \frac{1}{\nu}L_{\mathcal{G}}[k]\Big)\bar{\mu}(s_k) \tag{2.10}$$

at sampling instances s_k, and is constant between samples, i.e., at times $t\in(s_k, s_{k+1}]$.

Proof From the jump equation in (2.6), we can analyze the emulator centroid dynamics as

$$\bar{\mu}(s_k^+) = \frac{1}{\nu}(I_\nu\otimes\mathbb{1}')A_{\mathrm{jmp}}[k]\mu(s_k) = \Big(I_\nu - \frac{1}{\nu}L_{\mathcal{G}}[k]\Big)\bar{\mu}(s_k),$$

where (2.8b) was used for the second equality. Similarly, it follows from the flow equation in (2.6) that the centroid dynamics are invariant, i.e.,

$$\dot{\bar{\mu}}(t) = \frac{1}{\nu}(I_\nu \otimes \mathbb{1}')\dot{\mu}(t) = -\frac{\kappa}{\nu}(I_\nu \otimes (\mathbb{1}' L_{\hat{\mathcal{G}}}))\mu(t) = 0,$$

because $\mathbb{1}' L_{\hat{\mathcal{G}}} = 0$. ∎

The result of Lemma 2.3.2 says that the dynamics of the centroids are completely decoupled from the rest of the state of (2.6), i.e., δ, and are driven only by the interaction topology at sampling instances. This shouldn't come as a surprise in light of the discussion preceding Example 2.2.1. Note that dynamics (2.10) are in fact standard discrete consensus dynamics with gain $\kappa = \frac{1}{\nu}$ and switching over the induced graphs $\mathcal{G}[k]$. Despite similarities to (2.1), there is one important distinction: the parameters of (2.10) do not depend on the sampling intervals. This implies that whether the emulated centroids converge does not depend on the length of the individual sampling intervals or on any upper bound on them.

Having established that the emulator centroid dynamics are decoupled from those of the disagreement vector, we now focus our attention on analyzing the disagreement dynamics.

Lemma 2.3.3. *If $\mathbb{1}'\alpha_i = 1$ for all $i \in \mathbb{N}_\nu$, then*

$$\begin{cases} \dot{\delta}(t) = -\kappa(I_\nu \otimes L_{\hat{\mathcal{G}}})\delta(t), \quad \delta(0) = \delta_0 \\ \delta(s_k^+) = \delta(s_k) + B_{\text{jmp}} L_{\mathcal{G}}[k]\bar{\mu}(s_k) \end{cases}, \tag{2.11}$$

where $\delta_0 \coloneqq (I_\nu \otimes (I_\nu - P_{\mathbb{1}}))\mu_0$ satisfies $(I \otimes \mathbb{1}')\delta_0 = 0$ and

$$B_{\text{jmp}} \coloneqq \sum_{i=1}^{\nu} (e_i e_i') \otimes (\mathbb{1}/\nu - \alpha_i)$$

satisfies $(I \otimes \mathbb{1}')B_{\text{jmp}} = 0$.

Proof By $\mathcal{A}_2$, $\mathbb{1}' L_{\hat{\mathcal{G}}} = 0$, so $(I - P_{\mathbb{1}})L_{\hat{\mathcal{G}}} = L_{\hat{\mathcal{G}}}(I - P_{\mathbb{1}})$, whence the flow part of (2.11) follows directly from that of (2.6). The jump part of (2.6) leads then to

$$\begin{aligned} \delta(s_k^+) &= (I \otimes (I - P_{\mathbb{1}}))A_{\text{jmp}}\big(\delta(s_k) + (I \otimes \mathbb{1})\bar{\mu}(s_k)\big) \\ &= \delta(s_k) + (I \otimes (I - P_{\mathbb{1}}))A_{\text{jmp}}(I \otimes \mathbb{1})\bar{\mu}(s_k), \end{aligned}$$

which can be derived by the fact that $(I \otimes (I - P_{\mathbb{1}}))\delta = \delta$ and (2.8a). By (2.8b), we have that

$$(I \otimes (I - P_{\mathbb{1}}))A_{\text{jmp}}(I \otimes \mathbb{1}) = (A_{\text{jmp}} - I)(I \otimes \mathbb{1}) + L_{\mathcal{G}}[k] \otimes \mathbb{1}/\nu,$$

and then by (2.7), that

$$(I - A_{\mathrm{jmp}})(I \otimes \mathbb{1}) = -\sum_{i=1}^{\nu} \sum_{j \in \mathcal{N}_i[k]} \big(e_i(e_i - e_j)'\big) \otimes \alpha_i = -\sum_{i=1}^{\nu} (e_i e_i' L_{\mathcal{G}}[k]) \otimes \alpha_i$$
$$= -\Big(\sum_{i=1}^{\nu} (e_i e_i') \otimes \alpha_i\Big) L_{\mathcal{G}}[k].$$

Because $L_{\mathcal{G}}[k] \otimes \mathbb{1} = (I \otimes \mathbb{1}) L_{\mathcal{G}}[k] = (\sum_i (e_i e_i') \otimes \mathbb{1}) L_{\mathcal{G}}[k]$, we have $(I \otimes (I - P_{\mathbb{1}})) A_{\mathrm{jmp}} (I \otimes \mathbb{1}) = B_{\mathrm{jmp}} L_{\mathcal{G}}[k]$ and end up with the jump part of (2.11). ∎

Lemma 2.3.3 says that δ, similarly to μ, satisfies a hybrid dynamic equation. The difference of the equation for δ, (2.11), from that for μ, (2.6), is that the former has constant "A" matrices not only in its flow part, but also for jumps. The only varying part is the discrete "B" matrix, through which the exogenous input $\bar{\mu}$ affects δ. This matrix does not affect the stability of the system, so the stability and convergence analyses are greatly simplified. In fact, if $L_{\mathcal{G}}[k]\mu(s_k) = 0$, then (2.11) has no jumps and comprises effectively ν clones of the continuous-time consensus dynamics (1.5), except that $\delta(t)$ is kept orthogonal to $\operatorname{Im} \mathbb{1}$ for all t. The latter is one of the key properties leading to the main result of this chapter.

Theorem 2.1. *If $\boldsymbol{\mathcal{A}_{1,2}}$ hold, then agents* (1.2) *controlled by* (2.3) *with emulators* (2.2), (2.5) *converge asymptotically to* $\operatorname{Im} \mathbb{1}$ *for all initial conditions, all sampling sequences with uniformly bounded sampling intervals $s_{k+1} - s_k$, and all emulator update gains α_{ij} such that $\sum_{j \neq i} \alpha_{ij} = 1$, for all $i \in \mathbb{N}_\nu$. Moreover, the emulators remain bounded and agree asymptotically as well.*

Proof We need to show that the vector of centroids, $\bar{\mu}$, agrees and the vector of local disagreements, δ, vanishes asymptotically. So consider $\bar{\mu}$ first. By $\boldsymbol{\mathcal{A}_1}$, system (2.10) satisfies the conditions of [21, Lem. 2.29, Thm. 2.37] and thus $\bar{\mu}(t)$ is bounded and converges to $\operatorname{Im} \mathbb{1}_\nu$. This also implies that the sequence $\{L_{\mathcal{G}}[k]\bar{\mu}(s_k)\}_k$ vanishes asymptotically.

Now, move to δ. The first result that we need is the stability of the autonomous version of (2.11), under $L_{\mathcal{G}}[k]\bar{\mu}(s_k) = 0$ for all $k \in \mathbb{Z}_+$. This is a standard result for continuous-time disagreement dynamics. Namely, it is known [24, Sec. II-D] that $\|\delta(t)\|$ is bounded and vanishes exponentially, with the rate determined by the smallest nonzero eigenvalue of $L_{\hat{\mathcal{G}}}$, which is the algebraic connectivity of $\hat{\mathcal{G}}$. Thus, $\delta(t)$ is the state of an exponentially stable linear system, whose exogenous input is bounded and asymptotically vanishing. Hence, $\delta \to 0$ and is bounded. ∎

2.3.1 The choice of the complete graph as $\hat{\mathcal{G}}$

An obvious problem with implementing emulators is that their dimension equals the number of agents. Emulating all agents might not be feasible for large-scale networks.

Yet this problem can be resolved by an appropriate choice of the emulated connectivity graph $\hat{\mathcal{G}}$, which is in our power.

To this end, note that each agent broadcasts only the centroid of its emulator, which actually does not change between updates. The only obstacle preventing then to emulate only the centroid is the need in individual components of μ_i in the control law (2.3). But if the emulated graph is the complete graph, then following Example 2.2.1 control law (2.3) reads

$$u_i(t) = -\kappa\Big(\nu\mu_{ii}(t) - \sum_{j=1}^{\nu}\mu_{ij}(t)\Big) = -\kappa\nu\big(x_i(t) - \bar{\mu}_i(t)\big).$$

Hence, with this choice we do not need individual μ_{ij} to implement u_i either. This, in turn, allows to drop explicit emulators. The control law becomes then

$$u_i(t) = -\kappa\nu\big(x_i(t) - \bar{\mu}_i(s_k^+)\big), \quad \forall t \in (s_k, s_{k+1}], \tag{2.3′}$$

and the emulator updates (2.5) reduce to the updates of their centroids according to

$$\bar{\mu}_i(s_k^+) = \frac{1}{\nu}\Big((\nu - |\mathcal{N}_i[k]|)\bar{\mu}_i(s_{k-1}^+) + \sum_{j\in\mathcal{N}_i[k]}\bar{\mu}_j(s_k)\Big), \tag{2.5′}$$

where the condition $\sum_j \alpha_{ij} = 1$, required in Theorem 2.1, is used. Note that the control signal in (2.3′) is still not piecewise constant, as the local feedback is analog.

The controller defined by (2.3′) and (2.5′) has an intuitive interpretation. Namely, (2.3′) is the proportional analog servo system for agent (1.2) with the piecewise-constant $\bar{\mu}_i$ as its reference signal. This reference is then updated according to the discrete consensus protocol (2.5′) with reference signals of neighboring agents. This is a reasonable strategy in the case when local agents are easy to control, but the information about the outside world is hard to acquire. And this logic appears to extend seamlessly to the cases when agents have higher-order dynamics and unmeasurable states.

Remark (impulsive hold) 2.4. If we were allowed to use the impulsive control signals, we could cause the actual state of every agent to jump at each sampling time instance. In this case the choice $\alpha_i = \mathbb{1}/\nu$ would be possible. But with this choice, $B_{\text{jmp}} = 0$ for all k, rendering δ in (2.11) completely decoupled from $\bar{\mu}$. In that scenario, the selection of agreeing initial conditions for each emulator, i.e., $\mu_i(0) \in \operatorname{Im}\mathbb{1}$, would keep each emulator in agreement for all t and allow us to implement only centroids of each emulator again, now for every $\hat{\mathcal{G}}$ satisfying $\boldsymbol{\mathcal{A}_2}$. ▽

2.3.2 What happens if connectivity fails?

A natural question is what happens if $\boldsymbol{\mathcal{A}_1}$ fails. Since the graphs are induced by the sampling, this could occur, for example, due to repeated packet losses. Once more, we can simplify the analysis by leveraging the decoupled centroids-disagreements structure.

The first step in this direction is given in the following Lemma.

Lemma 2.3.4. *Consider discrete dynamics* (2.10) *for some initial conditions* $\bar{\mu}_0$*. For any sequence of graphs*

$$\bar{\mu}_i(s_{k+1}) \in \operatorname{conv} \bar{\mu}_0 \quad \forall i \in \mathbb{N}_\nu$$

and for all k*, where* $\operatorname{conv} \bar{\mu}_0$ *is the convex hull of the initial conditions vector.*

Proof Recall that $L_{\mathcal{G}} = D_{\mathcal{G}} - A_{\mathcal{G}}$ where $A_{\mathcal{G}}$ is a binary matrix with zero diagonal and $D_{\mathcal{G}}$ is a non-negative diagonal matrix whose entries are bounded by $\nu - 1$. This implies that for any graph, the elements of $\frac{1}{\nu} L_{\mathcal{G}}[k]$ are strictly smaller than 1 in magnitude, rendering

$$M[k] \coloneqq I - \frac{1}{\nu} L_{\mathcal{G}}[k]$$

a non-negative matrix. Moreover, since $L_{\mathcal{G}}[k]\mathbb{1} = 0$, we must have $M[k]\mathbb{1} = \mathbb{1}$ thus $M[k]$ is row-stochastic.

Consider $k = 1$, each element of $\bar{\mu}[2]$ is given by

$$\bar{\mu}_i[2] = m_i[1]' \bar{\mu}_0$$

where $m_i[1]'$ is the ith row of $M[1]$. Since $M[1]$ is non-negative and row-stochastic, $\bar{\mu}_i[2]$ is a convex combination of $\bar{\mu}_0$ and therefore lies within $\operatorname{conv} \bar{\mu}_0$. By induction $\bar{\mu}_i[k+1] \in \operatorname{conv} \bar{\mu}[k]$ for all i. From convexity we have

$$\operatorname{conv} \bar{\mu}[k] \subseteq \operatorname{conv} \bar{\mu}[k-1] \subseteq \cdots \subseteq \operatorname{conv} \bar{\mu}_0$$

which concludes the proof. ∎

Now exploiting once more our freedom in choosing $\hat{\mathcal{G}}$ as the complete graph, we can state the following result.

Theorem 2.2. *If* $\boldsymbol{\mathcal{A}_2}$ *holds and* $\hat{\mathcal{G}}$ *is chosen as the complete graph, then agents* (1.2) *controlled by* (2.3′) *and* (2.5′) *remain bounded within* $\operatorname{conv} \begin{bmatrix} x_0 \\ \bar{\mu}_0 \end{bmatrix}$ *for all initial conditions, gains* $\kappa > 0$*, and sampling sequences.*

Proof In the interval $t \in [0, s_1]$ the agents evolve continuously according to

$$x(t) = \mathrm{e}^{-\kappa \nu t} x_0 + (1 - \mathrm{e}^{-\kappa \nu t}) \bar{\mu}_0.$$

Since $|\mathrm{e}^{-\kappa \nu t}| \leq 1$ for all $t \geq 0$ and $\kappa > 0$, each $x_i(t)$ is a convex combination of the corresponding $x_{0,i}$ and $\bar{\mu}_{0,i}$, thus

$$x_i(t) \in \operatorname{conv} \begin{bmatrix} x_0 \\ \bar{\mu}_0 \end{bmatrix}, \quad \forall t \in [0, s_1].$$

Now consider time interval $t \in [s_1^+, s_2]$, and not that the state is given by

$$x(t) = \mathrm{e}^{-\kappa\nu(t-s_1)}x(s_1) + (1 - \mathrm{e}^{-\kappa\nu(t-s_1)})\bar{\mu}(s_1),$$

and similar reasoning reveals that

$$x_i(t) \in \operatorname{conv} \begin{bmatrix} x(s_1) \\ \bar{\mu}(s_1) \end{bmatrix}, \quad \forall t \in [s_1^+, s_2].$$

However, from the previous analysis and Lemma 2.3.4 we know that

$$x_i(s_1) \in \operatorname{conv} \begin{bmatrix} x_0 \\ \bar{\mu}_0 \end{bmatrix}, \quad \bar{\mu}_i(s_1) \in \operatorname{conv} \bar{\mu}_0 \quad \Longrightarrow \quad x_i(t) \in \operatorname{conv} \begin{bmatrix} x_0 \\ \bar{\mu}_0 \end{bmatrix}, \quad \forall t \in [0, s_2].$$

The rest follows immediately by induction. ∎

2.4 Illustrative example

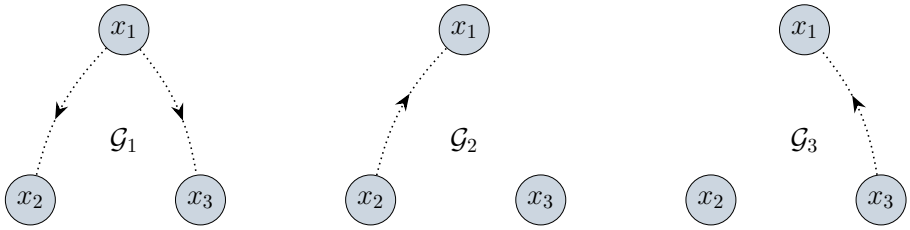

Figure 2.5: The three possible graphs for the example in Section 2.4.

To illustrate the proposed sampled-data protocol, consider the simple system comprised of $\nu = 3$ agents under the fixed (spatial) interaction topology

$$\mathcal{G} = \{(1,2), (1,3), (2,1), (3,1)\},$$

and asynchronous intermittent communication. The edge (i, j) indicates that the ith agent conveys its centroid to the jth one. By the asynchronous communication we mean that each agent transmits only at a subset of sampling instances. Consequently, in this example each $\mathcal{G}[k]$ may be the union of any nonempty subset of the graphs $\{(1,2),(1,3)\}$, $\{(2,1)\}$, and $\{(3,1)\}$, which are shown in Figure 2.5.

It should be clear that $\cup_k \mathcal{G}[k]$ contains a directed rooted tree if and only if it contains $\{(1,2),(1,3)\}$. In other words, the first agent serves as a fulcrum, facilitating information exchange between the other agents. Hence, $\boldsymbol{\mathcal{A}_1}$ holds if and only if the first agent transmits persistently, with uniformly bounded intervals. The sequence $\{k_n\}_n$ in $\boldsymbol{\mathcal{A}_1}$ may then comprise all indices of sampling instances, at which the first agent transmits.

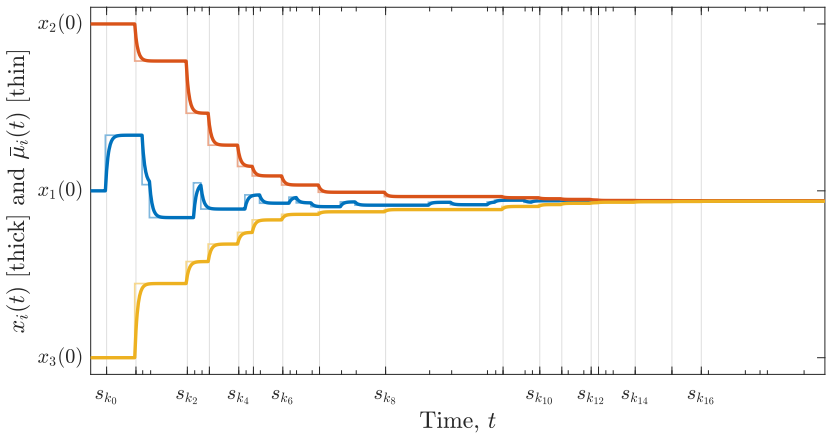

(a) $\mathcal{A}_1$ holds

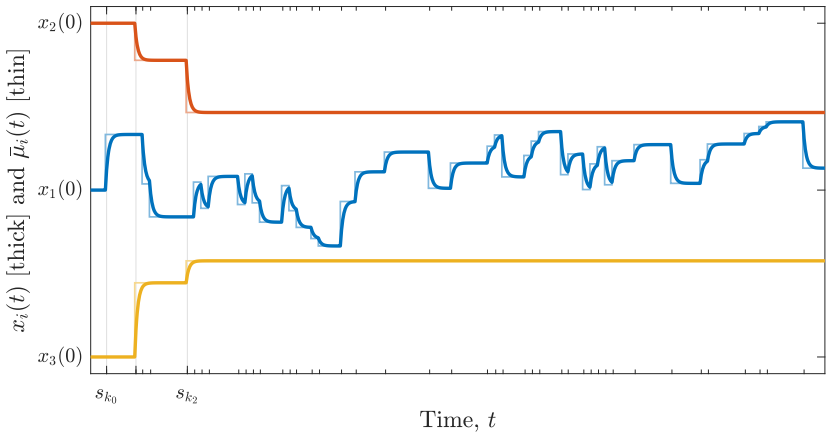

(b) $\mathcal{A}_1$ does not hold

Figure 2.6: State trajectories of agents x_i (thick lines) and centroids $\bar{\mu}_i$ (thin lines) for both examples. Minor ticks are sampling instances, major ticks indicate the connected subsequence from $\mathcal{A}_1$.

Figure 2.6 presents simulation results in the interval $t \in [0, 30]$, where the trajectories of the agents, $x_i(t)$, are depicted by thick lines and those of the centroids, $\bar{\mu}_i(t)$, are represented by thin lines. Sampling instances, shown by abscissa ticks, are a random variable such that $s_{k+1} - s_k \in 0.3\mathbb{N}_7$. Major ticks indicate the sub-sequence of sampling instances $\{k_n\}_n$ defined in $\mathcal{A}_1$. The emulators use the complete graph with $L_{\hat{\mathcal{G}}} = 3(I - P_{\mathbb{1}})$, as described in Subsection 2.3.1, with $\kappa = 3$.

First, we simulate the system for which $\mathcal{A}_1$ holds true. Specifically, the transmitting agent at each sampling instance s_k is a random pick from the set $\{1, 2, 3\}$. We can see from the plots in Figure 2.6(a) that the trajectories of the agents exhibit the behavior discussed at the end of Subsection 2.3.1, namely those of simple first order systems tracking piecewise-constant reference signals. Because agents are modeled as simple integrators, there is no local steady-state error. An increase (decrease) of κ would accelerate (slow down) local tracking. In any case, the centroids expectably converge to an agreement point, leading x_i to satisfy (1.1) asymptotically.

The situation is different when the first agent stops transmitting its information. Assume that this happens for the previous simulation after $t = s_{k_1}$, when the set of transmitting agents reduces to $\{2, 3\}$. The result shown in Figure 2.6(b) demonstrates that even multi-consensus, in which agents converge to a finite number of clusters, might not be reachable then. In our case the failure of the first agent to transmit creates, in a sense, a tug of war between the second and the third agents. The first agent gets stuck in the middle and keeps oscillating.

2.5 Concluding remarks

In this chapter we took our first steps into the realm of agreement problems. First by considering the classic consensus problem, and then by its non-trivial sampled-data counterpart. We were able to leverage ideas from optimal sampled-data control, namely designing our own A/D and D/A devices, to put forward a novel approach to solving the sampled-data consensus problem under intermittent and asynchronous sampling. The proposed architecture yields global asymptotic agreement under very mild connectivity assumptions. It was further shown that a particular choice of analog architecture can greatly reduce the complexity of the overall controller, resulting in a simple servo loop with a piecewise constant reference signal. Analysis of the controller's performance under (2.5) as well as other updating protocols is currently being investigated. Furthermore, the relatively weak assumptions required to guarantee convergence hint at potential synergy with event-triggering mechanisms.

In a broader perspective, the inherent separation between control and information processing offered by the proposed approach is particularly appealing. The methods introduced were derived for the consensus problem, but can potentially be extended to more general dynamics and multi-agent control goals. This is precisely the subject of the next chapter, which transitions from integrator consensus to more general agreement problems.

Chapter 3

From Sampled-Data Consensus to General Agreement

> Math is not thinking. Math is procedure.
> Memory is not thinking. Memory is storage.
> Thinking is thinking. Problem, solution.
>
> Andy Weir, *Project Hail Mary* [73]

Up until now we have considered only the simplest variant of $\boldsymbol{\mathcal{P}_1}$, that of identical integrator agents attempting to reach consensus to a constant. From here on, we shall consider a more general variant of agreement problem, involving general LTI agents and possibly non-constant agreement trajectories. As will be shown, the same insights used in Chapter 2 apply, after a fashion, to the general case. Moreover, the particular choices of transmitting centroids and updating them via an integrator consensus update map can be exploited significantly.

3.1 The general agreement problem

Consider ν homogeneous agents, each with linear dynamics given by

$$\Sigma_i : \begin{cases} \dot{x}_i(t) = Ax_i(t) + Bu_i(t) \\ y_i(t) = Cx_i(t) \end{cases} \tag{3.1}$$

for some $A \in \mathbb{R}^{n\times n}$ and $B \in \mathbb{R}^{n\times m}$, and $C \in \mathbb{R}^{p\times n}$, where x_i, u_i, and y_i are the ith state, control signal, and measured output, respectively. The global version of the dynamics can be written via Kronecker products as

$$\begin{cases} \dot{x}(t) = (I_\nu \otimes A)x(t) + (I_\nu \otimes B)u(t) \\ y(t) = (I_\nu \otimes C)x(t) \end{cases}. \tag{3.2}$$

The ensemble is subject to some set of communication constraints, manifesting as restrictions on the information each agent may use to generate its local control signal u_i. To this end, we define $z_i(t)$ as the local information the ith agent may transmit to its neighbors.

The spatial constraints are once more represented by neighborhood sets, $\mathcal{N}_i(t) \subset \mathbb{N}_\nu \setminus \{i\}$, where each $\mathcal{N}_i(t)$ denotes the neighbors of agent i at time t. This implies that $u_i(t)$ is some function of the local output, $y_i(t)$, and any $z_j(t)$ such that $j \in \mathcal{N}_i(t)$. Mathematically this may be written as $u_i(t) = \kappa(x_i, z_{\mathcal{N}_i(t)})$ where κ is some function and $z_{\mathcal{N}_i(t)}$ represent the neighbors of agent i at time t. The temporal constraints are represented by a strict monotonically increasing sequence of sampling instances $\{s_k\}$, $k \in \mathbb{Z}_+$, where agents may interact only on time instances $t = s_k$. We also use the convention that $t = s_k$ corresponds to the time at the receiving agent. When combined, the two restrict the ith control signal to be of the form

$$u_i(t) = \kappa(y_i(t), z_{\mathcal{N}_i(s_k)}),\ s_k \le t < s_{k+1}. \tag{3.3}$$

Consequently, at each s_k the collection of neighborhoods $\mathcal{N}_i[k]$ induces a directed graph, $\mathcal{G}[k]$, determining the permitted information exchange, as discussed in Example 2.1.3.

Remark (Scope of communication constraints) 3.1. Note that we only assume the existence of the sampling sequence $\{s_k\}$ and not how it is generated. For example, it can be time-triggered, event-triggered, stochastic, or periodic without loss of generality. Second, note that $\{s_k\}$ is a sequence of sampling instances for the entire ensemble, thus $\mathcal{N}_i[k]$ can be empty for certain agents at some $t = s_k$. This allows our framework to encompass asynchronous communication since $\{s_k\}$ is a sequence of all instances on which at least one of the agents received information. ▽

We consider the following objective in the spirit of [44] and (1.7).

$\boldsymbol{\mathcal{P}_2}$: Given $A_0 \in \mathbb{R}^{n\times n}$ such that $\operatorname{spec}(A_0) \cap \mathbb{C}_0 = \varnothing$ and its pure imaginary eigenvalues are all semi-simple, design u_i satisfying the spatio-temporal constraints and ensuring

$$\lim_{t\to\infty} \left\| x_i(t) - \mathrm{e}^{A_0 t} r_0 \right\| = 0, \quad \forall i \in \mathbb{N}_\nu, \tag{3.4}$$

for some constant $r_0 \in \mathbb{R}^n$ and all initial conditions $x_i(0)$ of agents (3.1).

It shall be emphasized that the matrix A_0 does not represent a leader node, but rather the shape of required agreement trajectories. Because setting $A_0 = 0$ recovers the consensus problem and setting $A_0 = A$ recovers the classical synchronization [42], $\boldsymbol{\mathcal{P}_2}$ may be viewed as a generalization of both. Moreover, clearly solving $\boldsymbol{\mathcal{P}_2}$ is equivalent to solving $\boldsymbol{\mathcal{P}_1}$ with an additonal constraint on the agreement trajectory. Despite the possible misnomer, unless stated otherwise, we shall use both "synchronization" and "agreement" interchangeably to refer to $\boldsymbol{\mathcal{P}_2}$, unless stated otherwise.

Agreement problems such as $\boldsymbol{\mathcal{P}_2}$ with a singular communication constraint are still well understood. For homogeneous LTI agents the problem reduces to designing a local

controller that is robust to certain perturbations, see [12] or [72, Ch. 8], in similar fashion to Proposition 2.1.2. In general, solution methods are similar to integrator consensus with some added technical complexity [36], [42]–[44]. Similarly, temporally constrained control laws with a full spatial structure are also known, even for intermittent and asynchronous sampling [74]. However, the same cannot be said when attempting to address problems subject to both constraints simultaneously, even when considering only LTI agents, see [48] and the references therein.

In such problems, the common practice is to use a sequential design. First, design a spatially constrained control law assuming continuous (analog) communication and then modify it to conform with the actual sampled communication. Such modifications usually introduce conservatism to the design. For example, zero-order hold (ZOH) discretization enforces sufficiently small and conservative gains even for integrator agents under synchronous and periodic sampling [24]. Other methods such as the input-delay approach, e.g. [75], [76], treat the sampling as a perturbation, making them inherently conservative. This creates a "design, discretize, robustify" approach which might introduce even more unnecessary conservatism compared to what was discussed in Chapter 2.

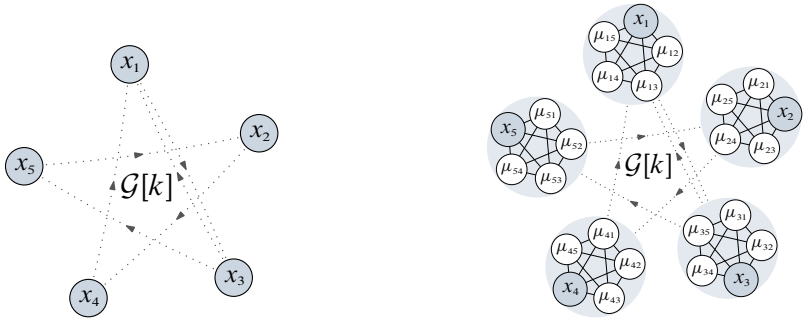

(a) Standard ZOH and ideal sampler.

(b) Emulation-based control.

Figure 3.1: Illustrations of sampled-data multi-agent control architectures: standard design vs emulation based design.

It is clear that the spatial and temporal constraints completely characterize the permitted information exchange that any permissible controller must respect. When compared with (3.3), it is clear that incorporating a ZOH synchronized with $\{s_k\}$ is just a particular option. Such controllers keep the control signal constant between updates, an additional constraint imposed by the designer rather than by the communication network. We opt for a different solution, one employing a generalized hold function [58] designed for the objective at hand. This is a well-known principle in lumped sampled-data control systems where, in the absence of an optimal solution, common wisdom dictates that the hold should attempt to reconstruct a "good" LTI continuous-time control law [68, §6.1]. In other words, it should locally *emulate* an analog closed loop, in an open-loop fashion, between samples. The conceptual difference between

the approaches is visualized in Figure 3.1. Figure 3.1(a) show the standard approach, where everything is discrete and the agents transmit their sampled states at $t = s_k$. In contrast, Figure 3.1(b) illustrates that the agents emulate an analog world between samples, resulting in time-varying control signal between samples.

Motivated by the revelations in Subsection 2.3.1, we now assume a priori that there are no spatial constraints on the emulated closed-loop. Solutions to this unconstrained problem are not readily available in the literature, hence out first step is to derive some "good" unconstrained control law to emulate. Before we begin, however, we assume throughout this chapter that

$\boldsymbol{\mathcal{A}_3}$: the triple (C, A, B) is stabilizable and detectable, and that

$\boldsymbol{\mathcal{A}_4}$: there is $\bar{F}$ such that $A_0 = A + B\bar{F}$.

Assumption $\boldsymbol{\mathcal{A}_3}$ is obviously needed for the existence of a stabilizing controller. The matching condition of $\boldsymbol{\mathcal{A}_4}$ is required for the existence of a local feedback law guaranteeing (3.4) for each agent, at least for all $\mathrm{j}\mathbb{R}$ modes of A_0. Its necessity will be show in the following sections.

The rest of this chapter is dedicated to solving $\boldsymbol{\mathcal{P}_2}$ with both full state and output only measurements. These results, sans some of the proofs, were published in [60] and [61], respectively.

3.2 Proposed architecture

3.2.1 The analog control law

The paradigm described hitherto served as the guiding principle in Chapter 2, where each agent locally emulated the consensus protocol over some agreed-upon spatial topology. Interestingly, the emulated topology could be chosen as the complete graph, i.e., centralized control law, and still result in a distributed controller respecting the communication constraints. This may be attributed to the fact that the emulators are local in nature and the particular structure of the update mechanism. Motivated by this, we shall design an unconstrained control law to emulate via the generalized hold.

To this end, consider first the simplified case of state-feedback, where we assume that $C = I$. To satisfy (3.4) the agents must track a common trajectory, implying that asymptotically the aggregate state must lie in the *agreement space*, $\mathrm{Im}(\mathbb{1}_\nu \otimes I_n)$. Introduce the signals

$$\bar{x} \coloneqq \frac{1}{\nu}(\mathbb{1}_\nu' \otimes I_n)x \quad \text{and} \quad x_\delta \coloneqq ((I_\nu - P_{\mathbb{1}}) \otimes I_n)x,$$

which may be interpreted as the *centroid* and *disagreement* signals, respectively, and satisfy $x = x_\delta + (\mathbb{1}_\nu \otimes I_n)\bar{x}$. The control objective (3.4) may then be equivalently

decomposed into two separate objectives, one for the disagreement,

$$\lim_{t\to\infty} x_\delta(t) = 0, \tag{3.5a}$$

and one for the centroid,

$$\lim_{t\to\infty} \left\| \bar{x}(t) - \mathrm{e}^{A_0 t} r_0 \right\| = 0. \tag{3.5b}$$

Because $(P_{\mathbb{1}} \otimes I_n)x = (\mathbb{1}_\nu \otimes I_n)\bar{x}$, the centroid and disagreement are orthogonal and the two objectives are independent, making it natural to propose some $u_\delta(t)$ and $\bar{u}(t)$ to independently satisfy (3.5). A state-feedback control in this vein would be

$$u(t) = (I_\nu \otimes F_\mathrm{d})x_\delta(t) + (\mathbb{1}_\nu \otimes \bar{F})\bar{x}(t) \tag{3.6}$$

for some gains F_d and $\bar{F}$, under which straightforward algebra reveals the following result.

Proposition 3.2.1. *The closed-loop dynamics of agents* (3.2) *controlled by* (3.6) *is given by*

$$\begin{aligned} \dot{x}_\delta(t) &= (I_\nu \otimes (A + BF_d))x_\delta(t) \\ \dot{\bar{x}}(t) &= (A + B\bar{F})\bar{x}(t) \end{aligned}.$$

Moreover, (3.5a) *holds if and only if* $A_d \coloneqq A + BF_d$ *is Hurwitz and* (3.5b) *holds whenever* $A_0 = A + B\bar{F}$.

Proof Rewrite (3.6) as

$$u(t) = \left((I_\nu - P_{\mathbb{1}}) \otimes F_\mathrm{d} + P_{\mathbb{1}} \otimes \bar{F}\right) x(t),$$

and note that

$$((I_\nu - P_{\mathbb{1}}) \otimes I_n)u = (I_\nu - P_{\mathbb{1}}) \otimes F_\mathrm{d} x(t) = (I_\nu \otimes F_\mathrm{d})x_\delta(t) \quad \text{and} \quad \frac{1}{\nu}(\mathbb{1}_\nu' \otimes I_n)u = \frac{1}{\nu}(\mathbb{1}_\nu' \otimes \bar{F})x(t) = \bar{F}\bar{x}(t).$$

Hence, by the mixed-product property of Proposition B.1.2

$$\dot{\bar{x}} = \frac{1}{\nu}(\mathbb{1}_\nu' \otimes I_n)\dot{x} = (A + B\bar{F})\bar{x}(t)$$

and

$$\dot{x}_\delta(t) = ((I_\nu - P_{\mathbb{1}}) \otimes I_n)\dot{x} = (I_\nu \otimes (A + BF_\mathrm{d}))x_\delta(t).$$

Note that the disagreement dynamics are in fact ν identical and independent copies of the same nth order system, while the centroid dynamics are simply one nth order system. Therefore, (3.5a) holds iff $A_\mathrm{d} \coloneqq A + BF_\mathrm{d}$ is Hurwitz and (3.5b) holds whenever $A_0 = A + B\bar{F}$. ∎

Proposition 3.2.1 shows that assumptions $\mathcal{A}_{\mathbf{3,4}}$ are indeed necessary for (3.6) to solve the problem. We require that our emulators solve the unconstrained problem, hence

these assumptions are necessary for the sampled-data case as well. It is possible to consider a more general dynamic controller, such as the one in [44], which would result in a different solvability assumption.

3.2.2 State-feedback sampled-data control law

Let $\mu_i(t) \in \mathbb{R}^{\nu n}$ denote the ith agent's emulation of the *entire* ensemble under control law (3.6). Accordingly, by $\mu_{ij}(t) \in \mathbb{R}^n$ we identify the ith agent's emulation of the jth agent's state with the convention that $\mu_{ii}(t) = x_i(t)$. As before, we can define the local disagreement of emulator i as

$$\Delta_i \coloneqq \big((I_\nu - P_{\mathbb{1}}) \otimes I_n\big)\mu_i$$

and the centroid of the ith emulator as

$$\bar{\mu}_i \coloneqq \frac{1}{\nu}(\mathbb{1}_\nu' \otimes I_n)\mu_i.$$

Between sampling instances μ_i emulates system (3.2), as if controlled by (3.6), and evolves continuously according to

$$\dot{\mu}_i(t) = \big(I_\nu \otimes A_{\mathrm{d}} + P_{\mathbb{1}} \otimes (B(\bar{F} - F_{\mathrm{d}}))\big)\mu_i(t), \tag{3.7}$$

for some initial conditions $\mu_i(0) = \mu_{i,0}$. Each agent is controlled by a local version of (3.6) based upon Δ_i and $\bar{\mu}_i$ instead of their analog counterparts, viz.

$$u_i(t) = (e_i' \otimes F_{\mathrm{d}})\Delta_i(t) + \bar{F}\bar{\mu}_i(t). \tag{3.8}$$

As previously discussed, (3.7) and (3.8) cannot solve $\boldsymbol{\mathcal{P}_2}$ on their own since each version of (3.7) evolves independently from the others. Hence, it must be accompanied by some information exchange mechanism that updates the local emulators while satisfying the spatio-temporal constraints.

We have seen that the cooperative aspect of $\boldsymbol{\mathcal{P}_2}$ can be reduced to a requirement on the centroid given by (3.5b). This mirrors the same logic and motivation we had in Section 2.2 for integrator consensus. Motivated by this similarity, we propose the same update mechanism based on the local *emulated centroids* given in (2.5). Namely, at sampling instances, each local emulator is updated according to a discrete system given by

$$\mu_{ij}(s_k^+) = \mu_{ij}(s_k) - \alpha_{ij} \sum_{l \in \mathcal{N}_i[k]} \left(\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)\right), \tag{2.5}$$

for all $i \neq j$ and some gains $\alpha_{ij} \in \mathbb{R}$. If gains $\alpha_i = \begin{bmatrix}\alpha_{i1} & \cdots & \alpha_{i\nu}\end{bmatrix}'$ are chosen such that $e_i'\alpha_i = 0$ for all $i \in \mathbb{N}_\nu$, then the closed-loop system of agents (3.2) controlled by (3.8) which is generated by (3.7) and (2.5) is given by

$$\begin{cases} \dot{\mu}(t) = \left(I_\nu \otimes \left(I_\nu \otimes A_{\mathrm{d}} + P_{\mathbb{1}} \otimes B(\bar{F} - F_{\mathrm{d}})\right)\right)\mu(t) \\ \mu(s_k^+) = (A_{\mathrm{jmp}}[k] \otimes I_n)\mu(s_k), \quad \mu(0) = \mu_0 \end{cases} \tag{3.9}$$

where $A_{\mathrm{jmp}}[k]$ is defined as in (2.7). Note that the agents communicate only through (2.5) which is spatially distributed, thus the controller respects the spatio-temporal constraints. Note that the update scheme is identical to the one we had for simple integrators, and indeed we must make exactly the same connectivity assumption $\boldsymbol{\mathcal{A}_1}$. Now, with a controller at hand, we are set to show that it solves $\boldsymbol{\mathcal{P}_2}$ with full state measurements.

3.3 Sampled-data solution to $\boldsymbol{\mathcal{P}_2}$ with state measurements

In Subsection 3.2.1 we saw that the disagreement and centroid dynamics were decoupled by (3.6), which in turn enabled $\boldsymbol{\mathcal{P}_2}$ to be reduced into two independent problems. Inspired by this, consider the following partition of the stacked emulators

$$\mu(t) = (I_\nu \otimes \mathbb{1}_\nu \otimes I_n)\bar{\mu}(t) + \Delta(t) \in \mathbb{R}^{\nu^2 n}.$$

Now, $\bar{\mu}(t)$ is an $(n\nu) \times 1$ *block vector*, where the ith $n \times 1$ block contains the centroid of the ith emulator. Similarly, $\Delta(t)$ is a block vector where blocks $\Delta_i(t)$ contain the local disagreement vector of emulator i. Recall that $x_i(t) = \mu_{ii}(t)$, hence the aggregate state is given by

$$x(t) = \bar{\mu}(t) + \sum_{i=1}^{\nu} \left((e_i(e_i' \otimes e_i')) \otimes I_n\right) \Delta(t).$$

The above allows us to pose equivalent conditions for the solution of $\boldsymbol{\mathcal{P}_2}$ in the same vein as those presented in (3.5). Namely, if the emulator disagreements asymptotically vanish,

$$\lim_{t\to\infty} \Delta(t) = 0, \tag{3.10a}$$

and the emulator centroids verify

$$\lim_{t\to\infty} \left\| \bar{\mu}(t) - \mathbb{1}_\nu \otimes \mathrm{e}^{A_0 t} r_0 \right\| = 0, \tag{3.10b}$$

then $\boldsymbol{\mathcal{P}_2}$ is satisfied.

3.3.1 Centroid-disagreement separation and synchronization

Unlike their analog counterparts, $\bar{\mu}(t)$ and $\Delta(t)$ are coupled through (2.5). However, for some choices of update gains α_i they take on a simple structure, allowing for a more streamlined analysis. Below is the high-order counterpart of Lemma 2.3.2 and Lemma 2.3.3.

Lemma 3.3.1. *If $\mathbb{1}_\nu'\alpha_i = 1$ for all $i \in \mathbb{N}_\nu$, then the disagreements dynamics are given by*

$$\begin{cases} \dot{\Delta}(t) = (I_{\nu^2} \otimes A_d)\Delta(t), \quad \Delta(0) = \Delta_0 \\ \Delta(s_k^+) = \Delta(s_k) + \big((B_{\text{jmp}} L_{\mathcal{G}}[k]) \otimes I_n\big)\bar{\mu}(s_k), \end{cases} \tag{3.11a}$$

with

$$B_{\text{jmp}} \coloneqq \sum_{i=1}^{\nu} (e_i e_i') \otimes (\mathbb{1}_\nu/\nu - \alpha_i),$$

and $L_{\mathcal{G}}[k]$ is the Laplacian matrix associated with $\mathcal{G}[k]$ [22, §2.3.5]. In addition, the centroid dynamics are given by

$$\begin{cases} \dot{\bar{\mu}}(t) = (I_\nu \otimes (A + B\bar{F})\bar{\mu}(t), & \bar{\mu}(0) = \bar{\mu}_0 \\ \bar{\mu}(s_k^+) = ((I_\nu - \frac{1}{\nu} L_{\mathcal{G}}[k]) \otimes I_n)\bar{\mu}(s_k) \end{cases} \tag{3.11b}$$

Proof The flow dynamics mirror the analog case. Note that $A_{\text{jmp}}[k]$ is defined identically to the jump map in [59, Lemma 2], hence the result follows from applying it to the jump map of (3.9). ∎

There are two immediate consequences of Lemma 3.3.1: i) the dynamics of $\bar{\mu}(t)$ are autonomous and do not depend on those of $\Delta(t)$ and ii) $\bar{\mu}(t)$ can be thought of as a discrete input affecting $\Delta(t)$ at time instances $t = s_k$. Consequently, finding conditions under which the centroids satisfy (3.10b) can be done independently of $\Delta(t)$. This is the purpose of the following result.

Lemma 3.3.2. *Consider* (3.11b) *and denote $A_0 \coloneqq A + B\bar{F}$. If $\mathcal{A}_1$ holds true, then there is r_0 such that*

$$\lim_{t\to\infty} \|\bar{\mu}(t) - \mathbb{1} \otimes \bar{\mu}_{ss}(t)\| = 0,$$

where the n-dimensional $\bar{\mu}_{ss}$ is such that $\bar{\mu}_{ss}(t) = \mathrm{e}^{A_0 t} r_0$. Moreover,

$$r_0 \coloneqq \left(\sum_{i=1}^{\nu} q_i \bar{\mu}_{i,0}\right) = (q' \otimes I_n)\bar{\mu}_0$$

where q is some constant vector which depends on the sequence of graphs.

Proof It is readily verified that $\bar{\mu}$ from (3.11b) satisfies

$$\bar{\mu}(s_k + \tau) = \mathrm{e}^{I \otimes (A_0 \tau)} \left(\prod_{j=1}^{k} \left(\left(I - \frac{1}{\nu} L_{\mathcal{G}}[j] \right) \otimes I \right) \mathrm{e}^{I \otimes (A_0 h_j)} \right) \bar{\mu}_0$$

for all k and $0 < \tau \leq h_{k+1}$, where $h_j \coloneqq s_j - s_{j-1}$. Because $\mathrm{e}^{I \otimes (A_0 h_j)} = I \otimes \mathrm{e}^{A_0 h_j}$ and $N \otimes I$ and $I \otimes M$ commute for all compatibly dimensioned M and N, we have

$$\bar{\mu}(s_k + \tau) = \left(\left(\prod_{j=1}^{k} \left(I - \frac{1}{\nu} L_{\mathcal{G}}[j] \right) \right) \otimes \mathrm{e}^{A_0 (s_k + \tau)} \right) \bar{\mu}_0.$$

If the connectivity assumption $\mathcal{A}_1$ holds, then [21, Lem. 2.29 and 2.30] there exists some constant $q \in \mathbb{R}^\nu$ such that

$$\lim_{k\to\infty} \prod_{j=1}^{k} \Big(I - \frac{1}{\nu} L_{\mathcal{G}}[j]\Big) = \mathbb{1} q'. \tag{3.12}$$

Therefore, if we choose $r_0 = (q' \otimes I)\bar{\mu}_0$ for $\bar{\mu}_{\mathrm{ss}}$, then

$$\lim_{k\to\infty} \big(\bar{\mu}(s_k + \tau) - \mathbb{1} \otimes \bar{\mu}_{\mathrm{ss}}(s_k + \tau)\big) = \lim_{k\to\infty} \left(\left(\prod_{j=1}^{k} \Big(I - \frac{1}{\nu} L_{\mathcal{G}}[j]\Big) - \mathbb{1} q'\right) \otimes \mathrm{e}^{A_0(s_k+\tau)}\right) \bar{\mu}_0 = 0$$

whenever $\mathrm{e}^{A_0 t}$ is bounded. The latter is guaranteed by the assumption that all pure imaginary eigenvalues of A_0 are semi-simple. ∎

Utilizing $\mathcal{A}_4$, the flow map ensures that each $\bar{\mu}_i(t)$ aligns precisely with the intended trajectory shape. However, if the initial conditions differ the trajectories would be different. This cannot be remedied by the non-interacting flows. On the other hand, the jump map mirrors discrete consensus dynamics. Under $\mathcal{A}_1$, this map will asymptotically steer a constant vector to a fixed consensus point within the agreement space. Thus, only the combined flow and jump dynamics under both assumptions guarantee the solution of $\mathcal{P}_2$ for $\bar{\mu}(t)$. Hence, Lemma 3.3.2 proves that under $\mathcal{A}_{4,1}$ and a proper choice of $\bar{F}$, $\bar{\mu}(t)$ asymptotically satisfy (3.10b). The final step would be to show the stability of (3.11a), which can be thought of as an LTI system with $\bar{\mu}$ as an impulsive input. Moreover, the "input matrix" for these impulses includes a, possibly different, graph Laplacian matrix at each k. In particular, for any graph $\mathbb{1}_\nu \in \ker L_{\mathcal{G}}[k]$, therefore if $\bar{\mu}(s_k) \in \operatorname{Im} \mathbb{1}_\nu \otimes I_\nu$ then (3.11a) will contain no jumps. This is a key property in proving the main result, which is stated below.

Theorem 3.1. *Consider agents* (3.1) *with* $C = I$*, controlled by* (3.8)*, generated by emulators* (3.9) *and update law* (2.5)*. If* $\mathcal{A}_{4,1}$ *holds, then control law* (3.8) *solves* $\mathcal{P}_2$ *for all gains* F_d *and* $\bar{F}$ *such that* $A_d = A + BF_d$ *is Hurwitz and* $A + B\bar{F} = A_0$ *and all emulator update gains* α_{ij} *such that* $\mathbb{1}_\nu' \alpha_i = 1$ *for all* $i \in \mathbb{N}_\nu$*. Moreover, the emulators asymptotically agree and remain bounded if* $\mathrm{e}^{A_0 t}$ *is bounded.*

Proof By assumption, $\mathbb{1}_\nu' \alpha_i = 1$ for all $i \in \nu$, therefore the condition for Lemma 3.3.1 holds. Define the centroid error $\bar{\epsilon}(t) \coloneqq \bar{\mu}(t) - \mathbb{1}_\nu \otimes \mathrm{e}^{A_0 t} r_0$ where r_0 is in Lemma 3.3.2. By assumption $A + B\bar{F} = A_0$, and A_{d} is Hurwitz and $\mathcal{A}_{4,1}$ holds, thus from Lemma 3.3.2 we know that $\bar{\mu}(t) \to \mathbb{1}_\nu \otimes \mathrm{e}^{A_0 t} r_0$, or equivalently that $\bar{\epsilon}(t) \to 0$ from every initial condition. Note that $\mathbb{1}_\nu \in \ker L_{\mathcal{G}}[k]$ for all k, thus we can rewrite the jump part of (3.11a) as

$$\Delta(s_k^+) = \Delta(s_k) + \big(B_{\mathrm{jmp}}[k] \otimes I_n\big)\big(L_{\mathcal{G}}[k] \otimes I_n\big)\bar{\epsilon}(s_k).$$

Since $\bar{\epsilon}(t) \to 0$ and is bounded, the sequence $\{(L_{\mathcal{G}}[k] \otimes I_n)\bar{\epsilon}(s_k)\}$ is bounded and vanishes asymptotically, reducing (3.11a) to an LTI system with a bounded and asymptot-

ically vanishing input. The stability of LTI systems with bounded and vanishing inputs is independent of the actual input, therefore since A_{d} is Hurwitz $\Delta(t) \to 0$. Combining with (3.10) yields

$$\lim_{t\to\infty} x(t) = \lim_{t\to\infty} \bar{\mu}(t) = \mathbb{1}_\nu \otimes \mathrm{e}^{A_0 t} r_0,$$

implying that the agents agree. Similarly, taking the limit for $\mu(t)$ yields

$$\lim_{t\to\infty} \mu(t) = \lim_{t\to\infty} (I_\nu \otimes \mathbb{1}_\nu \otimes I_n)\bar{\mu}(t) = \mathbb{1}_{\nu^2} \otimes \mathrm{e}^{A_0 t} r_0,$$

therefore the emulators also agree and remain bounded if $\mathrm{e}^{A_0 t}$ is bounded. ∎

3.3.2 Reduced order implementation

The obvious drawback of emulation-based control architectures is that each agent must locally emulate the entire group, yielding local controllers whose dimension grows linearly with $n\nu$. This may not be feasible for large networks of high-order agents. In an effort to circumvent that, consider a different representation of

$$u_i(t) = F_{\mathrm{d}} \mu_{ii}(t) + (\bar{F} - F_{\mathrm{d}})\bar{\mu}_i(t), \tag{3.8′}$$

which is obtained by substituting $\Delta_i(t) = \mu_i(t) - (\mathbb{1}_\nu \otimes I_n)\bar{\mu}_i(t)$. Control law (3.8′) requires two nth order states, the local emulated centroid and local emulated state, hinting that it might be possible to obtain a reduced order implementation.

Corollary 3.2. *If $\mathcal{A}_{4,1}$ holds and each agent can continuously measure its own state then the following nth order local controllers*

$$\begin{cases} \dot{\bar{\mu}}_i(t) = (A + B\bar{F})\bar{\mu}_i(t), \quad \bar{\mu}_i(0) = \bar{\mu}_{i,0} \\ \bar{\mu}_i(s_k^+) = \bar{\mu}_i(s_k) - \dfrac{1}{\nu} \displaystyle\sum_{l \in \mathcal{N}_i[k]} (\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)) \\ u_i(t) = F_d x_i(t) + (\bar{F} - F_d)\bar{\mu}_i(t) \end{cases} \tag{3.13}$$

solves $\mathcal{P}_2$ for all gains $F_d, \bar{F}$ such that $A + B\bar{F} = A_0$ and A_d is Hurwitz.

Proof By definition $\mu_{ii}(t) = x_i(t)$ which is locally available continuously by assumption, substituting $\mu_{ii}(t)$ with $x_i(t)$ in (3.8′) gives the first equivalence. Since $x_i(t)$ is locally available, to implement the control each agent needs to implement only $\bar{\mu}_i(t)$. By Lemma 3.3.1 we know that $\bar{\mu}(t)$ is independent of $\Delta(t)$ thus the rest of (3.13) follows immediately from considering the local version of (3.11b). ∎

The above implementation is still distributed and adheres to the spatial and temporal constraints, but now each local controller is only of dimension n regardless of the number of agents. This agrees with the intuition behind the analog control law from Subsection 3.2.1: the agents must track the centroid and drive the disagreements to

zero. The logic is reminiscent of a classic servo-regulation problem, where the control law has a stabilizing component acting on the state and a tracking component acting on the reference signal. This raises the natural question of how will this structure behave in the presence of disturbances, something synchronizing systems tend to do poorly [63].

Note that (3.13) can also be presented in a slightly different form. Using the coordinate transformation $z_i = \nu\bar{\mu}_i - x_i$, it can be rewritten locally as

$$\begin{cases} \dot{z}_i(t) = (A + B\bar{F})z_i(t) + B(\bar{F}x_i(t) - u_i(t)), \quad z_i(0) = z_{i,0} \\ z_i(s_k^+) = z_i(s_k) - \dfrac{1}{\nu}\displaystyle\sum_{l\in\mathcal{N}_i[k]}(z_i(s_k) - z_l(s_k) + x_i(s_k) - x_l(s_k)) \\ u_i(t) = F_{\mathrm{d}}x_i(t) + \dfrac{1}{\nu}(\bar{F} - F_{\mathrm{d}})(z_i(t) + x_i(t)) \end{cases}. \tag{3.13'}$$

This form is reminiscent of the Youla parametrization-based redesign approach in [69, Thm. 4.2], which also resulted in an emulation-based scheme. Implementation (3.13') will be beneficial later on when we consider the effects of transmission delays.

3.4 What changes in the output-feedback case

If the state x_i is not measurable, we can no longer realize the dynamics (3.9) in the intersample, its μ_{ii} component is not available. The use of a state observer is a conventional solution in such situations. There is certain ambiguity in how exactly an observer may be incorporated into the emulation and information exchange procedures. The extension proposed below is motivated mainly by the relative simplicity of analyzing the closed-loop dynamics with it.

Because measurement channels of agents (3.1) are uncoupled, we construct the local, i.e., uncoupled, analog observer

$$\dot{\hat{x}}_i(t) = A\hat{x}_i(t) + Bu_i(t) - L(y_i(t) - C\hat{x}_i(t)) \tag{3.14}$$

for some L such that $A + LC$ is Hurwitz. The observer-based counterpart of (3.6) is then straightforward, we just need to replace x_i and $\bar{x}$ with $\hat{x}_i$ and the centroid of the observer states of all agents. It is then readily seen that the resulting disagreement x_δ and centroid $\bar{x}$ still satisfy (3.5), the only change in their evolution is the addition of the aggregate observer error, which vanishes exponentially.

Moving to the spatially distributed sampled-data setting, we now substitute the control law (3.8) with

$$u_i(t) = F_{\mathrm{d}}\hat{x}_i(t) + (\bar{F} - F_{\mathrm{d}})\bar{\mu}_i(t), \tag{3.15}$$

where the observed state $\hat{x}_i$ is local and can thus be implemented in continuous time and the emulated centroid $\bar{\mu}_i$ is still generated by the n-dimensional hybrid system (3.13).

Remark 3.2. It can be shown the centroid of the observer states of all agents, say $\check{x} \coloneqq (1/\nu)\sum_{i=1}^{\nu}\hat{x}_i$, under the analog observer-based counterpart of (3.6) satisfies

$$\begin{bmatrix}\dot{\check{x}}\\ \dot{\bar{\epsilon}}\end{bmatrix} = \begin{bmatrix}A_0 & -LC\\ 0 & A+LC\end{bmatrix}\begin{bmatrix}\check{x}\\ \bar{\epsilon}\end{bmatrix},$$

where $\bar{\epsilon} \coloneqq \bar{x} - \check{x}$ is the centroid observation error. This relation may be used for alternative forms of the emulator. Exploring these alternatives might involve some involved technicalities and is thus left for future research. ▽

3.5 Synchronization analysis for output feedback

Combining plant (3.2) with the aggregate versions of (3.13),(3.14), and (3.15), the closed-loop dynamics read as the analog flow

$$\begin{cases}\dot{x} = (I\otimes A)x + (I\otimes BF_{\mathrm{d}})\hat{x} + (I\otimes B(\bar{F}-F_{\mathrm{d}}))\bar{\mu}\\ \dot{\hat{x}} = -(I\otimes LC)x + (I\otimes (A_{\mathrm{d}}+LC))\hat{x} + (I\otimes B(\bar{F}-F_{\mathrm{d}}))\bar{\mu}\\ \dot{\bar{\mu}} = (I\otimes A_0)\bar{\mu}\end{cases}$$

between sampling instances with the jump

$$\begin{cases}x(s_k^+) = x(s_k)\\ \hat{x}(s_k^+) = \hat{x}(s_k)\\ \bar{\mu}(s_k^+) = (I-(1/\nu)L_{\mathcal{G}}[k])\bar{\mu}(s_k)\end{cases}$$

at each s_k, where $L_{\mathcal{G}}[k]$ is the Laplacian matrix associated with the network connectivity graph $\mathcal{G}[k]$ at s_k. Note that $L_{\mathcal{G}}[k]\mathbb{1} = 0$ for all k.

Introduce now the emulation and observation errors

$$\varepsilon \coloneqq x - \bar{\mu} \quad \text{and} \quad \epsilon \coloneqq x - \hat{x},$$

respectively, the closed-loop dynamics can be rewritten in the more transparent form

$$\begin{bmatrix}\dot{\varepsilon}(t)\\ \dot{\epsilon}(t)\\ \dot{\bar{\mu}}(t)\end{bmatrix} = \begin{bmatrix}I\otimes A_{\mathrm{d}} & -I\otimes(BF_{\mathrm{d}}) & 0\\ 0 & I\otimes(A+LC) & 0\\ 0 & 0 & I\otimes A_0\end{bmatrix}\begin{bmatrix}\varepsilon(t)\\ \epsilon(t)\\ \bar{\mu}(t)\end{bmatrix} \tag{3.16a}$$

(here $A_{\mathrm{d}} = A + BF_{\mathrm{d}}$ and $A_0 = A + B\bar{F}$), with the jump

$$\begin{bmatrix}\varepsilon(s_k^+)\\ \epsilon(s_k^+)\\ \bar{\mu}(s_k^+)\end{bmatrix} = \begin{bmatrix}I & 0 & (1/\nu)L_{\mathcal{G}}[k]\otimes I_n\\ 0 & I & 0\\ 0 & 0 & (I-(1/\nu)L_{\mathcal{G}}[k])\otimes I_n\end{bmatrix}\begin{bmatrix}\varepsilon(s_k)\\ \epsilon(s_k)\\ \bar{\mu}(s_k)\end{bmatrix}. \tag{3.16b}$$

The signal ε is affected by both ϵ, via flow (3.16a), and $\bar{\mu}$, via jump (3.16b). At the same time, ϵ and $\bar{\mu}$ are completely decoupled. As such, we start the analysis with the last two signals.

It shall be clear that ϵ is an exponentially decaying signal. Therefore, it does not affect asymptotic properties of (3.16) and can be excluded from the analysis. Asymptotic behavior of $\bar{\mu}$ is generally more complex, but follows exactly the same dynamics it did in the state-feedback case. In particular, by Lemma 3.3.2 we know that although the νn-dimensional signal $\bar{\mu}$ is not decaying, all its n-dimensional block components are asymptotically equivalent. This leads to the following result, which is the main result of this section.

Theorem 3.3. *If F_d and L are such that $A + BF_d$ and $A + LC$ are Hurwitz and $\bar{F}$ is such that $A_0 = A + B\bar{F}$, then the control law defined by* (3.15), (3.14), *and* (3.11b) *solves $\boldsymbol{\mathcal{P}_2}$ for any sampling sequence $\{s_k\}$ satisfying $\boldsymbol{\mathcal{A}_1}$.*

Proof By Lemma 3.3.2 and the fact that $L_{\mathcal{G}}[k]\mathbb{1} = 0$ we have that

$$\lim_{k\to\infty} L_{\mathcal{G}}[k]\bar{\mu}_{\mathrm{ss}}(s_k) = 0,$$

The latter property implies that $\bar{\mu}$ asymptotically decouples from ε in (3.16b). Because the matrix A_{d} is Hurwitz and because ϵ vanishes exponentially, we have $\lim_{t\to\infty} \varepsilon(t) = 0$. This, in turn, yields

$$\lim_{t\to\infty} \|x(t) - \mathbb{1} \otimes \bar{\mu}_{\mathrm{ss}}(t)\| = 0,$$

which leads to (3.4). ∎

3.6 Directly emulating the observers

It is worth emphasising that despite the simplicity of the control law defined by (3.15), (3.14), and (3.11b) it is not merely a reapplication of the methodology from the state-feedback case. In fact, repeating the emulation process described in Section 3.3 with the simple change of $\mu_{ii} \equiv \hat{x}_i$ would result in a significantly different system. It can be shown that this process would result in the following counterpart of (3.16)

$$\begin{bmatrix} \dot{\varepsilon}(t) \\ \dot{\epsilon}(t) \\ \dot{\bar{\mu}}(t) \end{bmatrix} = \begin{bmatrix} I \otimes \bar{\bar{A}} & -I \otimes M & 0 \\ 0 & I \otimes (A + LC) & 0 \\ 0 & -I \otimes (\frac{1}{\nu} LC) & I \otimes A_0 \end{bmatrix} \begin{bmatrix} \varepsilon(t) \\ \epsilon(t) \\ \bar{\mu}(t) \end{bmatrix} \tag{3.16a′}$$

(here $\bar{\bar{A}} = A + BF_{\mathrm{d}} + B\bar{F}$, $A_0 = A + B\bar{F}$, and $M = BF_{\mathrm{d}} + (1/\nu)LC$), with the jumps

$$\begin{bmatrix} \varepsilon(s_k^+) \\ \epsilon(s_k^+) \\ \bar{\mu}(s_k^+) \end{bmatrix} = \begin{bmatrix} I & 0 & (1/\nu)L_{\mathcal{G}}[k] \otimes I_n \\ 0 & I & 0 \\ 0 & 0 & (I - (1/\nu)L_{\mathcal{G}}[k]) \otimes I_n \end{bmatrix} \begin{bmatrix} \varepsilon(s_k) \\ \epsilon(s_k) \\ \bar{\mu}(s_k) \end{bmatrix}. \tag{3.16b′}$$

Since now we do not have access to the actual states, this emulation leads to coupling between the flow of each $\bar{\mu}_i(t)$ and its observation error. Note that A_0 will generally have eigenvalues on the imaginary axis, hence not asymptotically stable, and that the stability of $\bar{\bar{A}}$ is not guaranteed. Thus while the estimation error, $\epsilon(t)$, is LTI and decays exponentially to zero, the same cannot be said for $\bar{\mu}(t)$ and $\varepsilon(t)$. In fact both are hybrid, non-autonomous, with an unstable flow and shift varying jumps.

By forgoing the straightforward derivation via emulation methodology we were able to decouple $\bar{\mu}$ for the state and the observer, cumulating with the simpler (3.16) rather than (3.16′). This significantly streamlined the proof and allowed us to avoid analyzing the aforementioned complicated hybrid system. The "price" we pay for the simplified analysis is that $\bar{\mu}$ is now completely decoupled, and in fact can be thought of as some sort of exosystem without direct feedback from the agents. It is worth mentioning that the results of Theorem 3.3 still hold for (3.16′), but the proof is significantly longer and more involved. It will be presented in the rest of this section for completion, but is not important for the rest of this thesis.

3.6.1 The analog control law

First, we must verify that the emulated control law can actually solve the prolbem in the ideal case. To this end, consider the ideal scenario where each agent locally implements (3.6) by replacing the state components with those of observer (3.14), yielding

$$\begin{cases} \dot{\hat{x}}_i(t) = A\hat{x}_i(t) + Bu_i(t) + L(C\hat{x}_i(t) - y_i(t)) \\ u_i(t) = F_d\hat{x}_{\delta,i}(t) + \bar{F}\hat{\bar{x}}(t) \end{cases} \tag{3.17}$$

where $\hat{x}_{\delta,i}((t)$ and $\hat{\bar{x}}$ are the disagreement and centroid of the observer's states. Following the same logic of partitioning the states to centroid and disagreements, the closed-loop dynamics are once more decoupled. Moreover, if $A + LC$ is Hurwitz then the estimation error $\epsilon(t) \coloneqq x(t) - \hat{x}(t)$ converges exponentially to zero. Thus if $\hat{x}$ synchronize asymptotically, so would $x(t)$.

Defining observation errors $\epsilon_\delta = x_\delta - \hat{x}_\delta$ and $\bar{\epsilon} = \bar{x} - \hat{\bar{x}}$ results in the following global observer disagreement dynamics

$$\begin{bmatrix} \dot{\hat{\delta}}(t) \\ \dot{\epsilon}_\delta(t) \end{bmatrix} = \begin{bmatrix} I_\nu \otimes (A + BF_d) & -I_\nu \otimes (LC) \\ 0 & I_\nu \otimes (A + LC) \end{bmatrix} \begin{bmatrix} \hat{\delta}(t) \\ \epsilon_\delta(t) \end{bmatrix} \tag{3.18a}$$

and observer centroid dynamics

$$\begin{bmatrix} \dot{\hat{\bar{x}}}(t) \\ \dot{\bar{\epsilon}} \end{bmatrix} = \begin{bmatrix} A + B\bar{F} & -LC \\ 0 & A + LC \end{bmatrix} \begin{bmatrix} \hat{\bar{x}}(t) \\ \bar{\epsilon}(t) \end{bmatrix}. \tag{3.18b}$$

Recall that $A + BF_d$ is Hurwitz, hence (3.18a) is stable iff $A + LC$ is Hurwitz, as par standard observer design, and satisfies condition (3.5a). However, $A + B\bar{F} = A_0$ is not

assumed to be Hurwitz, thus the fact that $\bar{\epsilon} \to 0$ does not imply that condition (3.5b) will be satisfied. The following proposition characterizes conditions under which the observer based feedback solves $\boldsymbol{\mathcal{P}_2}$.

Proposition 3.6.1. *If $A + BF_d$ and $A + LC$ are Hurwitz, $A_0 = A + B\bar{F}$, and* $\operatorname{spec} A_0 \cap \operatorname{spec}(A + LC) = \varnothing$*; then observer-based controller* (3.17) *solves $\boldsymbol{\mathcal{P}_2}$ for agents* (3.1)*. Moreover, $r_0 = (I - Q)\hat{\bar{x}}_0 + X\bar{x}_0$ where Q is the solution of the following Sylvester equation*

$$Q(A + LC) - A_0 Q = LC.$$

Proof Since $A+LC$ and $A+BF_d$ are Hurwitz condition (3.5a) is satisfied, thus we need to consider only the centroid dynamics. Consider a coordinate transformation of (3.18b) given by

$$T = \begin{bmatrix} I & Q \\ 0 & I \end{bmatrix}, T^{-1} = \begin{bmatrix} I & -Q \\ 0 & I \end{bmatrix},$$

applying it yields

$$T \begin{bmatrix} A_0 & -LC \\ 0 & A + LC \end{bmatrix} T^{-1} = \begin{bmatrix} A_0 & Q(A + LC) - A_0 Q - LC \\ 0 & A + LC \end{bmatrix}.$$

It is well known [77, Thm 1.1.5] that if $\operatorname{spec} A_0 \cap \operatorname{spec} A + LC = \varnothing$ then

$$Q(A + LC) - A_0 Q - LC = 0$$

always has a unique solution, thus the transformation decouples the centroid dynamics. Thus the trajectory of the transformed coordinate $\hat{\bar{x}}_Q = \hat{\bar{x}} + Q\bar{\epsilon}$ is given by

$$\hat{\bar{x}}_Q(t) = \mathrm{e}^{A_0 t}\hat{\bar{x}}_Q(0) = \mathrm{e}^{A_0 t}\left((I - Q)\hat{\bar{x}}_0 + Q\bar{x}(0)\right)$$

and in the original coordinates

$$\hat{\bar{x}}(t) = \hat{\bar{x}}_Q(t) - Q\bar{\epsilon}(t) = \mathrm{e}^{A_0 t} r_0 - Q\mathrm{e}^{A_L t}\bar{\epsilon}(0)$$

where $A_L = A + LC$. Since A_L is Hurwitz

$$\lim_{t \to \infty} \left\| \hat{\bar{x}}(t) - \mathrm{e}^{A_0 t} r_0 \right\| = 0$$

and similarly for the real centroid $\bar{x}(t)$. ∎

Proposition 3.6.1 shows that the observer based variation of the "ideal" control law can indeed solve $\boldsymbol{\mathcal{P}_2}$ without temporal constraints, making it suitable for emulation. Note that this time we have an additional constraint on the spectra of A_0 and $A + LC$, which is required to ensure the existence of a solution to a Sylvester equation. It is to be expected, and is indeed true, that a similar constraint will appear in the sampled-data variant as well.

3.6.2 Sampled-data synchronization

Now that we have established that the analog control law can indeed drive the emulators to agreement between sampling instances, we can consider the full hybrid dynamics of (3.16a′) and (3.16b′). The main difference is that the flow map of $\bar{\mu}$ in (3.16a) is autonomous, while in (3.16a′) it is also driven by $\epsilon(t)$. Even though $\epsilon(t)$ is completely continuous and exponentially stable, this coupling makes the hybrid trajectory of $\bar{\mu}$ much more complicated. In particular, each sampling instance would introduce a new term which is a product of the jump map at $t = s_k$ and an integral equation driven by ϵ. It is readily verified by combining (3.16a′) and (3.16b′) that $\bar{\mu}$ evolves according to

$$\begin{aligned}\bar{\mu}(s_k+\tau) &= \left(I_\nu \otimes \mathrm{e}^{A_0(s_k+\tau)}\right)\Bigg[\left(\prod_{l=1}^{k} M[l]\otimes I_n\right)\bar{\mu}(0) \\ &\quad - \frac{1}{\nu}\sum_{l=1}^{k}\left(\prod_{j=l}^{k} M[j]\otimes\left(\int_{s_{l-1}}^{s_l}\mathrm{e}^{-A_0 s}LC\mathrm{e}^{A_L s}\mathrm{d}s\right)\right)\epsilon(s_{l-1})\Bigg]. \\ &= \left(I_\nu \otimes \mathrm{e}^{A_0(s_k+\tau)}\right)\left(\eta[k]-\frac{1}{\nu}\theta[k]\right)\end{aligned} \tag{3.19}$$

where

$$M[k] \coloneqq I_\nu - \frac{1}{\nu}L_{\mathcal{G}}[k], \quad \eta[k] \coloneqq \left(\prod_{l=1}^{k} M[l]\otimes I_n\right)\bar{\mu}(0),$$

and

$$\theta[k] \coloneqq \sum_{l=1}^{k}\left(\prod_{j=l}^{k} M[j]\otimes\left(\int_{s_{l-1}}^{s_l}\mathrm{e}^{-A_0 s}LC\mathrm{e}^{A_L s}\mathrm{d}s\right)\right)\epsilon(s_{l-1}).$$

Note that if $\eta[k]$ and $\theta[k]$ converge to finite limits then $\bar{\mu}$ will asymptotically converge to a solution of

$$\dot{r}(t) = (I_\nu \otimes A_0)r(t), \quad r(0) = \lim_{k\to\infty}(\eta[k]-\frac{1}{\nu}\theta[k]),$$

but won't necessarily *synchronize*. In fact, $\bar{\mu}$ will synchronize to a solution of $\mathrm{e}^{A_0 t}r_0$ if and only if

$$\lim_{k\to\infty}(\eta[k]-\theta[k]) = \tilde{r}_0 = \mathbb{1}_\nu \otimes r_0$$

for some constant vector r_0. Since $\eta[k]$ is driven only by the initial conditions, we know that if $\boldsymbol{\mathcal{A}_1}$ holds it will converge to the agreement space. Thus we must prove two things: i) that $\theta[k]$ converges to a finite limit, and ii) that it converges to the agreement set. The following lemma provides conditions for the first requirement.

Lemma 3.6.2. *Let λ_i denote the eigenvalues of A_0 and μ_i the eigenvalues of $A+LC$. If L is chosen such that $\Re(\mu_i-\lambda_j)<0$ for all i,j and $\boldsymbol{\mathcal{A}_1}$ holds, then $\theta[k]$ converges absolutely to a finite limit.*

Proof First, note $\epsilon(t)$ is an exponentially stable continuous time LTI system, hence

$$\lim_{t\to\infty} \|\epsilon(t)\| = 0 \quad \text{and} \quad \exists \rho > 0 \ \text{ s.t } \ \|\epsilon(t)\| \le \rho \ \forall t \ge 0.$$

Next, by definition the Laplacian matrix, $L_{\mathcal{G}}[k]$, has row sum 0, non-negative diagonal entries bounded by $\nu - 1$, and either 0 or -1 on the off diagonal. Hence, $M[k]$ is non-negative and with row sum of 1, making it row-stochastic for all k. Since it is row-stochastic, by Proposition B.2.2 $\left\|\prod_{j=l}^{k} M[j]\right\|_2 \le \sqrt{\nu}$ for all l, k. When combined, we obtain

$$\|\theta[k]\| \le \sqrt{\nu}\rho \left\| \sum_{l=1}^{k} - \int_{s_{l-1}}^{s_l} \mathrm{e}^{-A_0 s} LC \mathrm{e}^{A_L s} \mathrm{d}s \right\| = \sqrt{\nu}\rho \left\| - \int_{0}^{s_l} \mathrm{e}^{-A_0 s} LC \mathrm{e}^{A_L s} \mathrm{d}s \right\|.$$

Next, note that if L is chosen such that $\Re(\mu_i - \lambda_j) < 0$ for all i, j, then

$$Q = \int_0^\infty \mathrm{e}^{-A_0 s} LC \mathrm{e}^{A_L s} \mathrm{d}s$$

exists and is the solution of the Sylvester equation

$$-A_0 Q + Q A_L + LC = 0,$$

since

$$-A_0 Q + Q A_L = \int_0^\infty \frac{\mathrm{d}}{\mathrm{d}s} \left(\mathrm{e}^{-A_0 s} LC \mathrm{e}^{A_L s} \right) \mathrm{d}s = -LC.$$

In particular, this shows that in the limit $\|\theta[k]\|$ is bounded by

$$\lim_{k\to\infty} \|\theta[k]\| \le \sqrt{\nu}\rho \, \|Q\|_2 .$$

Moreover, since the series of integrals converge they are bounded and there exists some constant $\beta > 0$ such that for any k

$$\|\theta[k]\| \le \sqrt{\nu}\beta \sum_{l=1}^{k} \|\epsilon(s_{l-1})\| .$$

Since $\|\epsilon(s_{l-1})\| \to 0$ exponentially, the series defined by $\|\theta[k]\|$ is positive and bounded, hence converges to a finite limit. Since $\|\theta[k]\|$ converges, this means that $\theta[k]$ converges absolutely. ■

Note that the condition $\Re(\mu_i - \lambda_j) < 0$ is slightly stronger then the one required for the analog case. This is because now we need to ensure that the integral converges, not just that the algebraic Sylvester equation has a solution. This is likely just conservatism due to the technical machinery of the proof.

Now that we have established that $\theta[k]$ converges, we can show that it, and the states, converge to the agreement space.

Theorem 3.4. *If F_d and L are such that $A + BF_d$ and $A + LC$ are Hurwitz and $\bar{F}$ is such that $A_0 = A + B\bar{F}$, and the conditions of Lemma 3.6.2 are satisfied, then the local hybrid controllers*

$$\begin{cases} \dot{\hat{x}}_i(t) = A\hat{x}_i(t) + Bu_i(t) + L(C\hat{x}_i(t) - y_i(t)) \\ \dot{\bar{\mu}}_i(t) = A_0\bar{\mu}_i(t) + \frac{1}{\nu}L(C\hat{x}_i(t) - y_i(t)) \\ \bar{\mu}_i(s_k^+) = \bar{\mu}_i(s_k) - \frac{1}{\nu}\sum_{l\in\mathcal{N}_i[k]}(\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)) \\ u_i(t) = F_d\hat{x}_i(t) + (\bar{F} - F_d)\bar{\mu}_i(t) \end{cases}, \tag{3.20}$$

solves $\boldsymbol{\mathcal{P}_2}$ for any sampling sequence $\{s_k\}$ satisfying $\boldsymbol{\mathcal{A}_1}$.

Proof Note that

$$\theta[k+1] = \left(M[k+1]\otimes I_n\right)\left(\theta[k] + \left(I_\nu \otimes \int_{s_k}^{s_{k+1}} \mathrm{e}^{-A_0\tau}LC\mathrm{e}^{A_L\tau}\mathrm{d}\tau\right)\hat{\epsilon}(s_k)\right)$$

and that

$$\theta[n+p] = \left(\prod_{j=n+1}^{n+p} M[j]\otimes I_n\right)\theta[n] + \sum_{l=n+1}^{n+p}\left(\prod_{j=l}^{k} M[j]\otimes\left(\int_{s_{l-1}}^{s_l}\mathrm{e}^{-A_0\tau}LC\mathrm{e}^{A_L\tau}\mathrm{d}\tau\right)\right)\hat{\epsilon}(s_{l-1}).$$

Since $\theta[k]$ converges absolutely, by Cauchy's convergence test for all $\epsilon > 0$ there exists some $N \in \mathbb{N}$ such that for all $n > N$ and $p \geq 1$

$$\left\|\sum_{l=n+1}^{n+p}\left(\prod_{j=l}^{k} M[j]\otimes\left(\int_{s_{l-1}}^{s_l}\mathrm{e}^{-A_0\tau}LC\mathrm{e}^{A_L\tau}\mathrm{d}\tau\right)\right)\hat{\epsilon}(s_{l-1})\right\| < \epsilon,$$

which in turn implies that for all $\epsilon > 0$ there exists some $N \in \mathbb{N}$ such that

$$\left\|\theta[N+p] - \left(\prod_{j=N+1}^{N+p} M[j]\otimes I_n\right)\theta[N]\right\| \leq \epsilon$$

for all $p \geq 1$.

Since $\boldsymbol{\mathcal{A}_1}$ holds, we know that for any finite N [21, Lem. 2.29 and 2.30] there exists some constant $q_N \in \mathbb{R}^\nu$ such that

$$\lim_{p\to\infty}\prod_{j=N+1}^{N+p} M[j] = \mathbb{1}q_N',$$

and since a limit is unique, this implies that $\theta[k] \to \mathrm{Im}\,\mathbb{1}\otimes I$. As we know $\eta[k]$ converges as well, we have from (3.19) that

$$\lim_{k\to\infty}\left\|\bar{\mu}(s_k+\tau) - \mathbb{1}\otimes\mathrm{e}^{A_0(s_k+\tau)}r_0\right\| = 0$$

where r_0 is defined by $\lim_{k\to\infty} \eta[k] - \frac{1}{\nu}\theta[k] = \mathbb{1} \otimes r_0$.

The final step is to show that $\varepsilon(t) = x(t) - \bar{\mu} \to 0$. To this end, consider first the auxiliary variable $\hat{\epsilon} = \hat{x} - \bar{\mu}$, with hybrid dynamics

$$\begin{cases} \dot{\hat{\epsilon}}(t) = (I_\nu \otimes (A + BF_d)\hat{\epsilon}(t) - (1 - \frac{1}{\nu})(I_\nu \otimes LC)\epsilon(t) \\ \hat{\epsilon}(s_k^+) = \hat{\epsilon}(s_k) + (\frac{1}{\nu}L_{\mathcal{G}}[k] \otimes I_n)\bar{\mu}(s_k) \end{cases}.$$

This is an LTI system with two inputs: a continuous $\epsilon(t)$ and impulsive $\bar{\mu}(s_k)$. As we have seen, $\epsilon(t) \to 0$ and $\bar{\mu}(s_k) \to \operatorname{Im} \mathbb{1} \otimes I$ which implies that $\{(L_{\mathcal{G}}[k] \otimes I_n)\bar{\mu}(s_k)\} \to 0$ as well. Since $A+BF_d$ is Hurwitz, we have a stable LTI system with two asymptotically vanishing inputs, thus $\hat{\epsilon}(t) \to 0$. No Note that $\varepsilon = \hat{\epsilon} + \epsilon$, hence we have

$$\lim_{t\to\infty} \|\varepsilon(t)\| = \lim_{t\to\infty} \|\hat{\epsilon}(t) + \epsilon(t)\| = 0$$

and the agents track $\bar{\mu}$ asymptotically, implying that

$$\lim_{t\to\infty} \left\| x(t) - \mathbb{1} \otimes \mathrm{e}^{A_0(s_k+\tau)} r_0 \right\| = 0,$$

hence $\boldsymbol{\mathcal{P}_2}$ is solved. ■

Curiously, (3.16) consistently outperformed (3.16′) in simulation despite the latter having continuous feedback from the agents. This is subject to current research.

3.7 Illustrative examples

To illustrate the proposed sampled-data protocol, consider a simple system comprised of $\nu = 3$ agents described by (3.1) with

$$\left[\begin{array}{c|c} A & B \\ \hline C & \end{array}\right] = \left[\begin{array}{cc|c} 4 & 9 & 2 \\ 1 & 4 & 1 \\ \hline \multicolumn{2}{c|}{C} & \end{array}\right]$$

where $C = I$ in the state feedback example, and $C = \begin{bmatrix} 1 & 0 \end{bmatrix}$ in the output feedback one. The goal is to synchronize to

$$A_0 = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix},$$

which corresponds to harmonic oscillations with the frequency 1 rad/sec, note that $\boldsymbol{\mathcal{A}_4}$ is satisfied via

$$\bar{F} = -\begin{bmatrix} 2 & 4 \end{bmatrix}.$$

We assume that communication between agents is intermittent and asynchronous, meaning that each agent transmits only at a subset of sampling instances. At each

sampling instance the connectivity graph $\mathcal{G}[k]$ is a union of any nonempty combination of the three graphs in Figure 3.2. The second condition of assumption $\boldsymbol{\mathcal{A}_1}$ is equivalent

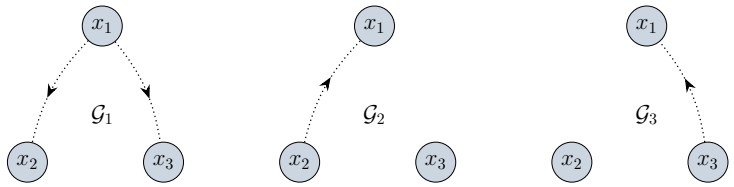

Figure 3.2: The three possible graphs for the examples in Section 3.7.

in this case to the existence of a subsequence of sampling instances at which $\mathcal{G}[k]$ contains $\mathcal{G}_1$.

3.7.1 State feedback

For the state feedback case we only require a stabilizing gain to satisfy the requirements of Theorem 3.1, for example

$$F_{\mathrm{d}} = \begin{bmatrix} -34.6 & 39.2 \end{bmatrix}.$$

The system is simulated for time interval $t \in [0, 30]$, the results of which are shown in Figure 3.3 and Figure 3.4. The sampling instances, shown by abscissa ticks, are a random variable such that $s_{k+1} - s_k \in 0.3\,\mathbb{N}_6$. Major ticks indicate the sub-sequence of sampling instances $\{k_p\}$ satisfying $\boldsymbol{\mathcal{A}_1}$. The synchronous trajectory as defined in Lemma 3.3.2 is plotted in lavender.

Figure 3.3 presents the time evolution of the agents states. It can be seen that each component of the state converges to a common trajectory as stated in Theorem 3.1. During the transients the state trajectories display spikes at time instances each agent receives new information due to the discontinuous jump of the corresponding $\bar{\mu}_i$. These spikes become progressively smaller the closer the emulated centroids are to one another, which would be explained in Chapter 4.

A second notable thing is that there is no general counterpart to Lemma 2.3.4, thus neither the emulated centroids not the actual states are uniformally bounded by the convex hull of the initial conditions. This is illustrated in the second coordinate of Figure 3.3, where both the yellow and red agents overshoot the initial conditions at the beginning.

Figure 3.4(a) shows the decay of the emulator disagreement norm, namely $\|\Delta_i(t)\|$, on a logarithmic scale. We can see that $\|\Delta_i(t)\|$ is not monotonically decreasing, which is due to the hybrid nature of the system. The signals $\|\Delta_i(t)\|$ sharply decrease between samples, but might jump up at $t = s_k$, when new information is brought in. Still, there are exponentially decreasing functions upperbounding the combined disagreements norms. The phase portrait of $\bar{\mu}_i(t)$ is given in Figure 3.4(b) and displays similar discontinuous behaviour, where each centroid sharply changes its trajectory when the emulators are updated, until they all converge to a common trajectory.

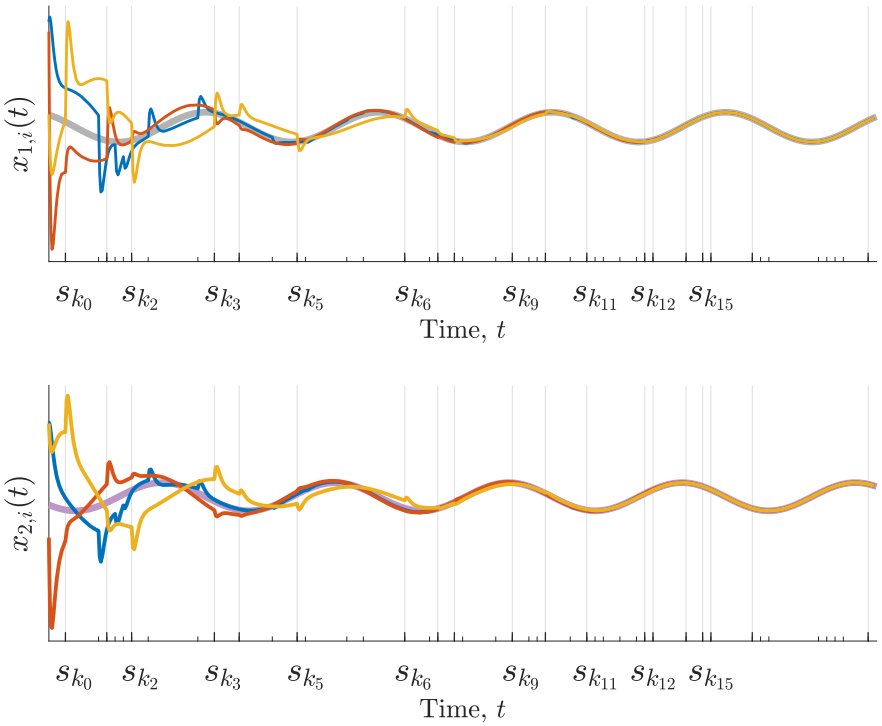

Figure 3.3: Evolution of agent states, $x_{1,i}$, $x_{2,i}$, and the synchronous trajectory for the example in Subsection 3.7.1. Minor ticks are sampling instances, major ticks indicate the connected subsequence from $\boldsymbol{\mathcal{A}_1}$.

3.7.2 Output feedback

For output measurements only, we must complement our state feedback with an observer gain. A good rule of thumb is to have the observer poles not too aggressive but still faster than the feedback poles. In the previous example we opted for very aggressive feedback gain, here we opt for more conservative choices

$$F_{\mathrm{d}} = -\begin{bmatrix} 7 & 1 \end{bmatrix} \quad \text{and} \quad L = -\begin{bmatrix} 19 \\ 11 \end{bmatrix},$$

which satisfy the requirements of Theorem 3.3, assigning the spectrum of A_{d} to $\{-3, -4\}$ and that of $A+LC$ to $\{-5, -6\}$. The simulation results, carried out over the time interval $t \in [0, 35]$, are presented in Figure 3.6. The sampling instances, shown by abscissa ticks, are a random variable such that $s_{k+1} - s_k \in 0.45\,\mathbb{N}_5$. Major ticks indicate the sampling instances at which $\boldsymbol{\mathcal{A}_1}$ is satisfied.

Figure 3.5(a) presents the time evolution of the agents states. It can be seen that each component of the state converges to a common trajectory solving $\boldsymbol{\mathcal{P}_2}$. Figure 3.5(b) portrays the time evolution of the emulated centroid states, while the real centroid, $\bar{x}(t)$, is plotted in dashed lavender line. This is to be expected, as the agents approach synchronization the only non-zero component of their state is the centroids. Coupled with the fact that Theorem 3.3 established that $\varepsilon(t) \to 0$ as $t \to \infty$, this

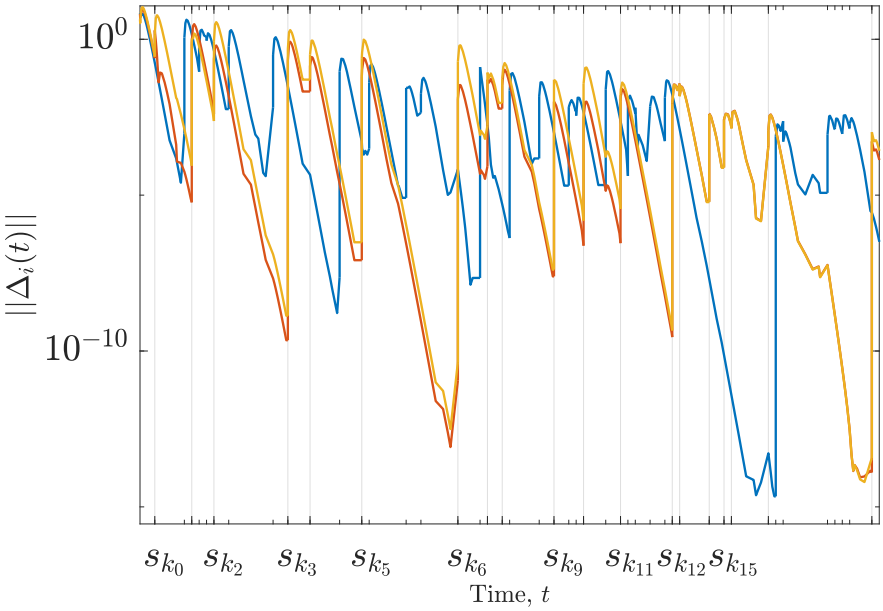

(a) Norm of emulator disagreements (logarithmic scale).

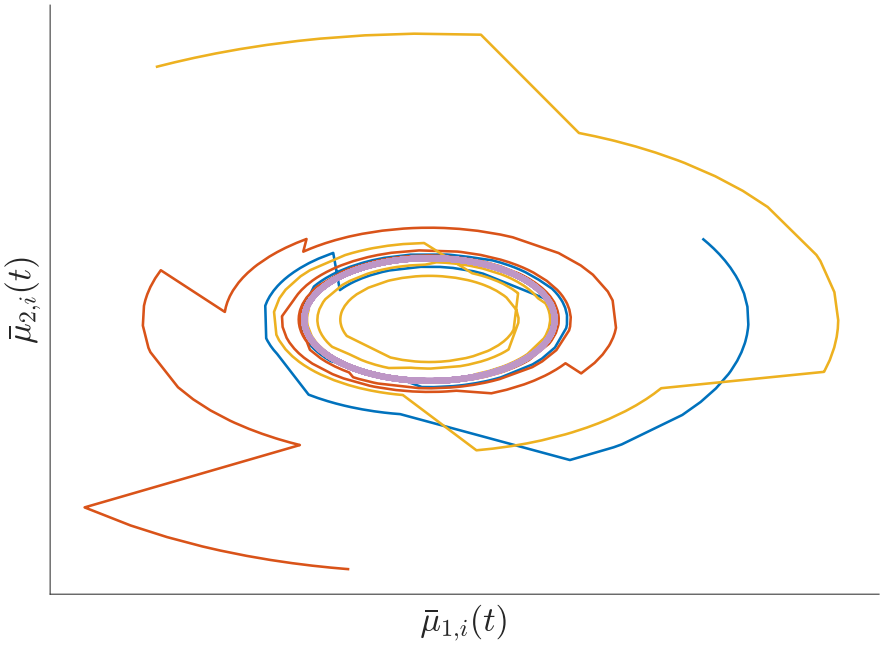

(b) Phase portrait of $\bar{\mu}_i(t)$.

Figure 3.4: The emulator's behavior for the example in Subsection 3.7.1: norm of emulated disagreements, and phase portrait of the centroids and the synchronous trajectory.

indicates that $\bar{\mu}(t) - \bar{x}(t) \to 0$ for all $i \in \mathbb{N}_3$.

Figure 3.6(a) shows the norm of the components of ε, i.e. the signals $x_i - \bar{\mu}_i$ for $i \in \mathbb{N}_3$, on a logarithmic scale. When no information arrives, these signals decay exponentially fast because each agents tracks the local emulated centroid $\bar{\mu}_i$. When new information about neighboring agents is received, each $\bar{\mu}_i$ updates, as the centroids

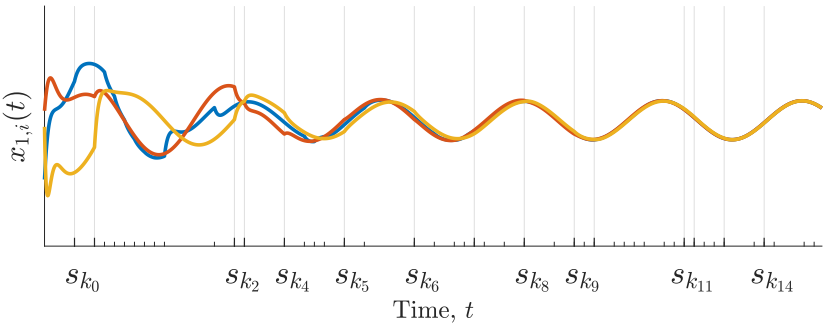

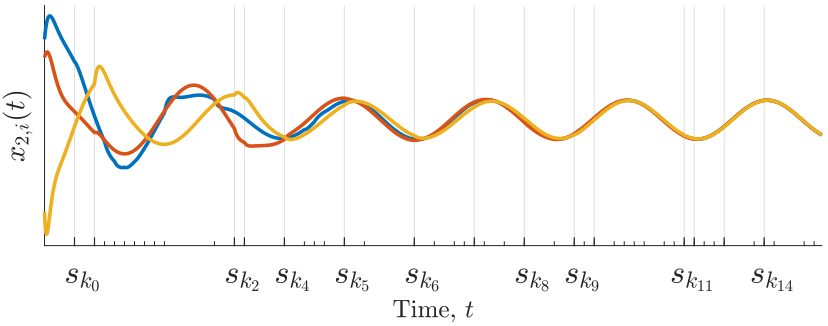

(a) Evolution of agent states.

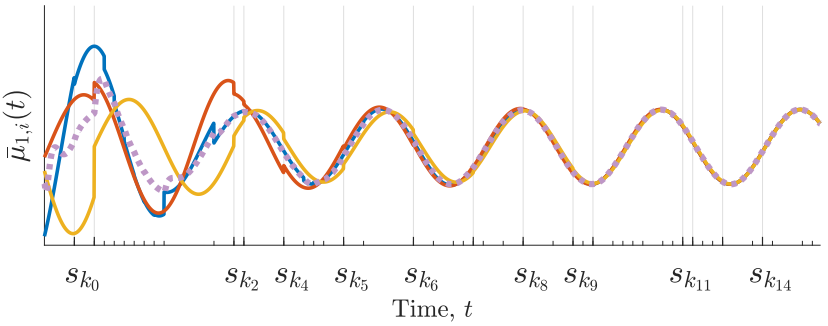

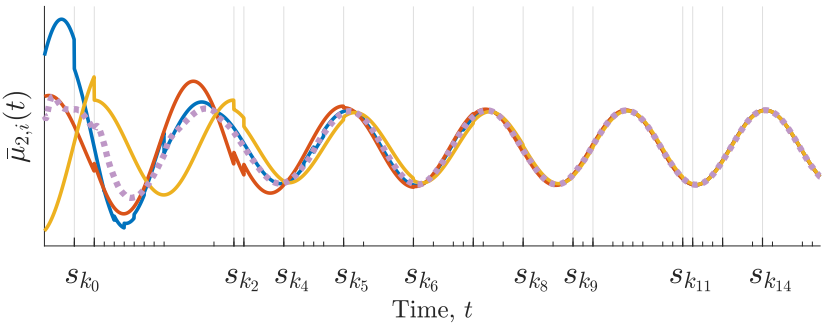

(b) Evolution of centroid states.

Figure 3.5: Evolution of agent states, centroids and synchronous trajectory (dotted) for the example in Subsection 3.7.2 w/ output measurements. Minor ticks are sampling instances, major ticks indicate the connected subsequence from $\boldsymbol{\mathcal{A}_1}$.

are drawn together by the jump map. This normally increases $\|x_i - \bar{\mu}_i\|$, for the local target jumps. Yet at the same time these targets at communicating agents approach each other, which is required to satisfy (3.5b).

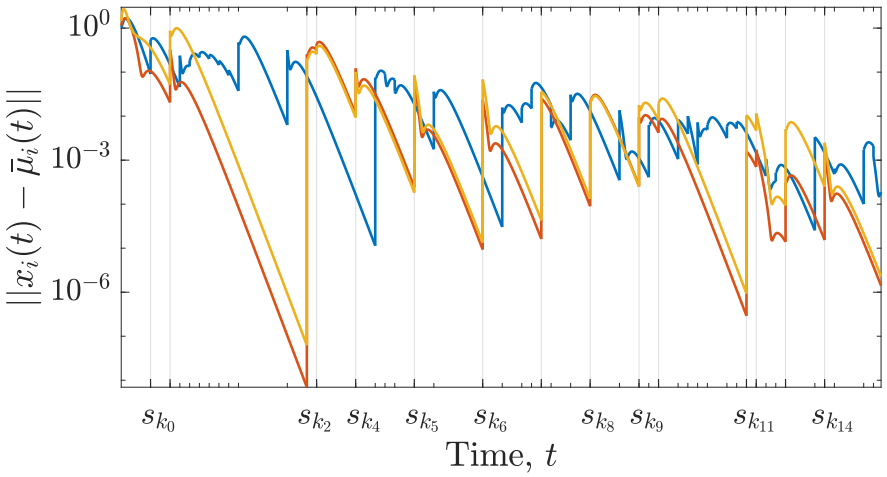

(a) Norm of the components of tracking error $\boldsymbol{\varepsilon}$ (logarithmic scale).

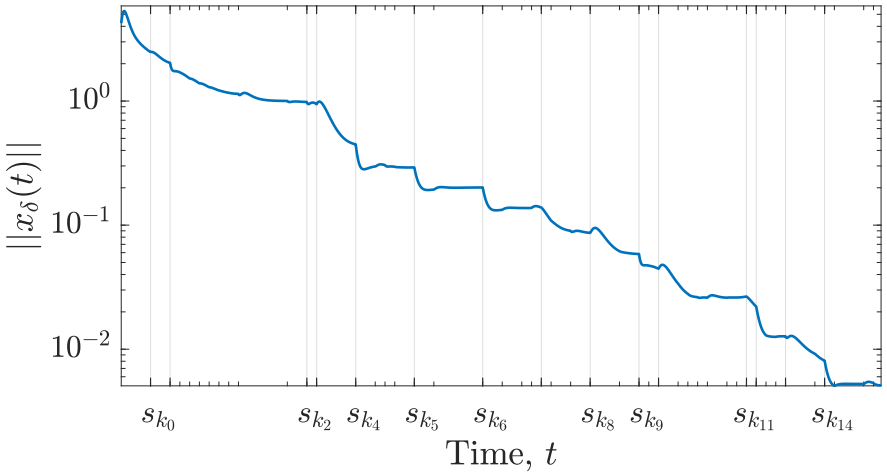

(b) Norm of the aggregate state disagreement (logarithmic scale).

Figure 3.6: Norms of tracking error and state disagreement w/ output measurements. Minor ticks are sampling instances, major ticks indicate the connected subsequence from $\boldsymbol{\mathcal{A}_1}$.

Finally, Figure 3.6(b) depicts the norm of x_δ on a logarithmic scale. In contrast to the components of $\boldsymbol{\varepsilon}$ from Figure 3.6(a), the quantity in Figure 3.6(b) decreases when new information is received. This behaviour indicates that the agent disagreements consistently decrease during information exchange, as required to satisfy condition (3.5a).

3.8 Concluding remarks

In this chapter, we addressed time-varying state synchronization of general LTI agents under complex communication constraints. The synchronization is not limited to trajectories generated by the open-loop dynamics, but rather to any dynamics reachable by local state-feedback. We were able to guarantee global asymptotic agreement under mild assumptions on the persistent connectivity of the graphs and sampling instances.

Moreover, the control parameters are independent of both the sampling sequence and spatial graphs. These properties are facilitated by a separation between the control law and the information processing mechanism, hinting at possible extensions to more general setups. In particular, extensions to output feedback, disturbance rejection, and systems affected by delays are currently being considered.

Chapter 4

Extensions to the Emulation Scheme

More wittinies, less math

Lewis Carrol, *Alice in Wonderland* [64]

In this chapter, we shall show several extensions and properties regarding the emulation scheme developed in Chapter 3. These include convergence analysis and conditions for exponential convergence, addition of time-varying weights to the update map, and augmentation of the update map with a predictor. The prediction scheme is designed to counter time-varying and heterogeneous transmission delays assuming that they are smaller then the sampling interval, and was published in [62].

4.1 Convergence rates and exponential convergence

Our first result concerns the convergence rate of the emulation scheme. Without temporal constraints, consensus-based synchronization algorithms converge exponentially, as their disagreement dynamics are stable LTI systems. Similarly, under the assumption of synchronous and periodic sampling, some discretizations method can also ensure exponential convergence since the equivalent discrete system is linear shift-invariant and stable. There are some results about exponential agreements of more general systems, but they still assume constant undirected graphs, synchronous sampling, and require solutions of complicated LMIs [78]. When the graphs are directed and time-varying, which is equivalent to asynchronous sampling, convergence is in general only asymptotic.

In hybrid systems where the flow and jump and not both contracting, even analyzing the convergence rate is non-trivial. One way to do so is by considering its *set distance* from the agreement space. Clearly, this distance is not monotonically decreasing at each sampling instance. However, it is indeed monotone along the subsequence $\{s_{k_p}\}$ defined in $\mathcal{A}_1$. This is formulated in the following proposition.

Proposition 4.1.1. *If the conditions of Theorem 3.1 hold, then*

$$d\left(\bar{\mu}(s_{k_{p+1}}), \operatorname{Im} \mathbb{1}_\nu \otimes \mathrm{e}^{A_0 t}\right) < d\left(\bar{\mu}(s_{k_p}), \operatorname{Im} \mathbb{1}_\nu \otimes \mathrm{e}^{A_0 t}\right)$$

where $\bar{\mu}(s_k)$ is the aggregation of (3.11b) *and $d(\cdot,\cdot)$ is the set (Hausdorff) distance.*

Proof The proof follows directly from applying the set-valued Lyapunov function of [31] to the aggregate solution of (3.11b) along the subsequence defined in $\boldsymbol{\mathcal{A}_1}$. ■

Proposition 4.1.1 is the reason the "spikes" at update times in x_i become progressively smaller, as noted in Subsection 3.7.1. The centroid updates are driven by matrices of the form $I - \frac{1}{\nu} L_{\mathcal{G}}[k]$, and by design for any vector x

$$(I - \frac{1}{\nu} L_{\mathcal{G}}[k])(x_\delta + (\mathbb{1} \otimes I)\bar{x}) = (I - \frac{1}{\nu} L_{\mathcal{G}}[k]) x_\delta.$$

Hence the closer the centroids are to agreement, the less affected they are by the jumps.

Recently, a new result proposed a special time-varying quadratic Lyapunov function that exploited the graph structure to ensure exponential convergence for time-varying graphs [79]. This result required a slightly stronger assumption than $\boldsymbol{\mathcal{A}_1}$, which is given below.

$\boldsymbol{\mathcal{A}_5}$: there is a strictly increasing sub-sequence of sampling indices $\{k_p\}$ such that for all $p \in \mathbb{Z}_+$ (i) the intervals $s_{k_{p+1}} - s_{k_p}$ are uniformly bounded and (ii) $\cup_{k=k_p+1}^{k_{p+1}} \mathcal{G}[k]$ is *strongly connected*.

Assumption $\boldsymbol{\mathcal{A}_5}$ requires that the union graphs are not only connected, but strongly connected. Below is a variation of [79, Thm. 1] using this assumption and a novel time-varying Lyapunov function.

Proposition 4.1.2. *Consider the discrete-time consensus protocol*

$$x[k+1] = (I - cL_{\mathcal{G}}[k])x[k] \coloneqq M[k]x[k]. \tag{4.1}$$

If $c \in \left(0, 1/\sup_k \max_{j \in \mathbb{N}_\nu} \sum_{i=1}^{\nu} m_{ij}[k]\right)$ and $\boldsymbol{\mathcal{A}_5}$ holds, then (4.1) *converges exponentially to consensus along the sequence $\{s_{k_p}\}$.*

Remark 4.1. Please note that Proposition 4.1.2 was modified to fit with our definition of the adjacency, and hence Laplacian, matrix. ▽

Proposition 4.1.2 guarantees that if $\boldsymbol{\mathcal{A}_5}$ discrete-time integrators would converge exponentially to consensus, provided that the coupling gain is sufficiently small. This result immediately applies to Theorem 2.1, which concerns integrator consensus.

Lemma 4.1.3. *If the conditions of Theorem 2.1 hold and $\boldsymbol{\mathcal{A}_5}$ replaces $\boldsymbol{\mathcal{A}_1}$, then* (2.10) *reaches consensus exponentially along the sequence $\{s_{k_p}\}$.*

Proof Recall that the centroid dynamics (2.10) are purely discrete and in-fact identical to those in (4.1). By construction, the maximal possible column sum of $L_{\mathcal{G}}[k]$ occurs when a node is an in-neighbor of all other nodes, and out-neighbor of none. For this case the column sum if $\nu - 1$, which is exactly the diagonal entry of $L_{\mathcal{G}}[k]$ for that column. The minimal column sum occurs in the dual case, and is then $-(\nu - 1)$. Consequently, for all k

$$1 - \frac{\nu - 1}{\nu} = \frac{1}{\nu} \leq \sum_{i=1}^{\nu} 1 - \frac{1}{\nu}\left[L_{\mathcal{G}}[k]\right]_{ij} \leq \frac{2\nu - 1}{\nu}$$

and

$$\nu \leq \frac{1}{1 - (1/\nu)\sup_k \max_{j\in\mathbb{N}_\nu} \sum_{i=1}^{\nu} \left[L_{\mathcal{G}}[k]\right]_{ij}} \leq \frac{2\nu - 1}{\nu} = 2 - \frac{1}{\nu}.$$

Clearly $c = 1/\nu$ satisfies the assumption in Proposition 4.1.2 for all $\nu \geq 1$, hence $\bar{\mu}(s_k)$ reach consensus exponentially fast along the subsequence $\{s_{k_p}\}$. ∎

In general, agreement conditions couple the dynamics, controllers, sampling (or switching sequence), and the graph structure. For simple integrators and consensus the dynamics are determined by the control, which is often just a scaled version of the Laplacian. Hence, convergence analysis can be reduced to some graph-theoretical conditions with some simple condition on the gain. Beyond this simple scenario the analysis can become exponentially more complex, cf. the double-integrator case in [80]. Hence, the graph-theoretic Lyapunov function used in the proof of Proposition 4.1.2 does not easily extend to general LTI agents.

However, the control structure described in Chapter 3 has two special properties: (i) the centroids are decoupled from the rest of the dynamics, and (ii) the update map (2.5′) follows *integrator consensus dynamics*. This is reoccurring in all the variations presented, along with the fact that the centroid dynamics upper bound the convergence rate of the actual states. In addition, in the proof of Lemma 3.3.2 we have seen that the flow dynamics of the centroids do not interfere with the jumps. Therefore, the equations governing the convergence rate of the states are essentially those of discrete integrators despite the general framework. This is the key feature which allows us to state the following result.

Theorem 4.1. *If the conditions of Theorem 3.1 hold and* $\mathcal{A}_5$ *replaces* $\mathcal{A}_1$*, then* (3.11b) *synchronizes exponentially along the sequence* $\{s_{k_p}\}$*.*

Proof Consider the coordinate transformation $z(t) \coloneqq (I_\nu \otimes \mathrm{e}^{-A_0 t})\bar{\mu}(t)$ (the same one from [42]). The dynamics of $z(t)$ are then given by

$$\begin{cases} \dot{z}(t) = 0 \\ z(s_k^+) = ((I_\nu - \frac{1}{\nu} L_{\mathcal{G}}[k]) \otimes I_n) z(s_k) \end{cases}$$

which are identical to those in considered in Lemma 4.1.3. This implies that there

exists constants $\delta_1, \delta_2 > 0$ and a constant vector $r^* \in \mathbb{R}^n$ such that for every $t \in \{s_{k_p}\}$

$$\|z(t) - \mathbb{1}_\nu \otimes r^*\| \leq \delta_1 \mathrm{e}^{-\delta_2 t} \|z(0) - \mathbb{1}_\nu \otimes r^*\| .$$

In the original coordinates this implies that

$$\left\|\bar{\mu}(t) - (\mathbb{1}_\nu \otimes \mathrm{e}^{A_0 t})(I_n \otimes r^*)\right\| \leq \left\|I_\nu \otimes \mathrm{e}^{A_0 t}\right\| \delta_1 \mathrm{e}^{-\delta_2 t} \|\bar{\mu}(0) - \mathbb{1}_\nu \otimes r^*\| .$$

Since all the unstable eigenvalue of A_0 are on the jω axis, this implies that its norm can diverge only polynomially fast, and thus still dominated by the convergence to agreement . Therefore there are constants $\delta_3, \delta_4 > 0$ such that

$$\left\|\bar{\mu}(t) - (\mathbb{1}_\nu \otimes \mathrm{e}^{A_0 t})(I_n \otimes r^*)\right\| \leq \delta_3 \mathrm{e}^{-\delta_4 t} \|\bar{\mu}(0) - \mathbb{1}_\nu \otimes r^*\| .$$

for all $t \in \{s_{k_p}\}$. ■

Note that in the proof we no longer assumed that the eigenvalues of A_0 on the jω axis are semi-simple as stated in $\boldsymbol{\mathcal{P}_2}$. This assumption was required in the proof of Lemma 3.3.2 to ensure that $\mathrm{e}^{A_0 t}$ is bounded, otherwise it can diverge at a polynomial rate. However, if the jump map converges exponentially it will always dominate the diverging flow. Indeed, if $\boldsymbol{\mathcal{A}_5}$ replaces $\boldsymbol{\mathcal{A}_1}$, this assumptions is no longer required in $\boldsymbol{\mathcal{P}_2}$.

It should be clear that $x(t)$ also synchronizes exponentially, albeit at a slightly different rate, since the tracking error evolves according to

$$\begin{cases} \dot{\varepsilon}(t) = I \otimes (A + BF_{\mathrm{d}})\varepsilon(t) \\ \varepsilon(s_k^+) = \varepsilon(s_k) - \dfrac{1}{\nu}(L_{\mathcal{G}}[k] \otimes I)\bar{\mu}(s_k) \end{cases} .$$

As before the flow is exponentially stable, but now impulsive input $\{(L_{\mathcal{G}}[k] \otimes I)\bar{\mu}(s_k)\}$ also vanishes exponentially. Hence, $x_i(t)$ will synchronize at an exponential rate upper bounded by occurrences of subsequence $\{s_{k_p}\}$ and the convergence rate of $A + BF_{\mathrm{d}}$. A similar exercise can be done for the output feedback dynamics of (3.16), but now the convergence rate will also be affected by the observer.

Remark 4.2. In both of the examples in Section 3.7 the sampling randomly switched between the three graphs in Figure 3.2. It is easily verified that the union of all three forms a strongly connected graph like the one required by $\boldsymbol{\mathcal{A}_5}$, hence all examples in-fact showed exponential convergence. This is particularly evident in the norm of the disagreements which represent the deviation of each agent from their emulated centroid. For example in Figure 3.4(a) there is a clear linearly decreasing trendline bounding the norms despite the non-monotonicity induced by the jumps. These are plotted on a logarithmic scale, hence the linear trendline indicates exponential decay. ▽

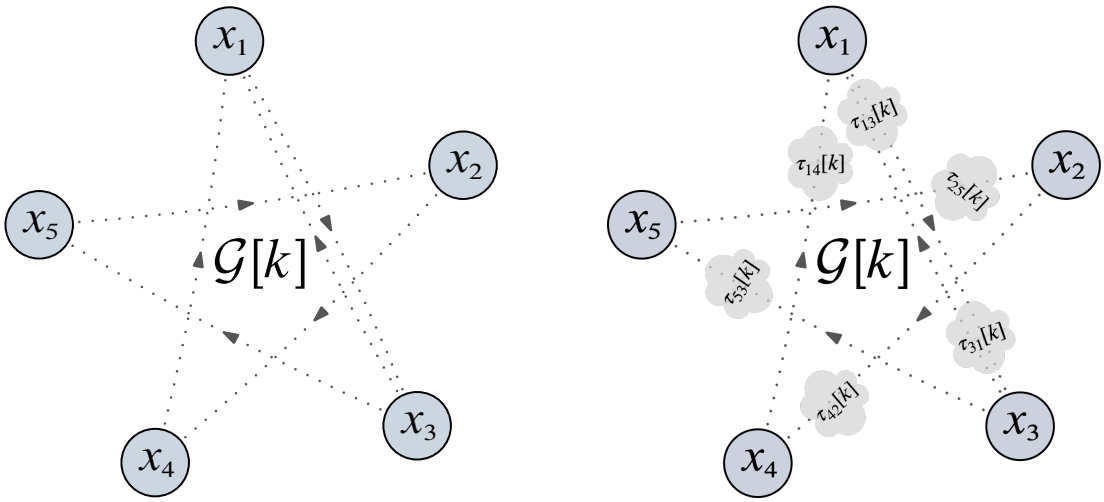

(a) Delay free communication: information is simultaneously sent and received at $t = s_k$.

(b) Communication with transmission delays: information sent from agent i to j at $t = s_k$ is received at $t = s_k + \tau_{ij}[k]$.

Figure 4.1: Sampled-data multi-agent communication with and without transmission delays.

4.2 Transmission delays

Consider $\boldsymbol{\mathcal{P}_2}$ under the same setup and assumptions as before, but with heterogeneous time-varying delays on the *communicated* information between the agents. Denote by $\tau_{ij}[k]$ the transmission delay from agent j to agent i at time instance $t = s_k$, this is visualized in Figure 4.1. We assume that

$\boldsymbol{\mathcal{A}_6}$: incoming information is time stamped and

$$s_k + \tau_{ij}[k] < s_{k+1}, \quad \forall i, j \in \mathbb{N}_\nu,\ k \in \mathbb{Z}_+.$$

The assumption above does not imply that the delays are known a priori, only that the receiving agent knows $\tau_{ij}[k]$ at $t = s_k + \tau_{ij}[k]$. The second part guarantees that there is no packet disorder, which is a reasonable assumption in MAS [49] and networked systems in general [81]. Note that the delays are allowed to vary between sampling instances as well as across communication channels. We seek to modify local controllers (3.13') to solve $\boldsymbol{\mathcal{P}_2}$ for all transmission delays satisfying $\boldsymbol{\mathcal{A}_6}$.

Since the agents interact only at discrete time instances and are decoupled otherwise, the delays modify only the discrete component of (3.13'). Essentially, $\boldsymbol{\mathcal{A}_6}$ splits the delay-free update of agent i at s_k into up to $|\mathcal{N}_i[k]|$ different updates spread over the interval $[s_k, s_{k+1})$ but still verifying

$$\bigcup_{j \in \mathcal{N}_i[k]} \mathcal{N}_i[t_{ij}[k]] = \mathcal{N}_i^{\mathrm{DF}}[k] \ \forall i \in \mathbb{N}_\nu,\ k \in \mathbb{Z}_+ \tag{4.2}$$

where $t_{ij}[k] \coloneqq s_k + \tau_{ij}[k]$ and $\mathcal{N}_i^{\mathrm{DF}}[k]$ denotes the delay-free neighborhood of agent i.

Packet loss 4.3. Note that (4.2) does not preclude the possibility of packet losses, but rather relegates them to the graphs induced by $\{s_k\}$. ▽

We now need to design an update rule

$$\bar{\mu}_i(t_{ij}[k]^+) = \kappa\left(\alpha_i\bar{\mu}_i(t_{ij}[k]) + \sum_{j\in\mathcal{N}_i[t_{ij}[k]]}\alpha_j\bar{\mu}_j(s_k)\right)$$

for some function $\kappa(\cdot)$.

The problem at hand is qualitatively different from the standard delay problems considered in the literature. Continuous-time delays are infinite dimensional systems, and therefore so are predictors used in delay compensation. In the proposed setup, the delays affect continuous information that is sent intermittently and used to update $\bar{\mu}_i$ in a discrete fashion. In discrete time, delays are finite dimensional and occur at discrete steps synchronized with the regular increments of the system. However, $\boldsymbol{\mathcal{A}_6}$ implies that the delayed information arrives and is processed *before* the next global sampling instance. Hence, the delay at hand does not fit into either of the standard descriptions. To understand how to construct a predictor for this hybrid type of delay, we shall first consider the special case of consensus of integrator agents.

4.2.1 Consensus of integrator agents

Consider the special case of first order integrator agents trying to achieve consensus. This corresponds to $\boldsymbol{\mathcal{P}_2}$ with $A = A_0 = 0$, and (3.13') simplifies to

$$\begin{cases}\bar{\mu}_i(s_k^+) = \bar{\mu}_i(s_k) - \dfrac{1}{\nu}\displaystyle\sum_{l\in\mathcal{N}_i[k]}(\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)) \\ u_i(t) = F_{\mathrm{d}}(x_i(t) - \bar{\mu}_i(s_k^+))\end{cases}.$$

The equation above is a generalization of the control law proposed in Subsection 2.3.1, for which the dynamics of $\bar{\mu}_i$ are purely discrete. This significantly simplifies the analysis and, in fact, makes any predictor redundant, as demonstrated in the following proposition.

Proposition 4.2.1. *Consider $\boldsymbol{\mathcal{P}_2}$ with $A = A_0 = 0$ and transmission delays. The control law*

$$\begin{cases}\bar{\mu}_i(t_{ij}[k]^+) = \bar{\mu}_i(t_{ij}[k]) - \dfrac{1}{\nu}\displaystyle\sum_{l\in\mathcal{N}_i[t_{ij}[k]]}(\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)) \\ u_i(t) = F_d(x_i(t) - \bar{\mu}_i(t_{ij}[k]^+))\end{cases} \tag{4.3}$$

will drive the agents asymptotically to consensus for all sampling sequences $\{s_k\}$ satisfying $\boldsymbol{\mathcal{A}_1}$, all time delays $\tau_{ij}[k]$ satisfying $\boldsymbol{\mathcal{A}_6}$, and all gains $F_d < 0$.

Proof Consider an ordered sequence $\{q_l[k]\}$ where each element is defined by

$$\begin{aligned} q_1[k] &= \min_{ij} t_{ij}[k] \\ q_l[k] &= \min_{ij} \{t_{ij}[k]\} \setminus \{q_j[k] : j < l\} \quad l = 2, 3, \ldots, |\mathcal{N}_i[k]|, \end{aligned} \tag{4.4}$$

i.e., the ordered time instances for the interval $[s_k, s_{k+1}]$ in which information arrives. By $\boldsymbol{\mathcal{A}_6}$, $\{q_j[k]\}$ has a finite (possibly different) number of elements for each k. An illustration of the relationship between $\{t_{ij}[k]\}$ and $\{q_l[k]\}$ is given in Figure 4.2. Assume without loss of generality that $q_p[k]$ is the last instance, since $\bar{\mu}_i$ is discrete

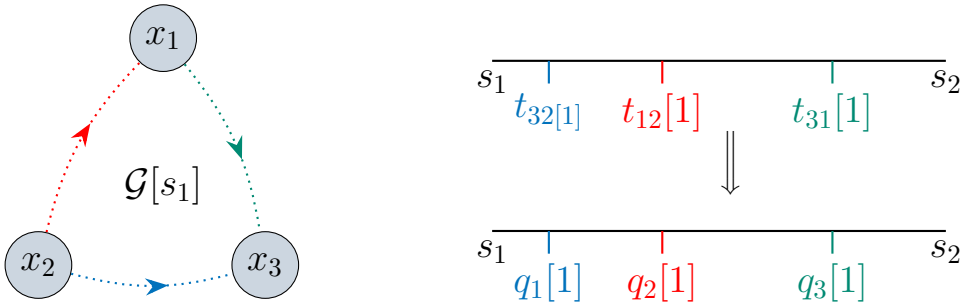

Figure 4.2: An example of the mapping between delayed times $t_{ij}[k]$ and the ordered sequence $q_p[k]$ in (4.4).

this implies that $\bar{\mu}_i(s_{k+1}) = \bar{\mu}_i(q_p[k]^+)$. Expanding the above we have

$$\begin{aligned} \bar{\mu}_i(s_{k+1}) &= \bar{\mu}_i(q_p[k]) - \frac{1}{\nu} \sum_{l \in \mathcal{N}_i[q_p[k]]} (\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)) \\ &= \bar{\mu}_i(s_k) - \frac{1}{\nu} \sum_{r=1}^{p} \sum_{l \in \mathcal{N}_i[q_r[k]]} (\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)) \,. \end{aligned}$$

By $\boldsymbol{\mathcal{A}_6}$ and (4.2), we know that

$$\bigcup_{r=1}^{p} \mathcal{N}_i[q_r[k]] = \mathcal{N}_i^{\mathrm{DF}}[k] \quad \Longrightarrow \quad \bar{\mu}_i(s_{k+1}) = \bar{\mu}_i^{\mathrm{DF}}(s_{k+1}).$$

This is true for all $i \in \mathbb{N}_\nu$ and $k \in \mathbb{Z}_+$. Therefore, if (3.13') will drive the delay-free system to consensus, (4.3) will as well. Note that for $A = A_0 = 0$, $\boldsymbol{\mathcal{A}_4}$ trivially holds, and that if $\boldsymbol{\mathcal{A}_1}$ holds for $\{s_k\}$ then it will also hold for the shifted sequence $\{s_{k+1}\}$. Thus, we can conclude that if the delay-free system will reach agreement for the sequence $\{s_k\}$ and its induced graphs, then the delayed system will for $\{s_{k+1}\}$. ■

Proposition 4.2.1 illustrates how the hybrid nature of the delay can render it redundant in certain cases. Since the updates are event-drive, i.e., an update occurs when new information arrives, and $\bar{\mu}_i$ is constant between updates, the transmission delays only amount to splitting one update into several smaller ones within the same time

interval. When combined with $\mathcal{A}_1$, which considers the union of the induced graphs over some subsequence, it is evident that the delays amount to a partition of the interval $[s_k, s_{k+1}]$ for which $\bigcup \mathcal{G}[t_{ij}[k]] = \mathcal{G}[k]$. Hence, from a consensus standpoint, there is no difference between the original problem and the delayed one. As such, there's no need to predict anything, simply to guarantee that the original sampling intervals remain disjoint, as required in $\mathcal{A}_6$.

The above reasoning does not hold when $\bar{\mu}_i(t)$ is no longer constant between updates, as in the general case of $\mathcal{P}_2$. However, this insight is the guiding principle in designing an appropriate predictor as will be done in the following section.

4.2.2 Synchronization of LTI agents

The key property exploited in Subsection 4.2.1 was that the value of $\bar{\mu}(s_{k+1})$ was the same as it would have been in the delay-free case for all k. In the following lemma, we propose an update rule that will guarantee this property for arbitrary A and A_0 satisfying $\mathcal{A}_4$.

Lemma 4.2.2. *If $\mathcal{A}_{4,6}$ hold and $A + B\bar{F} = A_0$, then under the update rule*

$$\bar{\mu}_i(t_{ij}[k]^+) = \bar{\mu}_i(t_{ij}[k]) - \frac{1}{\nu} e^{A_0 \tau_{ij}[k]} \sum_{j \in \mathcal{N}_i[t_{ij}[k]]} (\bar{\mu}_i(s_k) - \bar{\mu}_j(s_k)), \tag{4.5}$$

we recover the same $\bar{\mu}_i(s_{k+1})$ as in the delay-free system (3.13').

Proof Consider the ordered sequence $\{q_l[k]\}$ from (4.4) for an arbitrary agent with index i, and assume that it receives p delayed updates in the interval $[s_k, s_{k+1})$. Define

$$\theta_i[k,l] \coloneqq \sum_{j \in \mathcal{N}_i[q_l[k]]} (\bar{\mu}_i(s_k) - \bar{\mu}_j(s_k)),$$

to simplify the notation, we shall omit the argument k when it is clear from context or unimportant. Now consider $k = l = 1$, for which the update reads

$$\bar{\mu}_i(q_1^+) = e^{A_0 q_1} \bar{\mu}_{i,0} - \frac{1}{\nu} e^{A_0 \tau_1} \theta[1,1] = e^{A_0 \tau_1} \left(\bar{\mu}(s_1) - \frac{1}{\nu} \theta[1,1] \right),$$

where we used the general fact that $e^{A_0(q_{l+1} - q_l)} e^{A_0 \tau_l} = e^{A_0 \tau_{l+1}}$. From here, by induction

$$\bar{\mu}_i(q_p^+) = e^{A_0 \tau_p} \left(\bar{\mu}(s_1) - \frac{1}{\nu} \sum_{l=1}^{p} \theta[1,l] \right),$$

and once more applying $\mathcal{A}_6$ and (4.2) we obtain that

$$\bar{\mu}_i(q_p^+[1]) = e^{A_0 \tau_p} \bar{\mu}_i^{\mathrm{DF}}(s_1^+) = \bar{\mu}_i^{\mathrm{DF}}(q_p[1]).$$

For an arbitrary k and $l = 1$ the update reads

$$\begin{aligned}\bar{\mu}_i(q_1^+[k]) &= \mathrm{e}^{A_0(q_1[k]-q_p[k-1])}\bar{\mu}_i(q_p[k-1]^+) - \frac{1}{\nu}\mathrm{e}^{A_0\tau_1[k]}\theta[k,1] \\ &= \mathrm{e}^{A_0\tau_1[k]}\left(\bar{\mu}(s_k) - \frac{1}{\nu}\theta[k,1]\right),\end{aligned}$$

where we used the identities

$$\mathrm{e}^{A_0(q_1[k]-q_p[k-1])} = \mathrm{e}^{A_0\tau_1[k]}\mathrm{e}^{A_0(s_k-q_p[k-1])} \quad \text{and} \quad \mathrm{e}^{A_0(s_k-q_p[k-1])}\bar{\mu}_i(q_p[k-1]^+) = \bar{\mu}_i(s_k).$$

From here, by similar arguments, we can conclude that

$$\bar{\mu}_i(q_p^+[k]) = \mathrm{e}^{A_0\tau_p}\bar{\mu}_i^{\mathrm{DF}}(s_k^+) = \bar{\mu}_i^{\mathrm{DF}}(q_p[k]),$$

since there are no updates between $q_p[k]$ and s_{k+1} and the choice of i was arbitrary, the system evolves like its delay-free counterpart (3.13'). ■

One can view the Lemma 4.2.2 from a different angle. Consider the aggregation $\bar{\mu}(t) = \left[\ \bar{\mu}_1(t)', \ldots, \bar{\mu}_\nu(t)'\ \right]'$, then (4.5) in aggregate form is given by

$$\begin{aligned}\bar{\mu}(t_{ij}[k]^+) &= \bar{\mu}(t_{ij}[k]) - \frac{1}{\nu}\left(L_{\mathcal{G}}[t_{ij}[k]] \otimes \mathrm{e}^{A_0\tau_{ij}[k]}\right)\bar{\mu}(s_k) \\ &= \left((I_\nu - \frac{1}{\nu}L_{\mathcal{G}}[t_{ij}[k]]) \otimes I_n\right)\bar{\mu}(t_{ij}[k]),\end{aligned}$$

which is exactly the delay-free update rule for the sampling sequence $\{t_{ij}[k]\}$ instead of $\{s_k\}$. The predictor can be thought of as inducing a new sequence of graphs and sampling instances, whose union over the interval $[s_k, s_{k+1}]$ results in the same induced graph as the original sampling sequence and delay-free update mechanism. This is the key step in the proof of the main result.

Theorem 4.2. *If assumptions $\mathcal{A}_{1,3,4}$ hold and $\bar{F}, F_d$ are chosen such that $A_0 = A + B\bar{F}$ and $A + BF_d$ is Hurwitz, then the controller*

$$\begin{cases} \dot{\bar{\mu}}_i(t) = (A + B\bar{F})\bar{\mu}_i(t), \quad \bar{\mu}_i(0) = \bar{\mu}_{i,0} \\ \bar{\mu}_i(t_{ij}[k]^+) = \bar{\mu}_i(t_{ij}[k]) - \dfrac{1}{\nu}\mathrm{e}^{A_0\tau_{ij}[k]}\displaystyle\sum_{l\in\mathcal{N}_i[t_{ij}[k]]}(\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)) \\ u_i(t) = F_d x_i(t) + (\bar{F} - F_d)\bar{\mu}_i(t) \end{cases} \tag{4.6}$$

solves $\mathcal{P}_2$ for all heterogeneous and time-varying transmission delays satisfying $\mathcal{A}_6$.

Proof Consider first the aggregate delay-free state on the sequence $\{s_k\}$, denoted by $\bar{\mu}^{\mathrm{DF}}(s_k)$. If $\mathcal{A}_1$ holds, we know by Theorem 3.1 that $\bar{\mu}^{\mathrm{DF}}(s_k) \to \mathrm{Im}\,\mathbb{1}_\nu \otimes I_n$, and by Proposition 4.1.1 that it gets closer to that set along the sequence $\{k_p\}$ from $\mathcal{A}_1$. Applying Lemma 4.2.2 we know that $\bar{\mu}_i(s_{k+1}) = \bar{\mu}_i^{\mathrm{DF}}(s_{k+1})$ for all i and all k; hence,

both of them approach $\operatorname{Im} \mathbb{1}_\nu \otimes I_n$ at the same rate. Since the agreement set is an invariant set of both the continuous and discrete dynamics of (4.6), this implies that

$$\lim_{t\to\infty} \left\| \bar{\mu}_i^{\mathrm{DF}}(t) - \bar{\mu}_i(t) \right\| = 0, \quad \forall i \in \mathbb{N}_\nu.$$

Applying Theorem 3.1 implies that

$$\lim_{t\to\infty} \|\bar{\mu}_i(t) - x_i(t)\| = 0,$$

hence the states synchronize. ■

4.2.3 Implementability

Local controllers (4.6) are independent of the size of the system; however, update rule (4.5) makes use of $\bar{\mu}_i(s_k)$. Since the sequence $\{s_k\}$ is not assumed to be known, an immediate question arises as to whether the control law can be implemented. The following proposition states that it can be using a small buffer, and details how to update this buffer accordingly.

Proposition 4.2.3. *The update law* (4.5) *can be implemented using a buffer of size* 1.

Proof Each agent constructs its buffer as follows. Let $\begin{bmatrix} b_i' & t_i' \end{bmatrix}'$ denote the values of the ith buffer and corresponding timestamp, and denote by $t_i[k]$ the instance at which information is received and by s_k the time when it was sent.

1. If $t_i[k] = s_k$, assign
$$\begin{bmatrix} b_i \\ t_i \end{bmatrix} = \begin{bmatrix} \bar{\mu}_i(s_k) \\ s_k \end{bmatrix}.$$
2. If $s_k < t_i[k]$, check
   (a) If $t_i = s_k$, keep the current buffer.
   (b) If $t_i < s_k$ assign
$$\begin{bmatrix} b_i \\ t_i \end{bmatrix} = \begin{bmatrix} \mathrm{e}^{A_0(s_k - t_i[k])}\bar{\mu}_i(t_i[k]) \\ s_k \end{bmatrix}.$$

From $\mathcal{A}_6$ we know that if $t_i = s_k$ then we are still in the interval (s_k, s_{k+1}); hence, we need to keep the start of the interval in the buffer. Similarly, if $t_i < s_k$, this means that our buffer corresponds to the previous interval. Thus, there were no jumps in $[s_k, t_i[k]]$ and we can reconstruct $\bar{\mu}_i(s_k)$ like we would for a regular LTI system. ■

4.2.4 Numerical examples

To illustrate the proposed sampled-data protocol, consider two cases, both comprised of $\nu = 3$ identical agents. We assume that communication between agents is intermittent and asynchronous, meaning that each agent transmits only at a subset of sampling

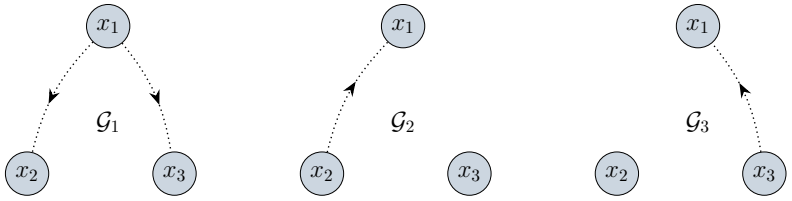

Figure 4.3: The three possible graphs for the examples in Subsection 4.2.4.

instances. At each sampling instance $\mathcal{G}[k]$ is a union of any nonempty combination of the three graphs in Figure 4.3. The sampling instances, shown by abscissa ticks on the bottom, are a random variable such that $s_{k+1} - s_k \in 0.3\,\mathbb{N}_6$, and the induced graphs satisfy $\boldsymbol{\mathcal{A}_1}$. Major ticks indicate instances where agent 1 transmits information, i.e. corresponding to $\mathcal{G}_1$ in Figure 4.3. The simulations were carried out with a time step of $\Delta t = 1 \times 10^{-3}$ on the time interval $t \in [0, 24]$. For each sampling interval, $h_k \coloneqq s_{k+1} - s_k$, a random integer m_k was drawn uniformly from the interval $[1, h_k/\Delta t]$, generating the delay $\tau_{ij}[k] = m_k \Delta t$, thus satisfying $\boldsymbol{\mathcal{A}_6}$. The major ticks at the top and corresponding dashed lines correspond to the delayed updates originating from agent 1 to agent 2. Both examples are simulated for the same delays, sampling sequence, and time interval.

The first simulation involves integrator agents as described in Subsection 4.2.1 with $F_\mathrm{d} = -5$. The agent's states can be seen in Figure 4.4(a), while the difference $\Delta_{\mu,i}(t) \coloneqq \bar{\mu}_i(t) - \bar{\mu}_i^{\mathrm{DF}}(t)$ is shown in Figure 4.4(b). It can be seen that indeed the agents asymptotically agree, and that $\Delta_{\mu,i}(t)$ repeatedly resets to zero after each agent finishes its "cycle" of delayed updates. Moreover, the trajectories are piecewise constant for this case since $\bar{\mu}_i(t)$ has no continuous-time dynamics as mentioned in the proof of Proposition 4.2.1.

The second example is comprised of identical agents with

$$\dot{x}_i(t) = \begin{bmatrix} 4 & 9 \\ 1 & 4 \end{bmatrix} x_i(t) + \begin{bmatrix} 2 \\ 1 \end{bmatrix} u_i(t)$$

trying to synchronize to $A_0 = \left[\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right]$. In this case

$$\bar{F} = -\begin{bmatrix} 2 & 4 \end{bmatrix} \quad \text{and} \quad F_\mathrm{d} = \begin{bmatrix} -34.6 & 39.2 \end{bmatrix}$$

satisfy the requirements of Theorem Theorem 4.2. The components of the agents' state are shown in Figure 4.5, and those of $\Delta_{\mu,i}$ are shown in Figure 4.6.

Once more, we can see that the agents' states synchronize to a common trajectory as in $\boldsymbol{\mathcal{P}_2}$ with A_0 corresponding to a sine wave with frequency 1. Furthermore, we again see that the difference between the delayed and delay-free system resets repeatedly after each "cycle" ends, and that the amplitude of the mismatch decays as the updates drive the systems closer to the agreement space. Note that this time $\bar{\mu}_i(t)$ is not piecewise constant between updates, since the synchronous trajectory is not constant.

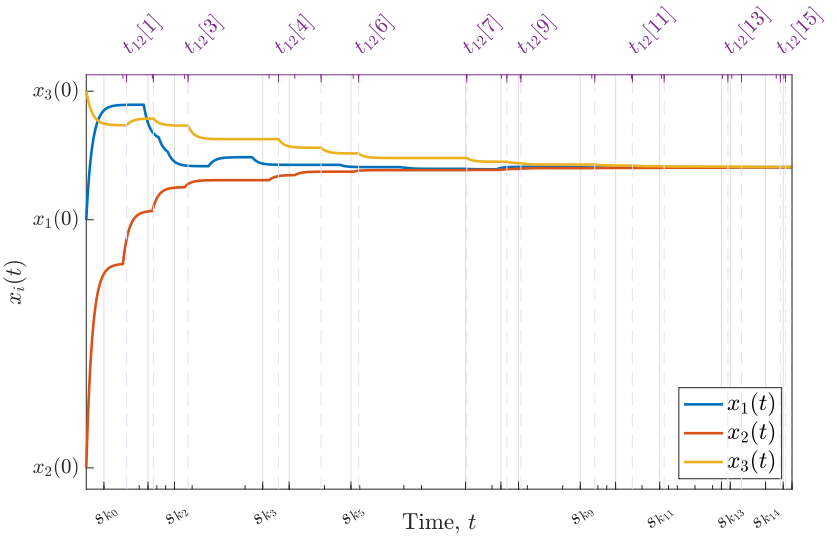

(a) Evolution of the agents' states for $A = A_0 = 0$.

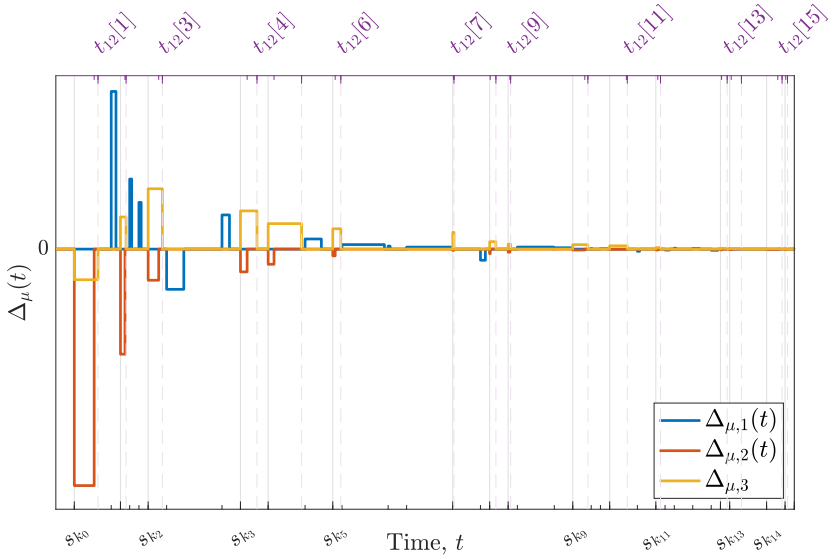

(b) The difference between $\bar{\mu}$ and its delay-free counterpart for $A = A_0 = 0$.

Figure 4.4: Simulations for the example with $A = A_0 = 0$ from Subsection 4.2.4. Minor ticks are sampling instances, major ticks indicate the connected subsequence from $\boldsymbol{\mathcal{A}_1}$. Major ticks in lavender correspond to the delayed counterparts of the regular major ticks.

4.3 Weighted update map

A common variation of the consensus protocol is the inclusion of *edge weights*. Locally, these weights modify protocol (1.6) to

$$u_i = \sum_{j \in \mathcal{N}_i} w_{ij}(y_j - y_i), \qquad i \in \mathbb{N}_\nu,$$

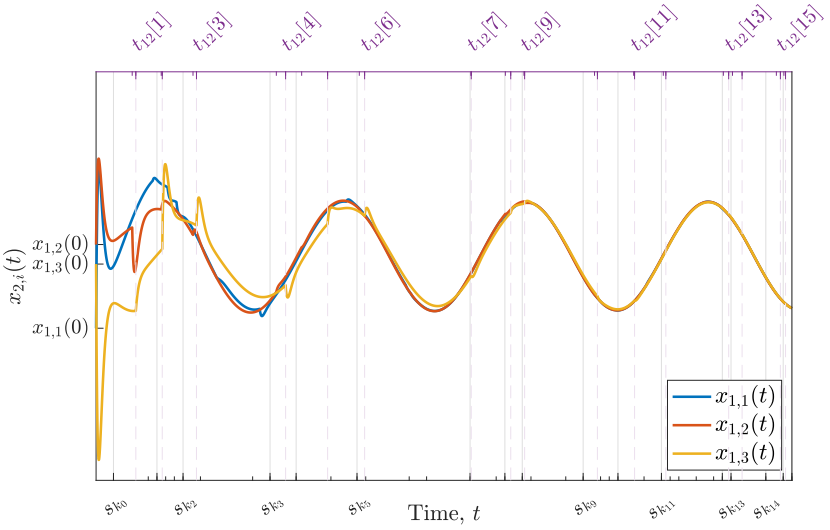

(a) Evolution of the agents' first state component for $A_0 \neq 0$.

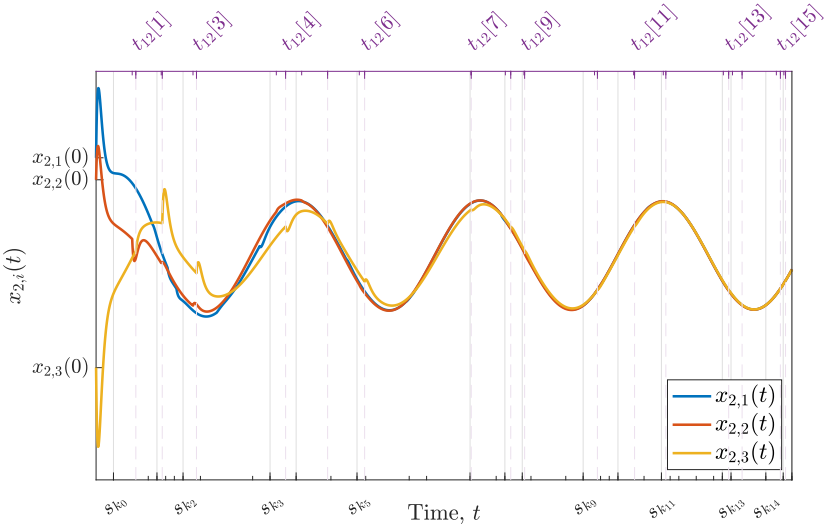

(b) [Evolution of the agents' second state component for $A_0 \neq 0$.

Figure 4.5: Evolution of the agents' states for the second example in Subsection 4.2.4, with $A_0 \neq 0$. Minor ticks are sampling instances, major ticks indicate the connected subsequence from $\boldsymbol{\mathcal{A}_1}$. Major ticks in lavender correspond to the delayed counterparts of the regular major ticks.

where w_{ij} are some non-negative weights. These weights can arise naturally from the modeling of the physical process [82], but can also be tuned to improve the convergence rate [83], [84] or attenuate the effects of measurement noise [85].

Despite being widely used, when constructing the consensus-based update maps in Chapter 2 and Chapter 3 we have opted to used the unweighted Laplacian. This choice was motivated by simplicity and in attempt to streamline the derivations, and is not an intrinsic requirement of the controllers. The following Proposition shows that under certain constraints on the weights, the original results hold verbatim even for a weighted update map.

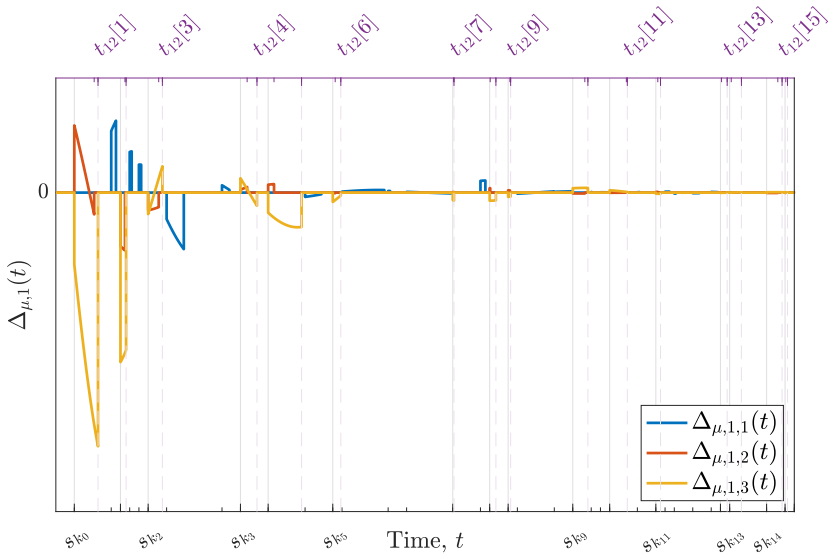

(a) Evolution of the first component of $\Delta_{\mu,i}$.

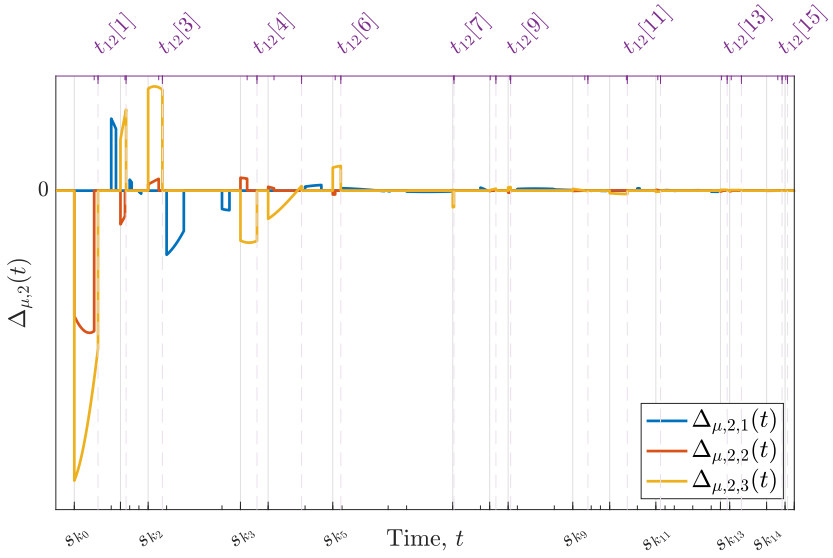

(b) Evolution of the second component of $\Delta_{\mu,i}$.

Figure 4.6: Evolution of the components of $\Delta_{\mu,i}$ for the second example in Subsection 4.2.4. Minor ticks are sampling instances, major ticks indicate the connected subsequence from $\mathcal{A}_1$. Major ticks in lavender correspond to the delayed counterparts of the regular major ticks.

Proposition 4.3.1. *Substituting update rule* (2.5) *with*

$$\mu_{ij}(s_k^+) = \mu_{ij}(s_k) - \alpha_{ij} \sum_{l \in \mathcal{N}_i[k]} (w_{il}[k](\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k))$$

for positive weights such that $\sum_l w_{il}[k] = 0$ *and* $w_{il}[k] < \nu$ *for all* k *does not affect the convergence properties of Theorem 3.1.*

Proof It is straightforward to verify that the weighted counterpart of (2.7) is

$$\tilde{A}_{\text{jmp}}[k] \coloneqq I_{\nu^2} - \frac{1}{\nu}\sum_{i=1}^{\nu}\sum_{j\in\mathcal{N}_i[k]}\left(e_i w_{ij}[k](e_i - e_j)'\right)\otimes(\alpha_i \mathbb{1}').$$

Reworking Lemma 2.3.1 for the new update map, (2.8a) obviously does not change since it requires only that $\mathbb{1}'(I - P_{\mathbb{1}}) = 0$, and the right hand side of (2.8b) is similarly unaffected. Now we must evaluate the summation on the left hand side of the Kronecker product, denote it by

$$L^w_{\mathcal{G}}[k] \coloneqq \sum_{i=1}^{\nu}\sum_{j\in\mathcal{N}_i[k]}\left(e_i w_{ij}[k](e_i - e_j)'\right).$$

Clearly both Lemma 2.3.2 and Lemma 3.3.1 still hold with this $L^w_{\mathcal{G}}[k]$ replacing the original Laplacian.

Under the assumptions on the weights the matrix $I - \frac{1}{\nu}L^w_{\mathcal{G}}[k]$ is still non-negative, row-stochastic, and all of its entries are uniformly bounded. Thus, by [21, Thm. 2.39] both (2.10) and (3.11b) will converge to the agreement space. In addition $L^w_{\mathcal{G}}[k]\mathbb{1} = 0$ for any graph and any weights satisfying the assumptions, thus the centroids-induced jumps asymptotically vanish and the proof for both Theorem 2.1 and Theorem 3.1 proceeds verbatim. ∎

At a glance Proposition 4.3.1 appears to be a small technical novelty. We considered a very general setting in which neither the graphs nor the sampling sequence are known, thus optimizing the weights seems unreasonable. However, the emulation controller can still be used under less stringent conditions. For example, perhaps the underlying graph is in fact known, but the unreliable sampling induces random subgraphs at each instance. Other possible scenarios are when only some communication links, i.e. edges, are unreliable; or when the sampling follows some pattern such as round-robin scheduling. In such situations the flexibility provided by Proposition 4.3.1 can be used to improve the performance. For example the approach in [86] can be used to provide improved performance with asymmetric link failures. Other relaxations include constant graphs and periodic sampling and dynamic induced graphs that are always undirected (i.e., synchronous sampling). For the former the weights can be chosen to minimize some quadratic cost [87], and for the latter there adaptive methods to counter adversarial attacks [88].

4.4 Concluding remarks

Up to this point, we have considered increasingly complex variations of the agreement problem. Moving from integrators to general LTI agents, from constant to time-varying trajectories, and from state to output feedback. We even touched upon heterogeneous

and time-varying transmission delays. Yet there are two obvious extensions we have, as of yet, not considered.

The first is transferring from homogeneous to heterogeneous agents. This involves some technical work, but does not impose a new conceptual problem. In fact, when considering heterogeneous agents it is well known that a necessary condition is the existence of a common *internal model* amongst all of the agents [44], [46]. The immediate solution is to identify $\bar{\mu}_i(t)$ with these models, and emulate some variation of the aforementioned continuous-time controllers. This avenue is not pursued in this thesis.

The second omission, is the agreement problem with external inputs - either measurement noise or load disturbances. Attenuating or even rejecting the effects of uncontrolled inputs are the bread and butter of classical control, hence this omission is particularly glaring. Surprisingly, such omissions are common within the multi-agent community, where it is common to assume the system is driven only by initial conditions. Such input-output analysis is the focus of the next chapter, which attempts to explain various odd behaviours of multi-agent systems with local disturbances.

Chapter 5

On Internal Stability of Diffusive-Coupling and the Dangers of Cancel Culture

> Something is rotten in the state of Denmark.
>
> W. Shakespeare, *Hamlet* [89]

In the previous chapters, we tackled different variations of sampled-data agreement problems. In all of them, we made a conscious assumption to explicitly treat each agent as if it has two distinct measurements: one local and one communicated. We made this choice since we wanted to explicitly impose different temporal constraints on the two input signals, but this viewpoint is uncommon in the literature. The common approach is to consider distributed control laws where only *relative* measurements are exchanged between neighbors. In other words, each agent has access only to the difference between its output and that of each of its neighbors. This restriction forces the agents to use *diffusive* control laws as in (1.8).

Yet even with our different approach, our entire work was centered around using the consensus protocol. The discrete updates were driven by a consensus-like protocol, and the emulated dynamics while different were inspired by it. This allowed us to utilize significant existing machinery to solve the resulting hybrid problems. This may be a double edges sword, inheriting both the useful technical machinery as well as possible underlying limitations of diffusive architectures. Hence, before continuing further on the sampled-data path, it would be beneficial to better understand diffusively coupled systems, their limitations, and how to circumvent them. This is the purpose of this chapter, which begins with a motivating example to illustrate concretely what we mean by "limitations". We then follow this example with an in depth system theoretic analysis of diffusive systems and fully characterize several issues that were previously reported but never explained in the literature.

5.1 Motivation

As discussed in Section 1.2, diffusively-coupled systems behave poorly when affected by disturbances and noise. In particular, persistent measurement noises or disturbances may result in trajectories exhibiting certain common traits associated with instability. Notably, the only systems to not exhibit this behavior are systems which are not purely diffusive, since they allow absolute measurements [90], [91] or an undisturbed leader [50]. This hints that there may be an *intrinsic* problem with the diffusive architecture. These "instability traits" can be illustrated by the classical consensus protocol, considered below for a set of integrator agents and with a static interaction network.

Example 5.1.1. Revisit the integrator consensus problem introduced in Section 1.1, which studies a group of independent integrator agents $\dot{x}_i(t) = u_i(t)$, where x_i and u_i are their states and control inputs, respectively. The goal is to reach asymptotic agreement between all agents, in the sense that

$$\lim_{t\to\infty}\big(x_i(t) - x_j(t)\big) = 0, \quad \forall i, j, \tag{5.1}$$

under the constraint that the ith agent has access only to states of its neighbors, whose indices belong to a set $\mathcal{N}_i$. This problem can be solved by the celebrated consensus protocol [24]

$$u_i(t) = -\kappa \sum_{j\in\mathcal{N}_i}\big(x_i(t) - x_j(t)\big), \quad \forall i, \tag{5.2}$$

which is diffusive state-feedback (1.3). From Proposition 2.1.2 we know that if the underlying graph is undirected and connected, then the control law (5.2) drives the agents to agreement exponentially fast. The state trajectories of four agents controlled by (5.2) with $\kappa = 1$ are shown in Figure 5.1 in the time interval $[0, t_d]$. Observe that on this time interval the states converge exponentially to the average of their initial conditions and the control signals all asymptotically vanish.

This might no longer be the case if the agents are affected by load disturbances d_i, viz.

$$\dot{x}_i(t) = u_i(t) + d_i(t). \tag{5.3}$$

An example of what happens in such situations is also shown in Figure 5.1. At the time instance $t = t_d$ one agent is affected by a unit step disturbance. As a result, all states cease to agree and start to diverge when $t > t_d$, whereas the control signals reach non-zero steady-state values. The apparent instability of the whole system, manifested in the unboundedness of the states, can be explained by the well-known fact that the consensus protocol has a closed-loop eigenvalue at the origin. To see this, note that a

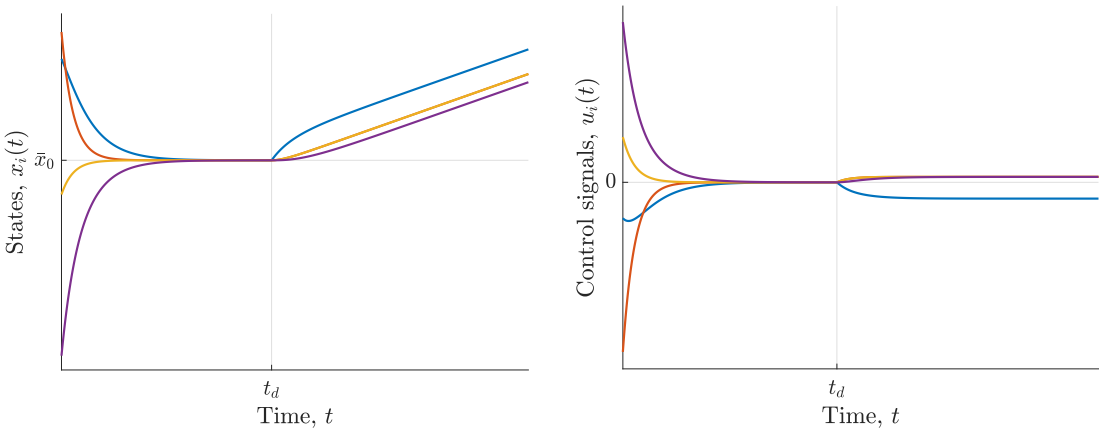

(a) Evolution of the states.

(b) Evolution of the control signals.

Figure 5.1: Simulation of protocol (5.2) with one agent perturbed by a step at $t = t_d$.

state-space realization of the aggregate system is given by

$$\begin{cases} \dot{x}(t) = -\kappa L_{\mathcal{G}} x(t) + d(t) \\ \begin{bmatrix} y(t) \\ u(t) \end{bmatrix} = \begin{bmatrix} I \\ -\kappa L_{\mathcal{G}} \end{bmatrix} x(t) \end{cases}, \tag{5.4}$$

and since $L_{\mathcal{G}}\mathbb{1} = 0$ for all $\mathcal{G}$, the system always has a pole at the origin. Nevertheless, the boundedness of the control signals under such conditions is intriguing. Situations wherein some signals in the closed-loop system are bounded while some others are not normally indicate unstable *pole-zero cancellations* in the feedback loop [92, Sec. 5.3]. However, controller (5.2) is static and thus has no zeros. ▽

The example above suggests that a deeper inspection of the *internal stability* property could offer insight into the behavior of diffusively-coupled systems. The internal stability of any feedback interconnection requires the stability of all possible input / output relations in the system, see [92], [93]. However, to the best of our knowledge, internal stability has not been explicitly studied in the context of diffusively-coupled architectures of MASs yet.

In this chapter, which is based on published work [63], we show that diffusively-coupled systems of LTI agents might not be internally stabilizable. Loosely speaking, this happens if the agents share common unstable dynamics, directions counting. This, for example, is always the case in a group of homogeneous unstable agents, like those discussed in Example 5.1.1. When restricting the result to finite-dimensional agents, we also explain the mechanism behind the shown internal instability. It is caused by *unstable cancellations* in the cascade of the aggregate plant and a diffusive controller. Important is that these cancellations are caused not by controller zeros, but rather by an intrinsic spatial deficiency of the diffusively-coupled configuration. These cancellations are intrinsic to the diffusive structure and cannot be affected by controller

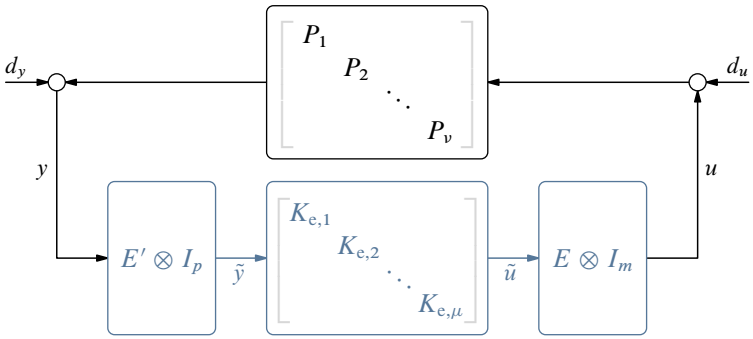

Figure 5.2: Block diagram of a general aggregated diffusively-coupled feedback setup (E is the incidence matrix of the connectivity graph $\mathcal{G}$).

dynamics. Consequently, the internal stability of feedback systems utilizing only relative measurements depends *solely* on the agent dynamics.

In addition to providing a rigorous analysis of the internal stability of diffusively-coupled systems, we show how the analysis is readily applied to common extensions found in the literature. In particular, we discuss more general symmetrically coupled MASs (i.e. not restricted to only diffusive coupling), asymmetric coupling (i.e. MASs over directed graphs), unstable systems with no closed right-half plane poles, and MASs over time-varying networks. This chapter relies heavily on several results from both linear systems and graph theory, the relevant preliminaries can be found in Appendix C and Appendix A respectively.

5.2 Problem formulation and general diffusive coupling

Consider ν continuous-time LTI agents P_i, each with m inputs and p outputs, who interact over a graph $\mathcal{G}$ with ν nodes and μ edges. In this formalism, agents i and j are neighbors if they are incident to the same edge.

A general diffusively-coupled MAS originated in [32], also known as the canonical cooperative control structure [72, Ch. 9], is presented in Figure 5.2. It comprises the block-diagonal aggregate plant $P := \operatorname{diag}\{P_i\}$ with ν blocks, a block-diagonal edge controller $K_e := \operatorname{diag}\{K_{e,j}\}$ with μ blocks, and pre- and post-processing based on the incidence matrix E associated with $\mathcal{G}$. To describe the logic of this setup we may disregard the exogenous signals d_y and d_u for the time being. The overall controller $K : y \mapsto u$ is thus defined as

$$K := (E \otimes I_m) K_e (E' \otimes I_p). \tag{5.5}$$

We now discuss how the controller K processes signals.

- The (νp)-dimensional aggregate output of the agents, y, is first processed by the transpose of the incidence matrix to produce a (μp)-dimensional vector $\tilde{y} = (E' \otimes I_p) y$ representing the relative outputs of neighboring agents.

- Each component of $\tilde{y}$, which is the relative measured coordinate along one edge, is then processed independently by an edge controller $K_{e,j}$, to produce a (μm)-dimensional "edge correction" signal $\tilde{u}$.
- The (νm)-dimensional aggregate control signal u is then produced by processing all $\tilde{u}_j$ by the incidence matrix, which sums up edge corrections for all edges connected to the corresponded node.

For example, if $\mathcal{G}$ is an undirected star graph on three nodes with node 3 as its center, then we can choose

$$E = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ -1 & -1 \end{bmatrix},$$

in which case

$$\tilde{y} = \begin{bmatrix} y_1 - y_3 \\ y_2 - y_3 \end{bmatrix} \quad \text{and} \quad u = \begin{bmatrix} \tilde{u}_1 \\ \tilde{u}_2 \\ -\tilde{u}_1 - \tilde{u}_2 \end{bmatrix}.$$

The consensus protocol (5.2) corresponds to the choice $K_e = -\kappa I$ in this case, as well as for any other choice of $\mathcal{G}$ and ν.

Now consider the exogenous signals d_u and d_y, which we refer to as disturbances. On the physical level they represent inevitable effects of the outside world on the controlled plant (agents). These signals are supposed to be bounded and independent of the signals generated by the controlled system. We introduce disturbances to define the notion of the internal stability for the system in Figure 5.2, which is the focus point of this chapter.

Definition 5.2.1. We say that the system in Figure 5.2 is *internally stable* if the 2×2 operator connecting exogenous signals d_u and d_y with internal signals u and y, i.e.

$$T_4 : (d_y, d_u) \mapsto (y, u) \tag{5.6}$$

is well defined and stable, see [94, Sec. 4].

The general question of interest in this chapter is *under what conditions on the agents P_i are there causal edge controllers $K_{e,j}$ internally stabilizing the diffusively-coupled system in Figure 5.2?* Note that the existence of edge controllers rendering the closed-loop operator well defined is obvious, just take $K_{e,j} = 0$ for all j. We shall thus focus on the stability of T_4.

Addressing the stability question in the most general, nonlinear and time-varying, case might be overly technical. We thus limit our attention to the class of LTI plants and edge controllers, whose transfer functions belong to the quotient field of H_∞, see [95, §A.7.1], which is a sufficiently general class. We further assume that

$\mathcal{A}_7$: there are right coprime $M_i, N_i \in H_\infty$ and left coprime $\tilde{M}_i, \tilde{N}_i \in H_\infty$ such that $P_i = N_i M_i^{-1} = \tilde{M}_i^{-1}\tilde{N}_i$ for all i,

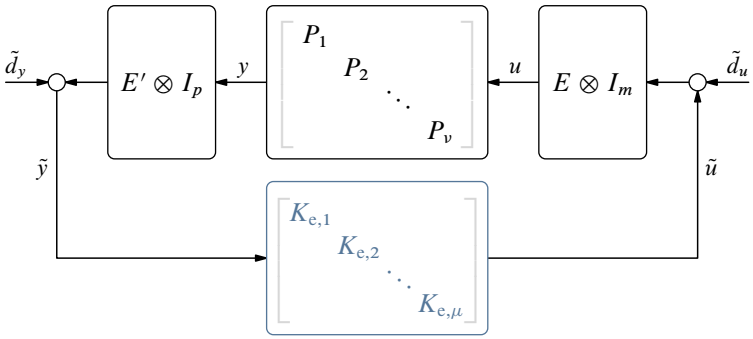

Figure 5.3: Block diagram of a diffusively-coupled feedback setup as edge stabilization. Here the incidence matrix is attached to the plant, generating the edge dynamics with a diagonal controller.

where coprimeness is understood as the existence of Bézout coefficients in H_∞, see Appendix C. The representation of P_i above is known as its coprime factorization. We hereafter refer to the transfer functions $M_i(s)$ and $\tilde{M}_i(s)$ as the right and left denominators of P_i, respectively, and the transfer functions $N_i(s)$ and $\tilde{N}_i(s)$ as its right and left numerator. Assumption $\mathcal{A}_7$ is practically nonrestrictive. It holds for all finite-dimensional agents with proper transfer functions and is equivalent to the stabilizability of P_i by feedback for agents with transfer functions from the quotient field of H_∞ [96]. Thus, if an agent fails to satisfy $\mathcal{A}_7$, we cannot expect any MAS that includes it to be stabilizable by diffusive coupling.

Remark 5.1. We choose the application points of exogenous disturbances for the internal stability analysis to be at the points where the agents, P, are connected with the controller K defined in (5.5). In this choice we follow the physical nature of the interconnection in Figure 5.2 and think of separating the blocks $E \otimes I$ and $E' \otimes I$ in the controller as merely a way to streamline the choice of the design parameters, which are the edge controllers in K_e. An alternative viewpoint is presented in Figure 5.3, where all *fixed* parts are regarded as the controlled plant,

$$P_e \coloneqq (E' \otimes I_p) P (E \otimes I_m), \tag{5.7}$$

much inline with the generalized plant philosophy [93, Sec. 3.8], see e.g. [33, Fig. 6] or [72, E9.6]. A natural definition of internal stability for it shall be based on the exogenous inputs $\tilde{d}_y$ and $\tilde{d}_u$, entering before and after the edge controller K_e. This would change the results, see Remark 5.3 at the end of Subsection 5.5.1. Still, we believe that the configuration in Figure 5.2 is the right way to address the internal stability of MASs. After all, it is the agents who interact with the environment. ▽

5.3 The internal stability of diffusive coupling

The main technical result of this chapter, whose proof is postponed to Subsection 5.3.1, is formulated as follows.

Theorem 5.1. *No LTI $K_{e,j}$ can internally stabilize the diffusively-coupled system in Figure 5.2 if there is $\lambda \in \bar{\mathbb{C}}_0$, common to all agents, such that*

$$\bigcap_{i=1}^{\nu} \ker [M_i(\lambda)]' \neq \{0\} \tag{5.8a}$$

or

$$\bigcap_{i=1}^{\nu} \ker \tilde{M}_i(\lambda) \neq \{0\}, \tag{5.8b}$$

where M_i and $\tilde{M}_i$ are denominators in the coprime factorizations of P_i under $\boldsymbol{\mathcal{A}_7}$.

Theorem 5.1, formulated in terms of coprime factors of agents, might appear somewhat abstract and technical. This is a consequence of considering a fairly general class of LTI agents under the mild assumption $\boldsymbol{\mathcal{A}_7}$. We show in the next section that if the class of admissible agents is limited to finite-dimensional ones, then more insightful statements can be provided. Nevertheless, the formulation in Theorem 5.1 becomes substantially more intuitive in some frequently studied special cases.

The first of them is the case of homogeneous agents, which is perhaps the best studied situation.

Corollary 5.2. *If the agents are homogeneous, i.e. $P_i = P_0$ for all $i \in \mathbb{N}_\nu$, and $P_0(s)$ has at least one pole in $\bar{\mathbb{C}}_0$, then no LTI $K_{e,j}$ can internally stabilize the system in Figure 5.2.*

Proof By Lemma C.1.6, if $\lambda \in \bar{\mathbb{C}}_0$ is a pole of $P_0(s)$, then both $M_0(\lambda)$ and $\tilde{M}_0(\lambda)$ are singular, whence the result follows. ∎

This result readily applies to the homogeneous consensus problem studied in Example 5.1.1 and more generally to the setups in Subsection 1.1.2. Note that closed-loop dynamics (5.4) can be rewritten as

$$\begin{cases} \dot{x}(t) = E \operatorname{diag}\{-\kappa\} E' x(t) + d(t) \\ \begin{bmatrix} y(t) \\ u(t) \end{bmatrix} = \begin{bmatrix} I \\ E \operatorname{diag}\{-\kappa\} E' \end{bmatrix} x(t) \end{cases},$$

which is in the form of Figure 5.2. The agents in (5.3) are homogeneous and $P_0(s) = 1/s$, has an unstable pole at the origin. Corollary 5.2 then agrees with the conclusion of Example 5.1.1 that the closed-loop system is not internally stable.

Another particular case for which the formulation is simplified is a MAS with single-input single-output (SISO) agents.

Corollary 5.3. *If the agents are SISO and all have a pole at the same $\lambda \in \bar{\mathbb{C}}_0$, regardless of multiplicities, then no LTI $K_{e,j}$ can internally stabilize the diffusively-coupled system in Figure 5.2.*

Proof By Lemma C.1.6, in this case $M_i(\lambda) = \tilde{M}_i(\lambda) = 0$ for all $i \in \mathbb{N}_\nu$, whence the result follows. ∎

A consequence of Corollary 5.3 is that the consensus protocol, as well as any other diffusively-coupled control laws, cannot internally stabilize a group of SISO agents if all of them contain an integral action. This result is reminiscent of that by [44] that states that a common internal model is a necessary condition for a diffusively-coupled system to synchronize their state trajectories. It highlights a contradiction or trade-off of sorts, where on the one hand, a common pole at the origin among agents is required for synchronization, and on the other hand, this common (unstable) pole is precisely the cause for lack of internal stability.

5.3.1 Proof of Theorem 5.1

We are now prepared to prove Theorem 5.1. Only the statement about the right coprime factor, i.e. (5.8a), is proved. The proof of (5.8b) follows by dual arguments.

The proof requires a technical result of [97], known as the matrix corona theorem, see also the proof of [94, Prop. 11] for a closer formulation.

Lemma 5.3.1. *If $G \in H_\infty^{n\times n}$, then*

$$G^{-1} \in H_\infty \iff \inf_{s\in\bar{\mathbb{C}}_0} \underline{\sigma}(G(s)) > 0.$$

It is readily seen that $M_P \coloneqq \operatorname{diag}\{M_i\}$ and $N_P \coloneqq \operatorname{diag}\{N_i\}$ are right coprime factors of $P = \operatorname{diag}\{P_i\}$. Because any internally stabilizing K in (5.5) is in effect stabilized by the plant, we only need to consider edge controllers for which K admits coprime factorizations over H_∞. So let $K = N_K M_K^{-1}$ for right coprime $M_K, N_K \in H_\infty$. By (5.5),

$$N_K(s) = (E \otimes I_m)K_{\mathrm{e}}(s)(E' \otimes I_p)M_K(s).$$

Because $\mathbb{1}'E = 0$, we have that $(\mathbb{1}'\otimes I_m)(E\otimes I_m) = 0$ as well and, hence, $(\mathbb{1}'\otimes I_m)N_K(s) = 0$ for all s at which $K_{\mathrm{e}}(s)$ is finite. But $K_{\mathrm{e}}(s)$ is in the quotient field of H_∞, meaning that the denominators of its entries are holomorphic in $\mathbb{C}_0$ and, by [98, Thm. 10.18], may have at most countable number of isolated zeros. As such, we can always find a region in $\mathbb{C}_0$ in which $(\mathbb{1}' \otimes I_m)N_K(s) = 0$. But the latter implies that

$$(\mathbb{1}' \otimes I_m)N_K = 0,$$

by the same [98, Thm. 10.18].

Now, return to the system in Figure 5.2. It is readily verified that the closed-loop system T_4 in (5.6) reads

$$T_4 = \begin{bmatrix} I \\ K \end{bmatrix} (I - PK)^{-1} \begin{bmatrix} I & P \end{bmatrix} =: \begin{bmatrix} S & T_{\mathrm{d}} \\ T_{\mathrm{c}} & T \end{bmatrix}, \tag{5.9}$$

where the blocks of T_4 are the four fundamental closed-loop transfer functions. Straightforward algebra yields that

$$T_4 = \begin{bmatrix} M_K & 0 \\ N_K & 0 \end{bmatrix} \begin{bmatrix} M_K & -N_P \\ -N_K & M_P \end{bmatrix}^{-1}. \tag{5.10}$$

This is a right coprime factorization of T_4, as attested by the Bézout equality (cf. (C.4a))

$$\begin{bmatrix} \tilde{M}_P & \tilde{N}_P \\ -Y_P & X_P \end{bmatrix} \begin{bmatrix} M_K & -N_P \\ -N_K & M_P \end{bmatrix} + \begin{bmatrix} X_K - \tilde{M}_P & Y_K + \tilde{N}_P \\ Y_P & X_P \end{bmatrix} \begin{bmatrix} M_K & 0 \\ N_K & 0 \end{bmatrix} = I,$$

where $\tilde{M}_P \coloneqq \operatorname{diag}\{\tilde{M}_i\}$ and $\tilde{N}_P \coloneqq \operatorname{diag}\{\tilde{N}_i\}$. By Lemma C.1.5, T_4 is stable if and only if

$$\begin{bmatrix} M_K & -N_P \\ -N_K & M_P \end{bmatrix}^{-1} \in H_\infty, \tag{5.11}$$

or

$$\inf_{s \in \bar{\mathbb{C}}_0} \underline{\sigma}\left(\begin{bmatrix} M_K(s) & -N_P(s) \\ -N_K(s) & M_P(s) \end{bmatrix} \right) > 0 \tag{5.12}$$

by Lemma 5.3.1. But (5.8a) implies that there is $v \neq 0$ such that $v' M_i(\lambda) = 0$ for all i or, equivalently, $(\mathbb{1} \otimes v)' M_P(\lambda) = 0$. Taking into account that $(\mathbb{1} \otimes v)' N_K = v'(\mathbb{1} \otimes I_m)' N_K = 0$, we end up with

$$\begin{bmatrix} 0 & (\mathbb{1} \otimes v)' \end{bmatrix} \begin{bmatrix} M_K(\lambda) & -N_P(\lambda) \\ -N_K(\lambda) & M_P(\lambda) \end{bmatrix} = 0, \tag{5.13}$$

which violates (5.12). We thus have that if (5.8a) holds, then there is no K_{e} that internally stabilizes the system in Figure 5.2.

5.4 Generalizations

Some possible generalizations of the result of Theorem 5.1 are outlined below.

5.4.1 Asymmetric coupling

Some MAS problems consider a *directed* interaction graph, making the notion of neighboring agents asymmetric. Controllers under such constrains are no longer diffusive in the sense discussed in Section 5.2. Still, a variant of Theorem 5.1 may apply.

For example, let an edge going from node i to node j indicate that the ith agent has access to $y_i - y_j$. The existence of the edge (i, j) does not imply that there is also

the edge (j,i). It is evident that the controller outlined in Figure 5.2 and (5.5) can no longer provide an appropriate distributed controller since, as discussed in Section 5.2, it sums up all the edge correction terms connected to each corresponding node. Nevertheless, several notable MAS control architectures over directed graphs still admit a decomposition similar that of (5.5).

Consider again the classic consensus protocol. It can be adapted to accommodate directed graphs by replacing the symmetric Laplacian, $L_{\mathcal{G}} = EE'$, with a directed counterpart such as the in and out degree Laplacians. To this end, define the following auxiliary matrices

$$[B_{\text{out}}]_{ie} = \begin{cases} 1 & \text{if vertex } i \text{ is the head of edge } e \\ 0 & \text{otherwise} \end{cases} \tag{5.14a}$$

and

$$[B_{\text{in}}]_{ie} = \begin{cases} 1 & \text{if vertex } i \text{ is the tail of edge } e \\ 0 & \text{otherwise} \end{cases}. \tag{5.14b}$$

Note that by construction $E = B_{\text{out}} - B_{\text{in}}$, hence both are binary matrices with both column and row sums equal to 1. The following proposition allows us to represents directed Laplacians using these two matrices and the regular incidence matrix.

Proposition 5.4.1. *Given a directed graph $\mathcal{G}$ and the in and out incidence matrices from* (5.14)*, the following relationships hold:*

1. *The adjacency matrix is given by*

$$A_{\mathcal{G}} = B_{in}B'_{out}.$$

2. *The degree matrices satisfy*

$$D^i_{\mathcal{G}} = B_{in}B'_{in} \quad \textit{and} \quad D^o_{\mathcal{G}} = B_{out}B'_{out}.$$

3. *The directed Laplacians satisfy*

$$L^i_{\mathcal{G}} = B_{in}E' \quad \textit{and} \quad L^o_{\mathcal{G}} = EB_{out}.$$

Proof The proof can be found in [72, E9.13], the only difference is that their adjacency is defined as the transpose of how we defined $A_{\mathcal{G}}$ in (A.1). Consequently, our Laplacians are defined slightly differently. ∎

Using the above, we can represent the directed Laplacian by the product $L^i_{\mathcal{G}} = B_{\text{in}}E'$. This suggests that a controller of the form

$$K_{\text{in}} \coloneqq (B_{\text{in}} \otimes I_m)K_{\text{e}}(E' \otimes I_p), \tag{5.15}$$

can be used to represent various control laws over directed graphs. For example setting $K_e = -I$ results in the aforementioned directed consensus protocol, while picking $K_e = I_\nu \otimes \tilde{K}$ for some gain $\tilde{K}$ yields the synchronizing controllers discussed in [72, Sec. 8.4].

The controller structure in (5.15) mirrors that in (5.5). If (5.8b) holds, then the proof of Theorem 5.1 applies verbatim to any MAS controlled by it. However, this is not the case for (5.8a), implying that some systems may be stabilizable only if the graph is directed, as illustrated in the following example.

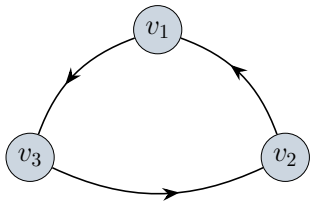

Figure 5.4: The communication graph used in Example 5.4.2.

Example 5.4.2. Consider a system of $\nu = 3$ first-order agents

$$P_1(s) = \begin{bmatrix} 1 & 0 \\ 1/s & 1 \end{bmatrix} \quad \text{and} \quad P_2(s) = P_3(s) = \begin{bmatrix} 1/s & 0 \\ 1 & 1 \end{bmatrix}.$$

Assume that their connectivity is represented by the directed cycle graph, which has three directed edges $(1,3)$, $(3,2)$, and $(2,1)$ as shown in Figure 5.4. This system can be described by (5.15) with

$$E = \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & -1 \\ -1 & 0 & 1 \end{bmatrix}, \quad B_{\text{in}} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix},$$

and arbitrary block-diagonal edge controllers. It is then a matter of standard algebra to verify that these plants admit denominators

$$\tilde{M}_1(s) = \begin{bmatrix} 1 & 0 \\ 0 & s/(s+1) \end{bmatrix}$$

and

$$\tilde{M}_2(s) = \tilde{M}_3(s) = \begin{bmatrix} s/(s+1) & 0 \\ 0 & 1 \end{bmatrix} = M_i(s), \quad \forall i \in \mathbb{N}_3.$$

Hence, condition (5.8a) holds for $\lambda = 0$, whereas condition (5.8b) holds for no λ. Thus, if the interconnection graph was undirected, then Theorem 5.1 would rule out the existence of internally stabilizing edge controllers. But in the directed case in form (5.15) with B_{in} having full rank, what matters is only (5.8b). Hence, we cannot rule out the existence of an internally stabilizing controller. And indeed, it can be verified

that

$$K_{\mathrm{e}}(s) = \operatorname{diag}\left\{I_2, \begin{bmatrix} 1 & 1 \\ 5/3 & 1 \end{bmatrix}, I_2\right\}$$

results in an internally stable interconnection, with the closed-loop poles in $\{-1/2, -3/4, -1\}$. ▽

Of course, following a similar procedure we may define the analogous K_{out} (corresponding for example to the out-degree directed consensus protocol) and consider only condition (5.8a), then again the proof holds unchanged.

Remark 5.2. The stabilizability of control architectures over directed graphs may nevertheless still require checking both conditions of Theorem 5.1. This thesis is based on an interpretation of the edge controller (5.15) as (dynamic) edge weights of the directed graph. A directed graph is called *weight balanced* if the accumulated weights of incoming and outgoing edges are equal for each node. It is known [22, Thm. 3.17] that the consensus protocol for integrator agents can reach an average agreement, i.e. $x_i(t) \to (1/\nu)\mathbb{1}'x(0)$ for all i, iff the underlying digraph is weight balanced and weakly connected. A key property to prove this result is that the Laplacian of a weight-balanced digraph, $L^o_{\mathcal{G}}$, satisfies $\ker L^o_{\mathcal{G}} = \ker(L^o_{\mathcal{G}})' = \operatorname{Im}\mathbb{1}$. Viewed within the context of Theorem 5.1, this implies that if edge controllers in (5.15) are chosen such that digraph is weight balanced, then both conditions of (5.8) must be checked anyway. ▽

5.4.2 Arbitrary symmetric coupling

The result of Theorem 5.1 still holds if the incidence matrix is replaced with a different coupling matrix, say $F \in \mathbb{R}^{\mu\times\nu}$, as long as there is a vector $0 \neq v \in \mathbb{R}^{\mu}$ such that $v'F = 0$. Such generalizations of a MAS were recently discussed in [99], but are also included in works considering, for example, distributed function calculation in MAS [100]. This corollary is formulated below.

Corollary 5.4. *Consider the setup in Figure 5.2, with some arbitrary coupling matrix $F \in \mathbb{R}^{\mu\times\nu}$ with some vector $0 \neq v \in \mathbb{R}^{\mu}$ such that $v'F = 0$. No LTI $K_{e,j}$ can internally stabilize the systems if there is $\lambda \in \bar{\mathbb{C}}_0$, common to all agents, such that*

$$\bigcap_{i=1}^{\nu} \ker\left[v_i M_i(\lambda)\right]' \neq \{0\} \tag{5.16a}$$

or

$$\bigcap_{i=1}^{\nu} \ker\left[v_i \tilde{M}_i(\lambda)\right] \neq \{0\}. \tag{5.16b}$$

where M_i and $\tilde{M}_i$ are denominators in the coprime factorizations of P_i under $\mathcal{A}_7$.

Proof The proof follows that of Theorem 5.1, the only change is that (5.13) now reads

$$\begin{bmatrix} 0 & (v \otimes \eta)' \end{bmatrix} \begin{bmatrix} M_K(\lambda) & -N_P(\lambda) \\ -N_K(\lambda) & M_P(\lambda) \end{bmatrix} = \begin{bmatrix} 0 & (v \otimes \eta)' M_P(\lambda) \end{bmatrix}.$$

The element above reads

$$(v \otimes \eta)' M_P(\lambda) = \eta'(v' \otimes I_m) M_p(\lambda) = \eta' \begin{bmatrix} v_1 M_1(\lambda) & \cdots & v_\nu M_\nu(\lambda) \end{bmatrix}$$

and it is equal zero if and only if $\eta \in \bigcap_{i=1}^{\nu} \ker [v_i M_i(\lambda)]'$. The rest follows using identical arguments to the diffusive coupling case. ∎

5.4.3 Unstable systems with no poles in $\bar{\mathbb{C}}_0$

It might happen that $P_i \notin H_\infty$ not because of poles, or other singularities, in $\bar{\mathbb{C}}_0$. For example, $P_i(s) = s/(s + 1 + se^{-s})$ has no singularities in $\bar{\mathbb{C}}_0$, but nonetheless does not belong to H_∞, see [101]. The proof still applies in this case, and all we need is to replace (5.8) with the assumption that there is a sequence $\{\lambda_j\}$ in $\mathbb{C}_0$ such that $\inf_{\{\lambda_j\}} v' M_i(\lambda_j) = 0$, or its dual version, holds for all $i \in \mathbb{N}_\nu$ and some $v \neq 0$.

5.4.4 Time-varying K

The main result also extends to the case of time-varying controllers. This is particularly relevant for varying interconnection topologies, i.e. those where $E_{\mathcal{G}(t)} = E(t)$ is the incidence matrix of the time-varying graph $\mathcal{G}(t)$. Still, the condition $\mathbb{1}' E(t)$ holds for any topology, rendering the denominator in (5.10) not stably invertible. We can then use [102, Theorem (i)] to show that under no choice of K_e the system is stabilizable, at least in the finite-dimensional case, whenever either one of the conditions in (5.8) holds.

5.5 Finite-dimensional agents

If the agents P_i are finite dimensional, the result of the previous section can be reformulated in a more insightful way. This is due to the ultimate connection between stability and pole locations, as well as clear definitions of cancellations in this case. So we proceed with assuming that all transfer functions $P_i(s)$ are real rational and proper ($\mathcal{A}_7$ always holds then).

Let $\mathrm{pdir_i}(G, \lambda)$ and $\mathrm{pdir_o}(G, \lambda)$ denote input and output direction of a pole λ in $G(s)$, see Appendix C for details and other related definitions. The result below reformulates the conditions of Theorem 5.1 via pole directions of agents.

Proposition 5.5.1. *If $P_i(s)$ are real rational and proper, then* (5.8a) *and* (5.8b) *are*

equivalent to the existence of $\lambda \in \bar{\mathbb{C}}_0$ *such that*

$$\bigcap_{i=1}^{\nu} \operatorname{pdir}_{\mathrm{i}}(P_i, \lambda) \neq \{0\} \tag{5.17a}$$

and

$$\bigcap_{i=1}^{\nu} \operatorname{pdir}_{\mathrm{o}}(P_i, \lambda) \neq \{0\}, \tag{5.17b}$$

respectively.

Proof Because $\lambda \in \bar{\mathbb{C}}_0$ is not a pole of $M_i(s)$, Lemma C.1.10 applies and (5.8a) reads $\cap_{i=1}^{\nu} \operatorname{zdir}_{\mathrm{o}}(M_i, \lambda) \neq \{0\}$. Then (5.17a) follows by Lemma C.1.11. The proof for (5.17b) is similar. ■

In other words, for the system in Figure 5.2 to not be stabilizable, the agents should not only have a common unstable pole, but also a common nontrivial direction of such a pole. Directions are obviously matched in the homogeneous and SISO cases addressed in Corollary 5.2 and Corollary 5.3, respectively. But the MIMO heterogeneous case may be less trivial.

Example 5.5.2. Consider a system with $\nu = 2$ first-order agents

$$P_1(s) = \begin{bmatrix} 1/s & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad P_2(s) = \begin{bmatrix} 1 \\ \alpha \end{bmatrix} \frac{1}{s} \begin{bmatrix} 1 & \beta \end{bmatrix}.$$

Directions of their pole at the origin are

$$\operatorname{pdir}_{\mathrm{i}}(P_1, 0) = \operatorname{pdir}_{\mathrm{o}}(P_1, 0) = \operatorname{Im} \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad \operatorname{pdir}_{\mathrm{i}}(P_2, 0) = \operatorname{Im} \begin{bmatrix} 1 \\ \beta \end{bmatrix},$$

and

$$\operatorname{pdir}_{\mathrm{o}}(P_2, 0) = \operatorname{Im} \begin{bmatrix} 1 \\ \alpha \end{bmatrix}.$$

There are nontrivial intersections between input and output directions of the agents if and only if $\beta = 0$ and $\alpha = 0$, respectively. The incidence matrix is $E = \left[\begin{smallmatrix} 1 \\ -1 \end{smallmatrix}\right]$ in this case. Choose the edge controller (there is only one edge in this example) as

$$K_{\mathrm{e}}(s) = \begin{bmatrix} (\alpha - \beta)\beta & -\alpha \\ \beta & 0 \end{bmatrix}.$$

The closed-loop characteristic polynomial, understood as the lowest common denominator of elements of $T_4(s)$ in (5.9), is then $(s+\alpha^2)(s+\beta^2)$. Thus, the closed-loop system is stable unless $\alpha = 0$ or $\beta = 0$, which agrees with (5.17). ▽

Also worth emphasizing is that conditions (5.17a) and (5.17b) might not be equivalent for MIMO agents, as illustrated by the example below.

Example 5.5.3. Return to the system studied in Example 5.4.2. Directions associated with the (unstable) pole at the origin are

$$\mathrm{pdir_i}(P_i, 0) = \mathrm{Im}\begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad \forall i \in \mathbb{N}_3$$

but

$$\mathrm{pdir_o}(P_1, 0) = \mathrm{Im}\begin{bmatrix} 0 \\ 1 \end{bmatrix} \neq \mathrm{Im}\begin{bmatrix} 1 \\ 0 \end{bmatrix} = \mathrm{pdir_o}(P_2, 0)\,.$$

Thus, in this case (5.17a) holds, whereas (5.17b) does not. This agrees with what we saw in Example 5.4.2 with respect to conditions (5.8). ▽

Another outcome of the finite dimensionality is that the formulation of Corollary 5.2 can be strengthened to an "if and only if" statement.

Corollary 5.5. *If the agents are homogeneous, i.e. $P_i = P_0$ for all $i \in \mathbb{N}_\nu$, and $P_0(s)$ is real rational and proper, then an LTI $K_{e,j}$ can internally stabilize the diffusively-coupled system in Figure 5.2 if and only if P_0 is stable.*

Proof If P_0 is unstable, then it has a pole in $\bar{\mathbb{C}}_0$ and Corollary 5.2 applies. If P_0 is stable, $K_e = 0$ does the job. ■

One should be careful not to conclude from the proof of Corollary 5.5 that only $K_e = 0$ can be used to guarantee internal stability. The case of $K_e = 0$ effectively decouples all the agents leading only to a "trivial" coordination (i.e. all agents converge to the origin). One can design edge controllers with additional external inputs to drive the relative states $\tilde{y}$ to non-trivial solutions using the methods, for example, described in [103]. For non-trivial agreement among the agents, the use of an unstable edge controller is possible provided that an appropriately defined external input is fed into the system at the point d_y in Figure 5.2.

5.5.1 Diffusive control laws and unstable cancellations

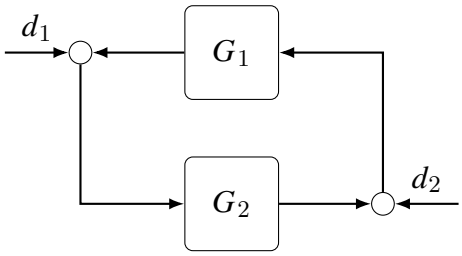

Figure 5.5: A generic feedback interconnection for internal stability.

The formulation of Proposition 5.5.1 is more intuitive than that of Theorem 5.1. Still, neither of them explains *why* no edge controller can stabilize the system in Figure 5.2 if agents share common unstable dynamics, directions counted. In this part

we aim at offering explanations. We argue that a key property to this end is intrinsic *unstable cancellations* between the plant and the controller.

Definition 5.5.4. We say that the cascade (series) interconnection G_2G_1 has cancellations if

$$\deg(G_2G_1) < \deg(G_1) + \deg(G_2)$$

where $\deg(G)$ is the *McMillan Degree* (cf. [104, Thm. 3.5]). We say that a pole of $G_1(s)$ and/or $G_2(s)$ is *canceled* if its multiplicity in $G_2(s)G_1(s)$ is smaller than the sum of its multiplicities in $G_1(s)$ and $G_2(s)$.

In other words, cancellations mean that some parts of the dynamics (modes) of either factor disappear in the cascade. Cancellations in the SISO case are always caused by the presence of zeros of $G_1(s)$ at the locations of poles of $G_2(s)$, or vice versa. As such, they are termed *pole-zero cancellations*. The situation is more complex in the MIMO case. For example, let

$$G_1(s) = \frac{1}{s}\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad G_2(s) = \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix},$$

with $\deg(G_1) = 2$ (two poles at the origin) and $\deg(G_2) = 0$ (no poles). The system G_2 is static and thus has no zeros either. Nevertheless, the transfer function

$$G_2(s)G_1(s) = \frac{1}{s}\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}$$

is first order, meaning that one of the poles of $G_1(s)$ is canceled. Such cancellations, brought on by the normal rank deficiency of $G_2(s)$, are a lesser-known phenomenon.

This, it seems, is exactly what happens in classic consensus, as shown in the following example.

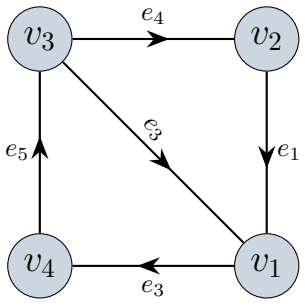

Figure 5.6: The communication graph used in Example 5.5.5

Example 5.5.5. Consider once more Example 5.1.1 of integrator agents attempting to achieve consensus, this time with $\nu = 4$ agents, and the corresponding undirected graph

as shown in Figure 5.6. We have $P(s) = \frac{1}{s} I_4$, and with $K_e = -I_5$ we obtain the following consensus protocol

$$E_{\mathcal{G}} = \begin{bmatrix} -1 & -1 & 1 & 0 & 0 \\ 1 & 0 & 0 & -1 & 0 \\ 0 & 1 & 0 & 1 & -1 \\ 0 & 0 & -1 & 0 & 1 \end{bmatrix} \quad \Longrightarrow \quad K = \begin{bmatrix} -3 & 1 & 1 & 1 \\ 1 & -2 & 1 & 0 \\ 1 & 1 & -3 & 1 \\ 1 & 0 & 1 & -2 \end{bmatrix}.$$

Clearly $\deg(P) = 4$ and $\deg(K) = 0$, yet it is readily verified that

$$P(s) E_{\mathcal{G}} K_e E_{\mathcal{G}}' = \frac{1}{s} \begin{bmatrix} -3 & 1 & 1 & 1 \\ 1 & -2 & 1 & 0 \\ 1 & 1 & -3 & 1 \\ 1 & 0 & 1 & -2 \end{bmatrix} = \begin{bmatrix} -3 & 1 & 0 & 1 \\ 1 & -2 & 1 & 0 \\ 1 & 1 & -1 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \frac{1}{s} & 0 & 0 & 0 \\ 0 & \frac{1}{s} & 0 & 0 \\ 0 & 0 & \frac{8}{s} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & 1 & -2 \\ 0 & 1 & 4 & -5 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

since the two static matrices are unimodular

$$\deg \left(P(s) E_{\mathcal{G}} K_e E_{\mathcal{G}} \right) = 3,$$

hence a pole at the origin was canceled. ▽

The result below states that such cancellations are present between the plant and the controller in Figure 5.2 whenever the conditions of Proposition 5.5.1 hold.

Proposition 5.5.6. *Let $P(s)$ and $K_e(s)$ be real rational and proper and let $\lambda \in \bar{\mathbb{C}}_0$ be a pole of $P(s)$.*

i) If (5.17a) *holds, then λ is canceled in $P(s)K(s)$.*

ii) If (5.17b) *holds, then λ is canceled in $K(s)P(s)$.*

Proof Bring in *minimal* realizations

$$P_i(s) = \left[\begin{array}{c|c} A_i & B_i \\ \hline C_i & D_i \end{array}\right] \quad \text{and} \quad K(s) = \left[\begin{array}{c|c} A_K & B_K \\ \hline C_K & D_K \end{array}\right]$$

so the realization

$$P(s) = \left[\begin{array}{c|c} A_P & B_P \\ \hline C_P & D_P \end{array}\right] := \left[\begin{array}{c|c} \operatorname{diag}\{A_i\} & \operatorname{diag}\{B_i\} \\ \hline \operatorname{diag}\{C_i\} & \operatorname{diag}\{D_i\} \end{array}\right]$$

is also minimal. To prove the first item of the Proposition it is then sufficient [104, Prop. 5.2] to show that λ is an uncontrollable mode of

$$P(s)K(s) = \left[\begin{array}{cc|c} A_K & 0 & B_K \\ B_P C_K & A_P & B_P D_K \\ \hline D_P C_K & C_P & D_P D_K \end{array}\right].$$

To this end, note that (5.5) implies $(\mathbb{1} \otimes I)' \begin{bmatrix} C_K & D_K \end{bmatrix} = 0$ and condition (5.17a) is equivalent to the existence of $0 \neq v \in \mathbb{C}^m$ such that $v = B_i'\eta_i$ for some η_i such that $\eta_i'(\lambda I - A_i) = 0$. The latter is equivalent to the existence of $\eta \neq 0$ such that

$$\eta'(\lambda I - A_P) = 0 \quad \text{and} \quad \eta' B_P = (\mathbb{1} \otimes v)'$$

for some $v \neq 0$. Therefore,

$$\begin{bmatrix} 0 & \eta' \end{bmatrix} \begin{bmatrix} A_K - \lambda I & 0 & B_K \\ B_P C_K & A_P - \lambda I & B_P D_K \end{bmatrix} = v'(\mathbb{1} \otimes I)' \begin{bmatrix} C_K & 0 & D_K \end{bmatrix} = 0$$

and the PBH test for the realization of PK fails for the mode at λ, proving the first item. The second item follows by similar arguments. ∎

Unstable pole-zero cancellations between a plant and a controller are a consensual taboo in feedback control. Textbooks treat them as a kind of a cardinal sin, which shall be avoided at all costs. The reason is that canceled dynamics do not really disappear. For example, poles of a SISO plant $P(s)$ canceled by zeros of a controller $K(s)$ always show up in the closed-loop disturbance sensitivity $T_{\mathrm{d}}(s)$, see (5.9). This is the very reason to require internal stability. Unstable cancellations due to deficient normal rank are less common and less studied. Nevertheless, they cause same repercussions. Namely, canceled dynamics shows up in at least one closed-loop relation, rendering the system prone to the effect of exogenous signals.

Assume, for example, that condition (5.17a), or (5.8a), holds for some $\lambda \in \bar{\mathbb{C}}_0$. It follows from the proof of Theorem 5.1 that there is then $v \neq 0$ such that (5.13) holds. Therefore,

$$\begin{bmatrix} 0 \\ \mathbb{1} \otimes v \end{bmatrix} \in \mathrm{zdir_o}\left(\begin{bmatrix} M_K & -N_P \\ -N_K & M_P \end{bmatrix}, \lambda \right) = \mathrm{pdir_i}(T_4, \lambda)$$

where the equality follows by Lemma C.1.11 and the fact that the factors in (5.10) are right coprime. By Lemma C.1.9 and (5.9)

$$T_4(s) \begin{bmatrix} 0 \\ \mathbb{1} \otimes v \end{bmatrix} = \begin{bmatrix} T_{\mathrm{d}}(s) \\ T(s) \end{bmatrix} (\mathbb{1} \otimes v)$$

has an unstable pole at $s = \lambda$. In other words, there is a load disturbance d_u in Figure 5.2 such that either y or u or both is unbounded. Likewise, it can be shown that if (5.17b) holds, then $\begin{bmatrix} S & T_{\mathrm{d}} \end{bmatrix} \notin H_\infty$, i.e. d_u or/and d_y might cause an unbounded y. This explains why the consensus protocol in Example 5.1.1 has an unstable load disturbance response.

It can be shown that if the consensus discussed in Example 5.1.1 can be attained, then all components of T_4 but T_{d} are stable, whereas $T_{\mathrm{d}}(s)$ has a pole at the origin. This agrees with the situation in SISO pole-zero cancellations discussed above. However, T_{d} is not necessarily unstable in a general MIMO case if either of the conditions in (5.17)

holds. The example below illustrates a different scenario.

Example 5.5.7. Consider a system with $\nu = 2$ agents

$$P_1(s) = \begin{bmatrix} s/(s+1) & 0 \\ 1/s & 1 \end{bmatrix} \quad \text{and} \quad P_2(s) = \begin{bmatrix} 1/s & 0 \\ 1 & s/(s+1) \end{bmatrix}$$

(both are second order). In this case there is only one edge. Select

$$K_{\mathrm{e}}(s) = K_{\mathrm{e},1}(s) = -\frac{1}{3}\begin{bmatrix} 1 & 0 \\ 0 & 2/s \end{bmatrix}.$$

It is then a matter of routine calculations to see that S, T_{d}, and T_{c} are stable, each having $(s+2)(2s+1)^2(3s+1)$ as the lowest common denominator of its entries. However, $T(s)$ has a pole at the origin in addition, rendering the whole T_4 unstable. ▽

Moreover, it may even happen that canceled dynamics of P are not excited by the (load) disturbance d_u, but rather only by d_y.

Example 5.5.8. Consider a system with $\nu = 2$ agents, yet again, now with the second order

$$P_1(s) = \begin{bmatrix} s/(s+1) & 1/s \\ 0 & 1 \end{bmatrix}, \quad P_2(s) = \begin{bmatrix} 1/s & 1 \\ 0 & s/(s+1) \end{bmatrix}$$

and the edge controller from Example 5.5.7. It can be calculated that in this case T, T_{d}, and T_{c} are stable, each having $(s+2)(2s+1)^2(3s+1)$ as the lowest common denominator of its entries. The sensitivity $S(s)$ has an additional pole at the origin. This implies that the responses to d_u are all stable, whereas the response of y to d_y is unstable. ▽

Remark 5.3. Stabilizability conditions for the setup in Figure 5.3 would be substantially different from those in Theorem 5.1 or Proposition 5.5.1. If we consider the class of LTI edge controllers K_{e}, then the stabilizability problem boils down to the question of existing decentralized fixed modes (DFMs) in P_{e} defined by (5.7), see [3, Sec. 2.2]. If controllers are allowed to be periodically time-varying, then even this condition is not restrictive [105]. However, this analysis has a snag in that the very construction of P_{e} might have unstable cancellations. For example, return to the case of $\nu = 3$ integrator agents with an undirected star interconnection graph discussed in Section 5.2. In this case $P(s) = (1/s)I_3$ has three poles at the origin, whereas

$$P_{\mathrm{e}}(s) = \begin{bmatrix} 1 & 0 & -1 \\ 0 & 1 & -1 \end{bmatrix}\left(\frac{1}{s}I_3\right)\begin{bmatrix} 1 & 0 \\ 0 & 1 \\ -1 & -1 \end{bmatrix} = \frac{1}{s}\begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$$

is a second-order transfer function. This P_{e} is easily stabilizable by decentralized edge controllers, e.g. by $K_{\mathrm{e}} = -I_2$. But this controller cannot see the canceled unstable mode, which remains a part of the closed-loop system. ▽

5.6 Concluding remarks

In this chapter we have studied the internal stability of MASs controlled by diffusively coupled laws. We have argued that internal stability, with entry points of exogenous signals at the connections between the agents and the controller, is a vital property in MASs and have proved that it can never be attained if the agents share common unstable dynamics, directions counted. In particular, this class always includes the case of homogeneous unstable agents or heterogeneous SISO agents with a common unstable pole, like an integral action. We have shown that the underlying reason for the lack of stabilizability is intrinsic cancellations of aligned unstable dynamics of agents by the diffusive coupling mechanism.

An immediate outcome of this analysis is that the consensus protocol described in Subsection 1.1.1 as well as its dynamic extensions cannot attenuate disturbance affecting the agreement mode, assuming all the agents share an unstable pole and direction. Since sharing such a pole is a necessary condition for reaching agreement [44], this implies that diffusive coupling is extremely *fragile* in non-ideal scenarios. Somehow, this uniformity must be carefully broken in order to avoid perturbations to the agreement mode. This is the underlying reason behind several of the different assumptions mentioned in Section 1.2 and Section 5.1. For example, in our notation the condition in [53, Prop. 7] reads

$$\left\|\bar{d}_u(t)\right\| \leq \gamma \ \forall t \geq 0, \quad \bar{d}_u(t) \coloneqq \frac{1}{\nu}\int_0^t (\mathbb{1}_\nu' \otimes I_n) d_u(s)\mathrm{d}s.$$

Since in that case the agents are integrators, the conditions requires that applying the projection of $d_u(t)$ on the agreement space as an input to an uncontrolled agent would result in a bounded trajectory. This agrees with our cancellations analysis, as the pole at the origin is always a source of instability.

Agreement, by definition, is an unstable phenomenon, as it requires the agents to converge to a non-zero trajectory from any initial conditions. Yet, there is more to gain from Theorem 5.1 beyond internal instability. The generic zero direction implies that in the all-ones direction there is no feedback. The diffusive controller cannot shape the trajectory or attenuate disturbances or noise inputs at all. This begs the question, is there an alternative architecture that can guarantee agreement, but still retain some measure of feedback in the all-ones direction. This is precisely the motivation and proposed solutions in the next chapter.

Chapter 6

Beyond Consensus: The Next Step?

> The most important step a man can take. It's not the first one, is it? It's the next one. Always the next step, Dalinar.
>
> Brandon Sanderson, *Oathbringer* [106]

Chapter 5 has exposed a crucial and intrinsic problem with diffusive controllers whenever external inputs are introduced. Similar observations have been noted in the literature, although without explanations and often brushed aside. Despite the focus on internal stability, there are other, less obvious, issues stemming from the aforementioned cancellations. Some have, to the best of our knowledge, never been clearly stated or explored in the literature. The most common class of diffusive controllers is by far variations of the consensus protocol, whose popularity exploded over the last twenty years. Hence, it would be instructive to conduct an input-output analysis of the consensus structure to illustrate some of the potential issues. Through such analysis, we can gain insight and propose modifications to improve performance.

6.1 The consensus protocol revisited

Consider a system comprised ν homogeneous SISO agents

$$\Sigma_i : y_i = P(u_i + d_i) + y_{0,i}, \quad \text{for all } i \in \mathbb{N}_\nu \tag{6.1}$$

where P is a given LTI model, u_i is a control input, d_i is a disturbance input, y_i is a measured regulated output, and $y_{0,i}$ is an initial condition response of the agent. The studied control problem for this system is to attain *consensus* among the agents, in the sense

$$\lim_{t\to\infty} y_i(t) = y_{\text{agt}}, \quad \text{for all } i \in \mathbb{N}_\nu \tag{6.2}$$

for an agreement variable $y_{\text{agt}} \in \mathbb{R}$, which is not prespecified. The consensus problem is not-trivial if $y_{\text{agt}} \neq 0$ and if the information about the neighbors is available to each Σ_i. A more general form of consensus protocol (1.3) at the ith agent may be expressed as the following diffusively-coupled law

$$u_i = k_i F \sum_{j \in \mathcal{N}_i} (y_j - y_i)$$

for some $k_i > 0$ and filter F, which are design parameters. Furthermore, it is not unreasonable to assume that measurements coming from neighboring agents are imperfect, e.g. corrupted by additive noise. In this case the control input

$$u_i = k_i F \sum_{j \in \mathcal{N}_i} (y_j + n_{ij} - y_i) \tag{6.3}$$

for some noise signals n_{ij}.

Assume that $\mathcal{G}$ is undirected and connected. It is then convenient to aggregate (6.1) and (6.3) for all indices i, which facilitates the use of the rich algebraic graph theory. So introduce the aggregate variables u, y, and y_0 (e.g. $u' = \begin{bmatrix} u_1 & \cdots & u_\nu \end{bmatrix}$), as well as the aggregate noise n, whose ith entry,

$$n_i = \sum_{j \in \mathcal{N}_i} n_{ij}, \tag{6.4}$$

sums up noises of all measurement channels of Σ_i. In this case the controlled system defines the relation

$$y = (I_\nu \otimes P)(u + d) + y_0 \tag{6.5}$$

and the consensus protocol reads

$$u = -KL_{\mathcal{G}} F y + KFn = -(KL_{\mathcal{G}} \otimes F) y + (K \otimes F) n \tag{6.6}$$

for $K \coloneqq \operatorname{diag}\{k_i\} > 0$, where $L_{\mathcal{G}}$ is the graph Laplacian of $\mathcal{G}$. A key technical result required to understand properties of the controlled systems is given in the lemma below.

Lemma 6.1.1. *The consensus protocol* (6.6) *results in*

$$y = U_K^{-1} \operatorname{diag}\{S_i\} U_K y_0 + U_K^{-1} \operatorname{diag}\{S_i P\} U_K d + U_K^{-1} \operatorname{diag}\{S_i P F\} U_K K n\big), \tag{6.7}$$

where S_i are systems with the transfer functions

$$S_i(s) \coloneqq \frac{1}{1 + \lambda_{K,i} P(s) F(s)} \tag{6.8}$$

(with $S_1 = 1$) and $U_K \in \mathbb{R}^{\nu \times \nu}$ is such that $U_K K^{1/2}$ is unitary and

$$U_K K L_{\mathcal{G}} U_K^{-1} = \operatorname{diag}\{\lambda_{K,i}\} \tag{6.9}$$

for $0 = \lambda_{K,1} < \lambda_{K,2} \leq \cdots \leq \lambda_{K,\nu}$.

Proof It is known from Proposition A.0.1 that for undirected and connected graphs $L_{\mathcal{G}} = L'_{\mathcal{G}} \geq 0$, with only a simple eigenvalue at the origin. This implies the existence of a sought U_K and the validity of properties of $\lambda_{K,i}$.

The rest of the proof essentially goes along developments in Proposition 2.1.2. Substituting (6.6) to (6.5) and using the mixed-product property from Proposition B.1.2, we end up with the controlled system

$$y = (I_\nu + KL_{\mathcal{G}} \otimes PF)^{-1}(y_0 + (I_\nu \otimes P)d + (I \otimes PF)Kn).$$

Note that (6.9) implies

$$KL_{\mathcal{G}} = U_K^{-1} \operatorname{diag}\{\lambda_{K,i}\} U_K,$$

and since $U_K = U_K \otimes 1$ applying the mixed-product property once more results in

$$(I_\nu + KL_{\mathcal{G}} \otimes PF)^{-1} = U_K^{-1} \operatorname{diag}\left\{\frac{1}{1 + \lambda_{K,i} P(s)F(s)}\right\} U_K,$$

which yields (6.7). ∎

By (6.7), the controlled dynamics decouples into ν independent subsystems

$$\eta_i' y = S_i \eta_i' y_0 + S_i P \eta_i' d + S_i PF \eta_i' Kn, \quad i \in \mathbb{N}_\nu$$

for every $\eta_i \in \operatorname{Im} U_K' e_i$. Choose

$$\eta_i \coloneqq \kappa U_K' e_i \quad \text{for } \kappa \coloneqq 1/\sqrt{\operatorname{tr}(K^{-1})}.$$

The signals $\eta_i' y$ reflect then important properties of the controlled system. Because $L_{\mathcal{G}} \mathbb{1} = 0$, U_k from (6.9) can be chosen such that $U_K^{-1} e_1 = \kappa \mathbb{1}$. Hence, $\eta_1 = \kappa U_K'(\kappa U_K \mathbb{1}) = \kappa^2 K^{-1} \mathbb{1}$ and

$$\eta_1' y = \kappa^2 \mathbb{1}' K^{-1} y = \sum_{i=1}^{\nu} \frac{y_i}{\operatorname{tr}(K^{-1}) k_i} =: \bar{y}_K.$$

This is a weighted average of the outputs y_i of all agents. It becomes the standard average if all gains k_i are equal.

Each $\eta_i' y$ for $i \geq 2$ is not as transparent as $\eta_1' y$. Yet the sum of their squares

$$\Delta \coloneqq \sum_{i=2}^{\nu} (\eta_i' y)^2 = \|\kappa U_K y\|^2 - \left\|\eta_1' y\right\|^2,$$

is meaningful. Indeed, by the unitary property of $U_K K^{1/2}$, we have $\|\kappa U_K y\|^2 = \left\|\kappa K^{-1/2} y\right\|^2$. Hence,

$$\Delta = \kappa^2 y' K^{-1/2} (I - \kappa^2 K^{-1/2} \mathbb{1}\mathbb{1}' K^{-1/2}) K^{-1/2} y.$$

Because $I - \kappa^2 K^{-1/2} \mathbb{1}\mathbb{1}' K^{-1/2}$ is an orthogonal projection matrix [65, Thm. 7.5], we end

up with

$$\Delta = \|\kappa K^{-1/2}(y - \mathbb{1}\bar{y}_K)\|^2 = \sum_{i=1}^{\nu} \frac{(y - \bar{y}_K)^2}{\operatorname{tr}(K^{-1})k_i}. \tag{6.10}$$

Thus, Δ can quantify the deviation of y_i from their weighted average, i.e. the disagreement between Σ_i.

The previous discussion implies that if $\eta_i' y$ vanish asymptotically for all $i \geq 2$, then the agents agree on the trajectory $\bar{y}_K$. This prompts the requirement to render all S_i for those indices decaying. At the same time, $S_1 = 1$ regardless of the choice of F. Thus,

$$\bar{y}_K = \bar{y}_{0,K} + \kappa^2 P \mathbb{1}' K^{-1} d + \kappa^2 P F \mathbb{1}' n \tag{6.11}$$

where $\bar{y}_{0,K} \coloneqq \kappa^2 \mathbb{1}' K^{-1} y_0$ is the weighted average of the responses of the agents to their initial conditions, $\kappa^2 \mathbb{1}' K^{-1} d$ is the weighted average disturbance, and $\mathbb{1}' n$ is the cumulated noise. In the conventionally assumed disturbance and noise free setting, consensus (6.2) is attained iff all initial condition responses $y_{0,i}$ converge, in which case $y_{\text{agt}} = \lim_{t\to\infty} y_{0,K}(t)$. Moreover, the disagreement Δ typically vanishes faster under larger gains k_i. However, the external signals change this situation dramatically.

6.1.1 Performance limitations of the consensus protocol

The best studied version of the consensus problem is that for $P(s) = 1/s$. The initial condition responses $y_{0,i}$ are then constant for all i. It is common to have $K = kI_\nu$ and $F = 1$. With these choices, $\lambda_{K,i} = k\lambda_i$, where λ_i are the eigenvalues of the Laplacian $L_{\mathcal{G}}$ sorted increasingly, the rows of U_K comprise an orthonormal eigenbasis of $L_{\mathcal{G}}$,

$$S_i(s) = \frac{s}{s + k\lambda_i} \quad \text{and} \quad S_i(s)P(s) = \frac{1}{s + k\lambda_i}$$

are stable for all $k > 0$, $\bar{y}_K = (1/\nu)\sum_{i=1}^{\nu} y_i$, and $\kappa^2 = k/\nu$.

If $d = n = 0$, then $\bar{y}_K(t) = (1/\nu)\sum_{i=1}^{\nu} y_{0,i}$ for all $t \geq 0$ and $\eta_i' y(t) = (1/\sqrt{\nu})\mathrm{e}^{-k\lambda_i t} e_i' U_K y_0$. Consequently,

$$\int_0^\infty \left\|y(t) - \mathbb{1}\bar{y}_{0,K}\right\|^2 \mathrm{d}t = \nu \int_0^\infty \Delta(t)\mathrm{d}t = \frac{1}{k}\sum_{i=2}^{\nu} \frac{\left\|e_i' U_K y_0\right\|^2}{2\lambda_i},$$

meaning that the energy of disagreement is inversely proportional to the gain k. In other words, increasing k reduces the disagreement between the agents.

This analysis is different from the standard ones in the literature, yet arrives to the same conclusion: increasing k improves the convergence rate of the system. If n or d are non-zero, however, the analysis changes in two major ways.

1. **Internal stability is compromised.** In this case $\bar{y}_K$ satisfies

$$\frac{\mathrm{d}\bar{y}_K(t)}{\mathrm{d}t} = \frac{1}{\nu}\sum_{i\in\mathbb{N}_\nu} \left(d_i(t) + kn_i(t)\right),$$

which is directly proportional to the average disturbance and the cumulative noise. Hence, if either of the averages is persistently non-zero the agreement variable will diverge. This is a special case of the results of Chapter 5, where the canceled integrator pole reappears. Similar results extend to general agent dynamics with unstable/marginally stable poles.

2. **There is a trade-off between noise sensitivity and nominal performance.** Assume that n_{ij} are independent zero-mean white processes with unit intensity and $d_i(t) = 0$. From the above equation, it is clear that $\bar{y}_K(t)$ is a Wiener process, whose mean is $\bar{y}_{0,K}$ and steady-state variance grows unbounded, proportionally to $k^2 t$. Likewise, consider the steady-state variance of disagreement, understood as the mean value of $\lim_{t\to\infty} \nu\Delta(t)$. In this case n_i are also white, with the intensity $|\mathcal{N}_i|$ and the covariance of n is $D_{\mathcal{G}}\delta(t-s)$. Because the effect of the initial conditions on $\eta_i' y$ vanishes, the steady-state variance of disagreement equals [107, Thm. 1.53] the square H_2-norm of

$$T_n = \kappa \begin{bmatrix} S_2 P & & 0 \\ & \ddots & \\ 0 & & S_\nu P \end{bmatrix} \begin{bmatrix} e_2' \\ \vdots \\ e_\nu' \end{bmatrix} U_K K D_{\mathcal{G}}^{1/2}, \tag{6.12}$$

where $D_{\mathcal{G}}$ is the degree matrix of $\mathcal{G}$. It is readily verified that

$$\|T_n\|_2^2 = k\kappa^2 \sum_{i=2}^{\nu} \frac{e_i' U_K D_{\mathcal{G}} U_K' e_i}{2\lambda_i}.$$

Therefore, noise sensitivity, both of the agreed variable and disagreement, deteriorates as k increases.

As established, the first issue is unavoidable when using consensus-like (and more generally, diffusive) controllers as it depends solely on the agent dynamics. Moreover, if we wish to reach non-trivial agreement this way the agents must have unstable poles. To illustrate this, consider now general P and F. In the Laplace domain $\bar{y}_K(s)$ satisfies

$$\bar{y}_K(s) = \frac{1}{\nu}\left(\mathbb{1}' y_0 + P(s)\mathbb{1}' d(s) + kP(s)F(s)\mathbb{1}' n(s)\right).$$

This represents uncontrolled initial conditions and disturbance responses, while the noise acts on a series interconnection of the uncontrolled dynamics $P(s)$ and controller $kF(s)$. Hence, to reach non-trivial agreement through controller (6.3) the uncontrolled response to initial conditions cannot converge to 0, ergo P must be unstable.

Even if we ignore stability concerns, the second issue also poses a problem. Say we assume that external signals cannot excite the unstable modes, hence we do not need to worry about unboundedness of the agreement trajectory. In this case we may want to choose k and F to somehow attenuate the noise. Revisiting the integrator consensus example, it is clear that unless $F(0) = 0$, $\bar{y}_K$ is still a Wiener process. This, however, is

not legitimate as it would result in non-decaying $S_i \eta_i' y_0$. In general, choosing F is not quite trivial since it affects $\bar{y}_K$ in an open-loop fashion and the other $\eta_i' y$ via feedback. For example, a naïve choice of first-order low-pass F with $F(s) = 1/(\tau s + 1)$, which still results in stable S_i, would increase each term in the expression of $\|T_n\|_2^2$ above by a factor of $1 + k\tau\lambda_i > 1$.

6.1.2 A different perspective

Perhaps the fundamental difficulty of understanding the consensus protocol, is that it attempts to solve a problem that is fundamentally different from those usually studied in control theory. First and foremost, reaching agreement from initial conditions is inherently an unstable phenomenon - a taboo in classical control problems. Second, the problem is often analyzed without external inputs. Even when there are external inputs, they are by large harmful signals to be attenuated. Consequently, many classical tools such as loop-shaping are absent from the realm of multi-agent systems.

Yet throughout this work we have alluded to the similarities between agreement and tracking, and between the consensus protocol to error feedback. In fact, we have seen in Example 2.2.1 that this is exactly the case for complete undirected graphs. Motivated by this, consider a slightly different outlook on this structure can be obtained by rewriting (6.3) as

$$u_i = -k_i F\,|\mathcal{N}_i|\,\big(y_i - \bar{y}_i + \frac{1}{|\mathcal{N}_i|} n_i\big), \quad \text{where } \bar{y}_i \coloneqq \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} y_j, \tag{6.3′}$$

and n_i is as in (6.4). This form is reminiscent of a servo problem in unity feedback where only the error is supplied to the controller [108, Sec. 1.3]. In this perspective, $k_i\,|\mathcal{N}_i|$ is the local feedback gain and $\bar{y}_i$, which is the average of measured neighbors, is the "reference" signal. A similar viewpoint was first proposed in [12, § III.A], and indeed agreement is achieved if and only if the underlying graph is connected and all the agents simultaneously solve this tracking problem.

Balancing the inherent tradeoffs between performance and robustness are the bread and butter of classical control. In servo-regulation, for example, tracking performance and robustness are difficult to simultaneously balance. This is, in part, a product of the error-based unity feedback control architecture. To see why, note that the output equation of the unity-feedback configuration in Figure 6.1(a) reads

$$y = (I + PR)^{-1}(PRr - PRn + Pd) \coloneqq Tr - Tn + T_d d.$$

Loop-shaping arguments would require that in the frequency domain $|T(\mathrm{j}\omega)| \approx 1$ where the spectrum of r is concentrated, $|T(\mathrm{j}\omega)| \ll 1$ where the spectrum of n is concentrated, and $|T_d(\mathrm{j}\omega)| \ll 1$ where the spectrum of d is concentrated. Yet all of these functions are coupled with only one tuning parameter, R. Hence, if the spectra of the external

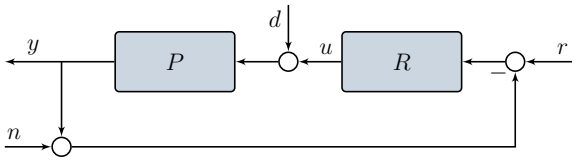

(a) Error-based unity feedback.

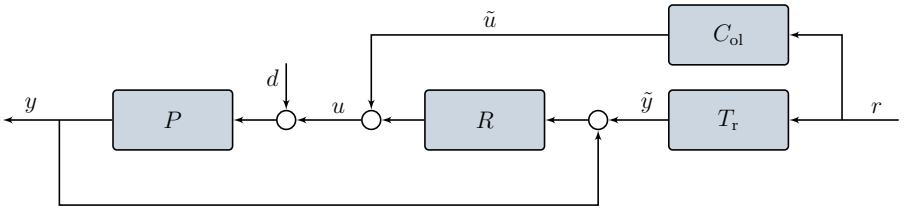

(b) Two Degrees of Freedom control.

Figure 6.1: Classical servo-regulation control architectures: 1 degree-of-freedom and 2 degrees-of-freedom.

signal are not well separated there are strict limitations on achievable performance.

One classical solution to this is employing a two degrees-of-freedom (2DOF) architecture, using separate controllers for outputs and reference signals. Indeed, variations of 2DOF architectures, first introduced over 70 years ago [54], have been extensively studied [108, Sec. 2.9]. The term "two-degrees-of-freedom control" is used to refer to several slightly different control architectures. Here we consider the architecture shown in Figure 6.1(b), which can completely decouple the disturbance and tracking design. This is accomplished by designing the control law in the following fashion. First, define a signal $\tilde{y}$ to represent the required output behavior. This is often modeled as the response of some system T_r to the reference signal r. Then, the signal $\tilde{u}$ is designed to achieve this behavior in an open-loop fashion, i.e. satisfying

$$\tilde{y} = P\tilde{u} = T_\mathrm{r} r.$$

If designed correctly, the control law reads

$$u = \tilde{u} + R(y - P\tilde{u}) = Ry + (I - RP)\tilde{u},$$

and using the identities $\tilde{y} = P\tilde{u}$ and $(I - RP)^{-1}P = P(I - PR)^{-1}$ we obtain

$$y = \tilde{y} + (I - PR)^{-1} P d.$$

Therefore, the disturbance response depends only on R, and the tracking response on $\tilde{u} = P^{-1} T_\mathrm{r} r$.

It is reasonable to assume that, as in servo regulation, there might be an alternative architecture to consensus that will simplify the design. The next section outlines some

preliminary work in this vein, motivated by the analogy to the servo problem.

6.2 A two-degrees-of-freedom approach

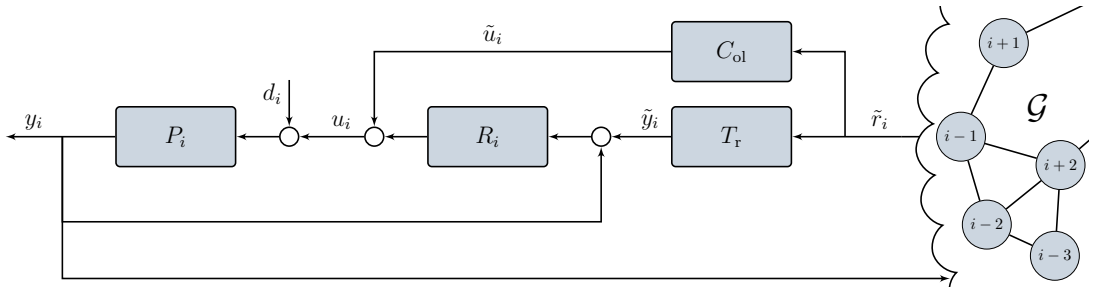

Figure 6.2: A Two-Degrees-of-Freedom consensus protocol inspired by Figure 6.1(b) with a network-generated reference signal.

Motivated by the parallels to tracking problems, we wish to derive a 2DOF variant of the consensus protocol, as illustrated in Figure 6.2. To this end, consider a system comprised ν possibly heterogeneous $p \times m$ agents

$$\Sigma_i : y_i = P_i(u_i + d_i) + y_{0,i}, \quad \text{for all } i \in \mathbb{N}_\nu \tag{6.13}$$

where P_i is a given LTI model, u_i is a control input, d_i is a disturbance input, y_i is a measured regulated output, and $y_{0,i}$ is an initial condition response of the ith agent.

The main difference between servo-regulation and agreement problems is that the latter lacks a well-defined reference signal. Consequently, there are various ways to select both r and $\tilde{y}$. It can be done in an open-loop way, with the agents exchanging controller variables and agreeing on their common rendezvous point or trajectory. This would lead to a control structure akin to that of [44], where the agents exchange the state of some common internal model. Alternatively, it can be done in a closed-loop way, where $\tilde{y}$ is generated using only the measured neighbors outputs. Consider the latter approach, which can be motivated by the classic consensus protocol. To this end, assume that $F = 1$ and $n_i = 0$, then the aggregate form of (6.3′) reads

$$u = -(KD_{\mathcal{G}} \otimes I)(y - \bar{y}), \quad \text{where } \bar{y} \coloneqq (A^{\star}_{\mathcal{G}} \otimes I_p)y$$

and $K \coloneqq \mathrm{diag}\{k_i\}$. The above indicates that for consensus $r = \bar{y}$, which is already distributed according to the graph structure. A standing assumption is that the noise, n, is generated at the network level. Hence, we assume it is additive to $\bar{y}$. This implies that a good model for the required behavior is a *filtered* version of $\bar{y}$, i.e.

$$\tilde{y} = (I \otimes T_{\mathrm{r}})\big((A^{\star}_{\mathcal{G}} \otimes I)y + n\big) \tag{6.14}$$

where T_r is the additional degree of freedom and n represents additive noise induced by the network (a la (6.4)). Note that we assume a uniform filter T_r for all the agents, even heterogeneous ones.

Let $P \coloneqq \mathrm{diag}\{P_i\}$ and $R \coloneqq \mathrm{diag}\{R_i\}$ denote the aggregate plant and local controller respectively. Substituting (6.14) into the 2DOF architecture yields the *2DOF consensus protocol*

$$u = Ry + (I_{\nu m} - RP)\,\tilde{u}, \quad \tilde{u} = P^{-1}\tilde{y} \tag{6.15a}$$

and resulting closed-loop dynamics

$$y = U_D^{-1}\,\mathrm{diag}\{\hat{S}_i\}U_K\,(T_\mathrm{d}d + (I \otimes T_\mathrm{r})n + Sy_0) \tag{6.15b}$$

with

$$\hat{S}_i \coloneqq \left(I_p - \alpha_i T_\mathrm{r}\right)^{-1}, \quad T_\mathrm{d} \coloneqq (I_{\nu p} - PR)^{-1}P, \quad \text{and} \quad S \coloneqq (I_{\nu p} - PR)^{-1}.$$

and $U_D \in \mathbb{R}^{\nu\times\nu}$ is such that $U_D D_\mathcal{G}^{-1/2}$ is unitary

$$U_D A_\mathcal{G}^\star U_D^{-1} = \mathrm{diag}\{\alpha_i\}.$$

Note that the control signal has a *decentralized*, i.e. block-diagonal, component while the other component is *distributed* with respect to the communication graph $\mathcal{G}$ since

$$U_D^{-1}\,\mathrm{diag}\{\hat{S}_i\}U_D = \left(I_{\nu p} - A_\mathcal{G}^\star \otimes T_\mathrm{r}\right)^{-1}.$$

Recall that in agreement problems the objective is to reach agreement driven only by initial conditions, thus the nominal performance of (6.15) depends on $U_D^{-1}\,\mathrm{diag}\{\hat{S}_i\}U_D S$. Like in the control law, the nominal dynamics have a decentralized component in S, and a distributed component in $U_D^{-1}\,\mathrm{diag}\{\hat{S}_i\}U_D$. Similarly, the disturbance dynamics depend on the distributed $U_D^{-1}\,\mathrm{diag}\{\hat{S}_i\}U_D$ and on the decentralized T_d.

Unlike 1DOF consensus protocols, the closed-loop dynamics in (6.15b) explicitly separate the network and local dynamics. The network component depends only on T_r, which is uniform across agents, while local dynamics captured by S and T_d are decentralized by construction. This inherent separation naturally accommodates agent heterogeneity provided that their network component, T_r, is homogeneous. The following theorem formalizes these observations.

Theorem 6.1. *Consider heterogeneous agents driven only by initial conditions, interacting over an undirected and connected graph $\mathcal{G}$, and controlled by* (6.15b)*. If each local controller R_i stabilizes its corresponding plant P_i, then the agents reach asymptotic agreement if and only if*

$$\hat{S}_i \coloneqq \left(I_p - \alpha_i T_r\right)^{-1} \in H_\infty, \quad \forall \alpha_i \in \mathrm{spec}\, A_\mathcal{G}^\star \setminus \{1\}$$

and $\hat{S}_1 = \left(I_p - T_r\right)^{-1}$ has all poles in the closed left half-plane.

Proof By assumption $\mathcal{G}$ is undirected and connected, therefore from Proposition A.0.1 there is a nonsingular U_D such that

$$U_D A_{\mathcal{G}}^{\star} U_D^{-1} = \operatorname{diag}\{\alpha_i\},$$

and since $A_{\mathcal{G}}$ is symmetric there exists one such that $U_D D_{\mathcal{G}}^{-1/2}$ is unitary. Defining $\hat{y} = (U_D \otimes I)y$ and pre-multiplying (6.15b) by $U_D \otimes I$ yields

$$\hat{y} = \left(I_{\nu p} - \operatorname{diag}\{\alpha_i\} \otimes T_{\mathrm{r}}\right)^{-1} \hat{y}_0, \text{ with } \hat{y}_0 \coloneqq (U_D \otimes I) S y_0.$$

By assumption R_i internally stabilizes P_i, therefore $\hat{y}_0(t)$ is bounded and asymptotically decays to zero. Now the system from input $\hat{y}_0(t)$ to $\hat{y}$ is a block-diagonal system, therefore each $\hat{y}_i$ depends only on $\hat{y}_{0,i}$ as

$$\hat{y}_i = \hat{S}_i \hat{y}_{0,i}.$$

For the first direction, assume that $\hat{S}_i$ is stable for all $\alpha_i \neq 1$ and that for $\alpha_1 = 1$ all of its poles are in the closed left half-plane. Then, for every $\epsilon > 0$ there exists a time $t_\epsilon \geq 0$ such that for all $t > t_\epsilon$

$$\|\hat{y}(t) - e_1 \otimes \hat{y}_1(t)\| < \epsilon,$$

where $\hat{y}_1(t)$ is the time response of the first block of $\hat{y}$. Since no coordinate of $\hat{y}$ diverges exponentially, the transformations are well defined and invertible. Returning to the original coordinates, we obtain

$$\|y(t) - \mathbb{1}_\nu \otimes \hat{y}_1(t)\| < \epsilon$$

because we can choose U_D such that $U_D^{-1} e_1 = \mathbb{1}_\nu$.

For the other direction, suppose the agents reach asymptotic agreement. Then there exists a trajectory $y_{\mathrm{agt}}(t)$ such that, for all $\epsilon > 0$, there is a $t_\epsilon \geq 0$ with

$$\left\|y(t) - \mathbb{1}_\nu \otimes y_{\mathrm{agt}}(t)\right\| < \epsilon \quad \forall t > t_\epsilon.$$

The remainder of the proof follows by reversing the above steps. ∎

Theorem 6.1 provides clear conditions for agreement but does not explicitly specify the resulting agreement trajectory. Since S is stable, the trajectory is determined solely by the unstable poles of $\hat{S}_1$. Hence, T_{r} must be designed to both solve a simultaneous stabilization problem against the eigenvalues of $A_{\mathcal{G}}^{\star}$ and satisfy certain interpolation constraints.

Still, pole cancellations can occur in the series interconnection $U_D^{-1} \operatorname{diag}\{\hat{S}_i\} U_D S$, altering the agreement trajectory. Such cancellations, however, are outside the feedback loop and thus do not jeopardize stability. The following proposition provides a simple

necessary condition for these cancellations.

Proposition 6.2.1. *Let R, P, and T_r be finite-dimensional systems, and denote by p_i the imaginary-axis poles of $\hat{S}_1$. If p_i is not a pole of $U_D^{-1}\operatorname{diag}\{\hat{S}_i\}U_D S$, then it must be a zero of S_i for all i.*

Proof It is known [104, Prop. 5.2] that given a cascade interconnection $U_D^{-1}\operatorname{diag}\{\hat{S}_i\}U_D S$, a pole p_i of $U_D^{-1}\operatorname{diag}\{\hat{S}_i\}U_D$ is cancelled if and only if

$$\operatorname{pdir}_{\mathrm{i}}\Big(U_D^{-1}\operatorname{diag}\{\hat{S}_i\}U_D, p_i\Big) \cap \operatorname{zdir}_{\mathrm{o}}(S, p_i) \neq \{0\}.$$

Bring in a minimal realization (A, B, C, D) of $\operatorname{diag}\{\hat{S}_i\}$, By definition

$$\operatorname{diag}\{\hat{S}_i\} = D + C(sI - A)^{-1}B \iff U_D^{-1}\operatorname{diag}\{\hat{S}_i\}U_D = \hat{D} + \hat{C}(sI - A)^{-1}\hat{B}$$

where

$$\hat{D} = (U_D^{-1} \otimes I)D(U_D \otimes I),\ \hat{C} = (U_D^{-1} \otimes I)C, \text{ and } \hat{B} = B(U_D \otimes I).$$

Now let p_i be an unstable pole of $U_D^{-1}\operatorname{diag}\{\hat{S}_i\}U_D$, then it must be a pole only of $\hat{S}_1$ since the other components are stable. Thus

$$v_i \in \operatorname{pdir}_{\mathrm{i}}\Big(\hat{S}_1, p_i\Big) \implies (e_1 \otimes v_i) \in \operatorname{pdir}_{\mathrm{i}}\Big(\operatorname{diag}\{\hat{S}_i\}, p_i\Big),$$

and this is true for all unstable p_i. By definition

$$\begin{aligned}\operatorname{pdir}_{\mathrm{i}}\Big(U_D^{-1}\operatorname{diag}\{\hat{S}_i\}U_D, p_i\Big) &= \hat{B}'\ker(p_i I - A)' \\ &= (U_D' \otimes I)B'\ker(p_i I - A)' \\ &= (U_D' \otimes I)\operatorname{pdir}_{\mathrm{i}}\Big(\operatorname{diag}\{\hat{S}_i\}, p_i\Big),\end{aligned}$$

and consequently

$$(e_1 \otimes v_i) \in \operatorname{pdir}_{\mathrm{i}}\Big(\operatorname{diag}\{\hat{S}_i\}, p_i\Big) \iff ((U_D' e_1) \otimes v_i) \in \operatorname{pdir}_{\mathrm{i}}\Big(U_D^{-1}\operatorname{diag}\{\hat{S}_i\}U_D, p_i\Big).$$

We know that $U_D' e_1 = \gamma$, where γ is the normalized left eigenvector associated with $\alpha_1 = 1$. Moreover, we know that

$$\gamma' = \frac{1}{\sqrt{\operatorname{tr}(D_{\mathcal{G}})}}\mathbb{1}_\nu' D_{\mathcal{G}},$$

and $D_{\mathcal{G}}$ is a diagonal matrix with only positive entries, therefore all the components of γ are nonzero. This implies that for p_i to be canceled we must have $(\gamma \otimes v_i) \in \operatorname{zdir}_{\mathrm{o}}(S, p_i)$, since all the coordinates of γ are non-zero, this implies that p_i must be zero of all S_i.■

Proposition 6.2.1 has an important implication for robustness. In consensus-like protocols, the resulting agreement trajectory is generally vulnerable to persistent disturbances in the agreement direction, as these disturbances excite common unstable

poles. Intentionally introducing heterogeneity in local controllers, however, can exploit the cancellation properties described above to improve robustness against such disturbances. This insight follows a conjuncture made in the concluding remarks of Chapter 5, and is demonstrated in the following example.

Example 6.2.2. Consider a group of ν identical SISO agents, each with a pole at the origin, aiming to achieve consensus. Standard consensus protocols with disturbance-rejection mechanisms (for example [50]) typically cannot reject step disturbances, resulting in linear divergence of outputs. Suppose one agent, say agent 1, is not affected by DC disturbances. By intentionally designing local controllers such that R_i is a PI controller for all $i > 1$, all $T_{\mathrm{d},i}$ (except the first) have a zero at the origin, thus effectively rejecting DC disturbances. Despite this heterogeneity, agents still reach consensus since, according to Proposition 6.2.1, no cancellations between the network and local dynamics occur. ▽

In the absence of cancellations in $U_D^{-1}\operatorname{diag}\{\hat{S}_i\}U_D S$, the agreement trajectory is entirely determined by T_{r}, as is the response to network noise n. Consequently, the design of the network filter T_{r} depends only on the graph $\mathcal{G}$, the desired agreement trajectory, and the spectrum of network-induced noise. Crucially, it does not depend on the agents' dynamics, P. This significantly simplifies the design of T_{r} for prescribed agreement trajectories, as illustrated in the following example.

Example 6.2.3. Assume for simplicity that the agents are SISO, and that their goal is to reach consensus, i.e. agreement to a constant. Assuming that the are no unstable cancellations in the local loops, by Theorem 6.1 reaching consensus depends only on T_{r} and the eigenvalues of $A_{\mathcal{G}}^{\star}$. Since we know that $\operatorname{spec} A_{\mathcal{G}}^{\star} \in [-1, 1]$ (c.f Proposition A.0.1), we can easily derive some simple prototype network filters to guarantee consensus for any connected graph. Some examples are provided below.

1. **First order low-pass:** we have

$$T_{\mathrm{r}}(s) = \frac{k}{\tau s + 1} \implies \hat{S}_i = (1 - \alpha_i T_{\mathrm{r}}(s))^{-1} = \frac{\tau s + 1}{\tau s + 1 - \alpha_i k}.$$

   Clearly for $k = 1$ $\hat{S}_1$ always has a simple pole at the origin, and $\hat{S}_i$ for $i > 1$ will be stable for any $\tau > 0$ and any undirected graph.

2. **Second order low-pass:** similarly, for arbitrary 2nd order filter we have

$$T_{\mathrm{r}}(s) = \frac{k}{(\tau_1 s + 1)(\tau_2 s + 1)} \implies \hat{S}_i(s) = \frac{(\tau_1 s + 1)(\tau_2 s + 1)}{\tau_1 \tau_2 s^2 + (\tau_1 + \tau_2)s + (1 - \alpha_i k)},$$

   which again requires $k = 1$ to ensure an integrator and is otherwise stable for any graph and $\tau_i > 0$. Note that the same holds for a second order underdamped

system since

$$T_{\mathrm{r}}(s) = \frac{k\omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2} \implies \hat{S}_i(s) = \frac{s^2 + 2\zeta\omega_n s + \omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2(1 - \alpha_i k)}.$$

3. **Third order system:** here we have

$$T_{\mathrm{r}}(s) = \frac{k\omega_n^2}{(\tau s + 1)(s^2 + 2\zeta\omega_n s + \omega_n^2)} \implies$$
$$\hat{S}_i(s) = \frac{(\tau s + 1)(s^2 + 2\zeta\omega_n s + \omega_n^2)}{\tau s^3 + (2\zeta\omega_n\tau + 1)s^2 + (\tau\omega_n^2 + 2\zeta\omega_n)s + \omega_n^2(1 - \alpha_i k)}.$$

Applying the Routh-Hurwitz criterion results in

$$\omega_n(2\zeta(\tau\omega_n)^2 + (4\zeta^2 + 1)\tau\omega_n + 2\zeta) > \tau(1 - \alpha_i)$$

under the assumption that $k = 1$ and that all parameters are positive. This case is slightly more challenging, but can still be addressed with some conservatism.

Define

$$f(\tau\omega_n) = 2\zeta(\tau\omega_n)^2 + (4\zeta^2 + 1)\tau\omega_n + 2\zeta.$$

The polynomial $f(\tau\omega_n)$ has roots at -2ζ and $-1/(2\zeta)$; hence for $\tau\omega_n > 0$ it is strictly positive. Since $(1 - \alpha_i) \leq 2$ and $f(0) = 2\zeta$, the original inequality is satisfied if

$$2\zeta\omega_n > 2\tau,$$

which is easily achieved. ▽

Similar procedures can be carried out for more complex filters, such as ones with zeros, as well as for other agreement trajectories.

Example 6.2.4. Assume that the goal is to reach agreement on some harmonic signal at frequency ω_0. All $T_{\mathrm{r}}(s)$ ensuring that $\hat{S}_1(s)$ has poles at $\pm\mathrm{j}\omega_0$ can be parametrized via

$$T_{\mathrm{r}}(s) = 1 + \frac{s^2 + \omega_0^2}{(s+1)^2} Q(s)$$

for any stable and proper $Q(s)$. Ensuring stability, however, is more difficult. Rewriting $\hat{S}_i(s)$ we have

$$\hat{S}_i(s) = \frac{1}{1 - \alpha_i} \frac{1}{1 - \frac{\alpha_i}{1-\alpha_i} \frac{s^2+\omega_0^2}{(s+1)^2} Q(s)}$$

which is equivalent to the following robust control problem:

$$\text{Find } Q \in H_\infty \text{ to stabilize} \quad \tilde{P}(s) = k\frac{s^2 + \omega_0^2}{(s+1)^2}, \quad k \in [-0.5, \infty).$$

Optimizing on Q may not be trivial, but it follows that any static Q satisfying

$$Q < \frac{1}{-0.5\omega_0^2}$$

solves the equivalent problem. ▽

Finally, to see the potential of the method in combating noise, consider the counterpart of signals $\eta_i' y$ as in Section 6.1. Note that $A_{\mathcal{G}}$ is symmetric just like $L_{\mathcal{G}}$, hence we can construct vectors η_i that decouple $\eta_i' y$ into ν independent subsystem by repeating the procedure we did for $KL_{\mathcal{G}}$ but for $A_{\mathcal{G}}^\star = D_{\mathcal{G}}^{-1} A_{\mathcal{G}}$. Namely, we may choose

$$\eta_i \coloneqq \kappa_2 U_D' e_i, \quad \kappa_2 \coloneqq 1/\sqrt{\mathrm{tr}(D_{\mathcal{G}})}$$

resulting in

$$\eta_i' y = \hat{S}_i \left(e_i' U_D T_{\mathrm{d},i} d + T_{\mathrm{r}} \hat{n} + e_i' U_D S y_0 \right), \quad \hat{n}_i \coloneqq e_i' U_D n \quad i \in \mathbb{N}_\nu. \tag{6.16}$$

In general $e_i' U_D$ is not known, hence each $\eta_i' y$ depends on some linear combination of the local loops. There are two notable exception to this, one being the agreement direction and the other the noise response. For the agreement direction we can choose U_D such that

$$U_D^{-1} e_1 = \kappa_2 \mathbb{1} \quad \text{and} \quad e_1' U_D = \kappa_2 \mathbb{1}' D_{\mathcal{G}},$$

resulting in

$$\eta_1' y = \kappa_2^2 \hat{S}_1 \mathbb{1}' D_{\mathcal{G}} \left(T_{\mathrm{d}} d + (I \otimes T_{\mathrm{r}}) n + S y_0 \right).$$

Assuming that $D_{\mathcal{G}}$ is globally known, this implies that we can calculate precisely the dynamics generating $\eta_1' y$. As for the noise response, since T_{r} is uniform among all agents things simplify in similar fashion to classical consensus. Namely, we can construct an explicit expression for the steady-state variance of the disagreements as in (6.12) through the squared H_2 norm of

$$\hat{T}_n = \kappa_2 \begin{bmatrix} \hat{S}_2 T_{\mathrm{r}} & & 0 \\ & \ddots & \\ 0 & & \hat{S}_\nu T_{\mathrm{r}} \end{bmatrix} \begin{bmatrix} e_2' \\ \vdots \\ e_\nu' \end{bmatrix} U_D D_{\mathcal{G}}^{1/2}. \tag{6.17}$$

Similar standard consensus, each $\hat{S}_i$ depend on the corresponding α_i, and to minimize (6.17) we must a-priori know the graph just like in (6.12). However, the noise response in the agreement direction always depends on $\hat{S}_1 T_{\mathrm{r}}$ which is independent of the graph. This is a stark contrast with 1DOF agreement protocols, where there was no feedback at all in the agreement direction. Using the 2DOF architecture we can directly design T_{r} to attenuate the effects of the noise in the agreement direction, and still implicitly attenuate the disagreement variance.

6.3 Numerical examples

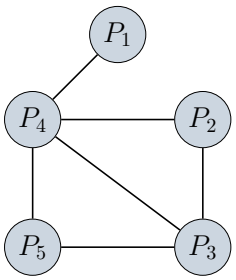

Figure 6.3: The underlying communication graph for the examples in Section 6.3.

The following two examples illustrate the flexibility and potential of the 2DOF protocol. In all examples, a group of $\nu = 5$ integrator agents, $P_0 = 1/s$, attempts to reach static consensus. The agents interact over the undirected graph shown in Figure 6.3, whose spectral properties are given by

$$\gamma = \frac{1}{12}\begin{bmatrix} 1 \\ 2 \\ 3 \\ 4 \\ 2 \end{bmatrix}, \quad \operatorname{spec} D_{\mathcal{G}}^{-1} A_{\mathcal{G}} = \left\{ \frac{\pm\sqrt{33} - 3}{12}, -0.5, 0, 1 \right\}.$$

In both examples we compare two architectures: (i) classic 1DOF consensus using (6.6), and (ii) the 2DOF protocol (6.15). The controllers are tuned to achieve similar nominal performance as measured by settling time. Then, we show how we can improve the behavior of the agreement mode under non-nominal conditions.

6.3.1 Attenuating agreement mode drift

One of the main issues with the consensus protocol is the lack of feedback in the agreement direction as illustrated in (6.11). When afflicted with white-noise as in Subsection 6.1.1, this imposes a direct trade-off between the nominal performance and the variance of the output. In particular, when considering the agreement variable under white-noise we see that the variance diverges at a rate proportional to the gain k. In this example, assume that the agents are controlled via (6.6) with $K = kI_\nu$. For the given graph, numerical simulations indicate that a choice of $k = 2.65$ yields a settling time of $t_s \approx 1.433[s]$. Using (6.11), when driven by additive white noise this design results in Wiener process with a linearly increasing drift with a slope of $k/\nu = 0.53$. Moreover, by (6.12) the steady-state variance of the disagreements equals $\|T_n\|_2^2 = 0.9858$.

To illustrate the potential of the 2DOF scheme, we shall design controllers T_r and R to ensure consensus with the same nominal performance and improved noise sensitivity.

To this end, consider a third order network filter

$$T_{\mathrm{r}}(s) = \frac{\omega_n^2}{(\tau s + 1)(s^2 + 2\zeta\omega_n s + \omega_n^2)},$$

resulting in noise response

$$\hat{S}_i(s)T_{\mathrm{r}}(s) = \frac{\omega_n^2}{\tau s^3 + (2\zeta\omega_n\tau + 1)s^2 + (\tau\omega_n^2 + 2\zeta\omega_n)s + \omega_n^2(1-\alpha_i)}.$$

Minimizing the slope of the Wiener process' drift amounts to minimizing the squared H_2 norm of

$$s\hat{S}_1(s)T_{\mathrm{r}}(s) = \frac{\omega_n^2}{\tau s^2 + (2\zeta\omega_n\tau + 1)s + (\tau\omega_n^2 + 2\zeta\omega_n)}$$

which is a simple second order system with a canonical companion realization

$$s\hat{S}_1(s)T_{\mathrm{r}}(s) = \left[\begin{array}{cc|c} 0 & 1 & 0 \\ -\frac{\tau\omega_n^2 + 2\zeta\omega_n}{\tau} & -\frac{2\zeta\omega_n\tau + 1}{\tau} & 1 \\ \hline \frac{\omega_n^2}{\tau} & 0 & 0 \end{array}\right].$$

For this simple structure we can analytically calculate the squared H_2 norm by solving a Lyapunov equation, resulting in

$$\left\| s\hat{S}_1(s)T_{\mathrm{r}}(s) \right\|_2^2 = \frac{\omega_n^3}{(2\omega_n\tau + 2\zeta)(2\omega_n\tau\zeta + 1)}. \tag{6.18}$$

From Example 6.2.3 we know that this third order filter will ensure consensus for any undirected and connected graph if

$$\zeta\omega_n > \tau,$$

and when combined with (6.18) we have a non-linear minimization problem. Note that a heuristic minimization strategy would be to keep ω_n small and τ large, while selecting ζ to enforce the stability constraint. This, however, could lead to slow poles and a dominant zero at $1/\tau$ in $\hat{S}_i(s)$ which would impact the nominal convergence rate. After some trial and error with different bounds on the parameters, we obtained

$$\begin{cases} \omega_n & = 3 \\ \tau & = 5 \\ \zeta & = 2 \end{cases} \quad \Longrightarrow \quad T_{\mathrm{r}}(s) = \frac{9}{(5s+1)(s^2 + 12s + 9)}. \tag{6.19}$$

For this T_{r} we can calculate the disagreement norm via (6.17) as well as estimate the slope of the variance via (6.18), resulting in

$$\left\|\hat{T}_n\right\|_2^2 = 0.1605 \quad \text{and} \quad \|s\hat{S}_1(s)T_{\mathrm{r}}(s)\|_2^2 = \frac{27}{34 \cdot 61} \approx 0.0116.$$

After some trial and error, a uniform local controller

$$R_0(s) = -\frac{7.586s + 16}{s + 0.4143}$$

achieves a nominal settling time of $t_s \approx 1.432[s]$ which is comparable to the standard protocol.

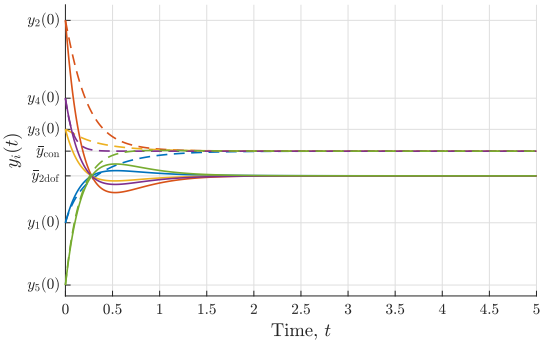

(a) Evolution of the outputs under nominal conditions.

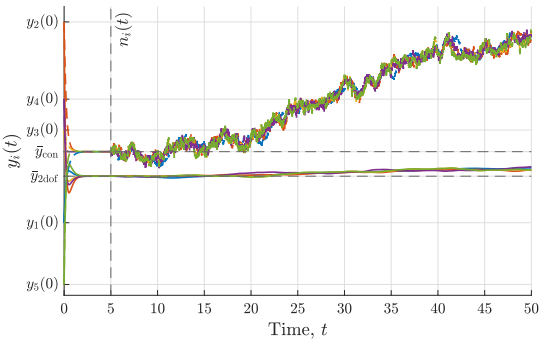

(b) Evolution of the outputs w/ additive white measurement noise.

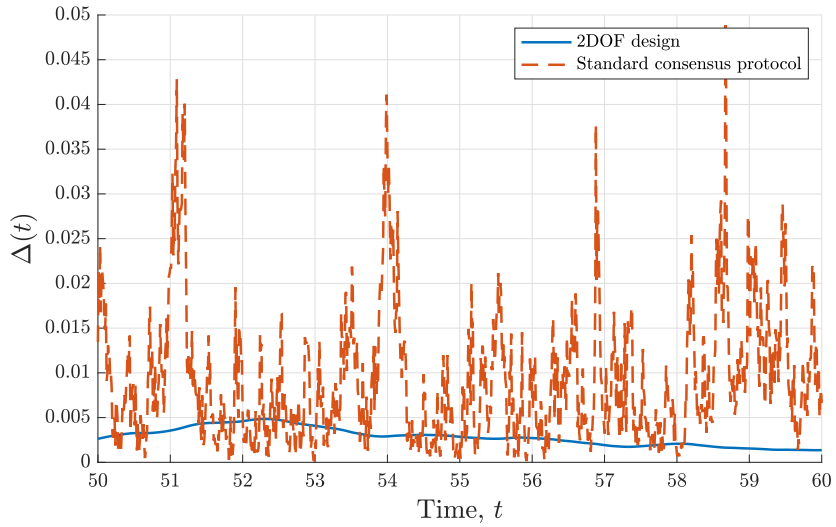

(c) The squared disagreement norm, $\Delta(t)$, at steady-state for both designs.

Figure 6.4: Simulations of the control designs for the example in Subsection 6.3.1 (solid: 2DOF design, dashed: standard consensus protocol).

Figure 6.4(a) shows the nominal behavior of both designs, which indeed have comparable settling time. Note that the designs converge to different consensus points, classical consensus to the average of initial conditions and the 2DOF to some weighted average which also depends on T_r and S. Figure 6.4(b) shows the same setups, now with white noise with intensity $|\mathcal{N}_i|$ applied at $t = 5[s]$. Since both designs have a pole at the origin for $\eta_1' y$, both behave as a Wiener process with linearly diverging variance. However, the 2DOF design has noticeably smaller drift compared to the system controlled by classic consensus. In fact, denoting the nominal consensus values by $\bar{y}_{2dof}$

and $\bar{y}_{\text{con}}$, after 60 seconds we have errors of

$$\|y_{\text{2dof}}(60) - \bar{y}_{\text{2dof}}\mathbb{1}\|_2 = 0.64 \quad \text{and} \quad \|y_{\text{con}}(60) - \bar{y}_{\text{con}}\mathbb{1}\|_2 = 9.1698,$$

respectively. The improved performance also extends to the disagreements as predicted by comparing the H_2 norm of T_n and $\hat{T}_n$. This is illustrated in Figure 6.4(c), where $\Delta(t) = \sum_{i=2}^{\nu}(\eta_i' y)^2$ is shown in steady state for both designs.

6.3.2 Consensus with disturbance rejection

Consider now again Example 5.1.1, which served as a motivating example for the instability discussed in Chapter 5. Namely, assume that one agent in the group is perturbed by a step disturbance at $t_d = 5[s]$. Since this common instability is necessary for the agents to reach agreement, the 2DOF framework is not internally stable either. Despite this, as in the previous example, we can attenuate the divergence rate of the output by shaping T_{r}. For example, Figure 6.5 compares the designs discussed in the previous example with a delayed step applied to the first agent. Indeed the output trajectories of both designs diverge linearly in response to the step disturbance, but the 2DOF design does so significantly slower.

Still, we can obtain even better results. Contrary to the response to noise, the disturbance response does not depend strictly on T_{r} but also on T_{d}. Following Example 6.2.2, we know by Proposition 6.2.1 that if there is at least a single "safe" agent, the 2DOF architecture can reject disturbances. This requires the agents to simply design local controller to reject the particular disturbance using the celebrated internal model principle [109]. For the particular case of a step disturbance, any PI controller would ensure perfect rejection of the step disturbance. Hence, a 2DOF protocol with an appropriate network filter T_{r}, local PI controllers R_i for $i = 1, \ldots, \nu - 1$, and any stabilizing controller R_ν, would still achieve asymptotic consensus.

To illustrate this, consider once more classic consensus with $k = 2.65$ and the 2DOF protocol with network filter (6.19), and assume that the fifth agent is not affected by DC disturbances. Following the logic outlined in Example 6.2.2 we design the following local controllers

$$R_i(s) = -\frac{4.74s + 8.777}{s}, \quad i \in \mathbb{N}_4, \quad \text{and} \quad R_5(s) = -\frac{7.586s + 16}{s + 0.4143}.$$

Since the first four agents have local PI controllers, they will asymptotically reject step disturbances. Despite this, we know from Proposition 6.2.1 that $\hat{S}T_{\text{d}}$ would still have a pole at the origin in the agreement direction, as required for consensus. Consequently, the agents would converge to consensus, but as long as agent 5 is safe, the outputs would not diverge. This is illustrated in Figure 6.6 where agent 1 suffers from a step disturbances at $t_d = 5[s]$. This simple way to ensure disturbance rejection is in stark contrast to the unavoidable fragility of diffusive coupling discussed in Chapter 5.

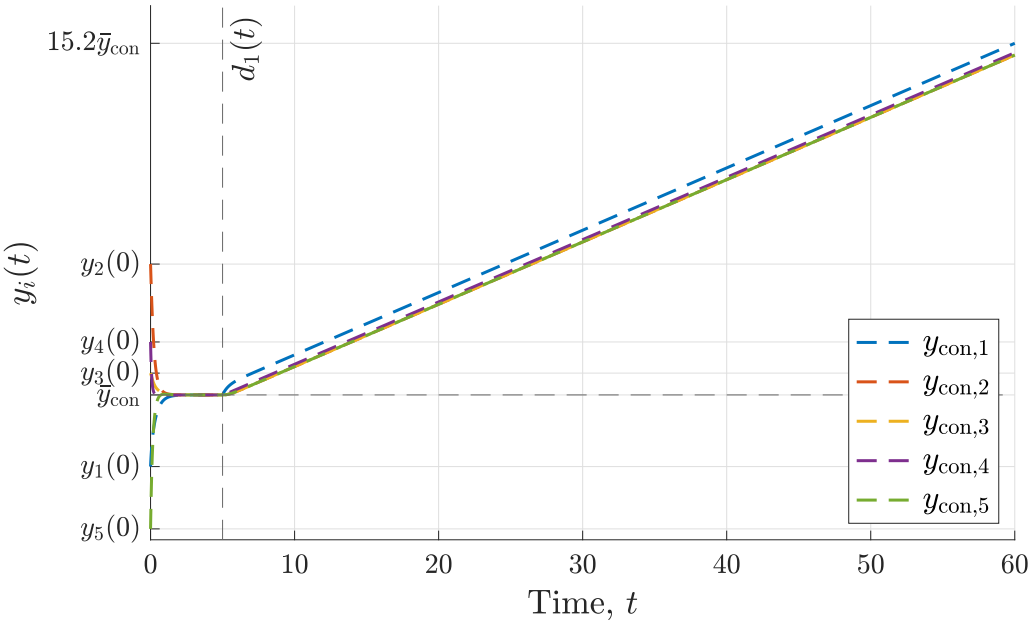

(a) Output trajectories of standard consensus protocol.

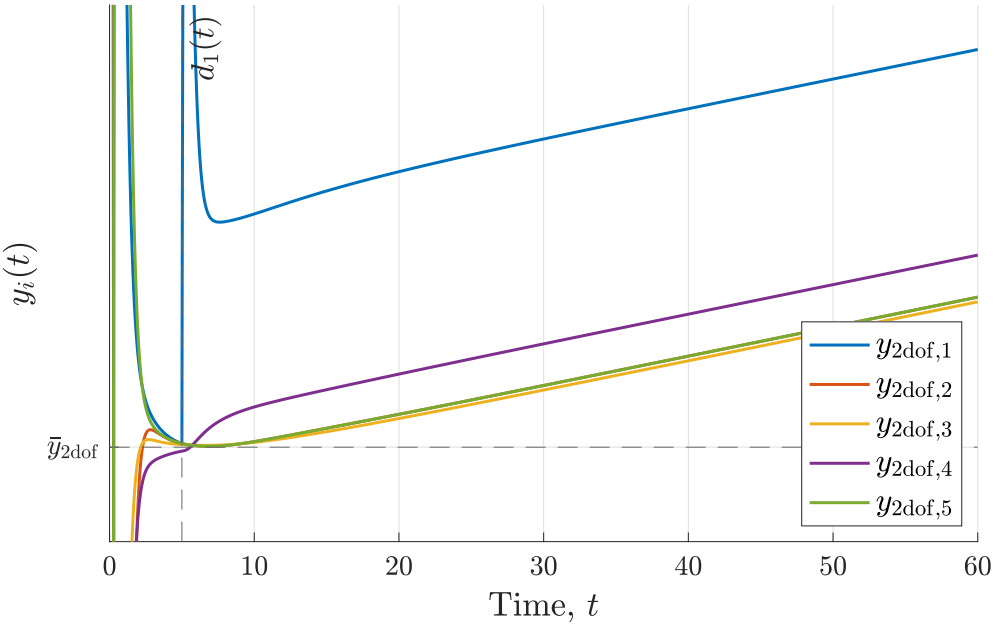

(b) Output trajectories of the 2DOF design.

Figure 6.5: The output trajectories of controllers from the example in Subsection 6.3.1 for a step disturbance applied to agent 1 at $t = 5$.

6.4 Concluding remarks

The 2DOF architecture developed here is an intriguing and novel alternative to consensus-like protocols. The separation between the local loop and network filter is a powerful tool, which allowed us to treat heterogeneous agents using similar tools to those employed in homogeneous 1DOF consensus protocols. Combined with the clean sepa-

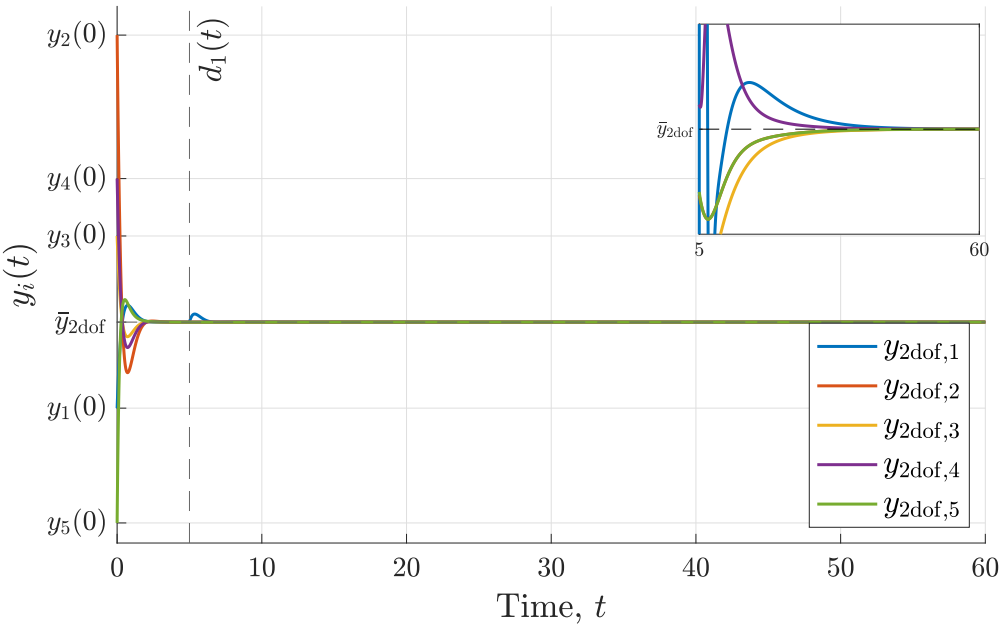

Figure 6.6: The output trajectories of the 2DOF design with PI controller from the example in Subsection 6.3.2 for a step disturbance applied to agent 1 at $t = 5$.

ration between local dynamics and the network noise, the architecture allows for "off the shelf" design of network filters. Such filters can be designed a-priori to achieve some prescribed noise attenuation: explicitly in the agreement direction and implicitly for the disagreements. In addition, controller heterogeneity can be exploited to reject local disturbances - even those exciting unstable agreement poles. As shown in Chapter 5, this is strictly impossible under standard diffusive coupling. Combined with the parallels to classical servo problems, 2DOF consensus protocols seem like a promising alternative to their classical counterpart.

Chapter 7

Conclusions and Future Directions

> Learning does not make one learned: there are those who have knowledge and those who have understanding. The first requires memory and the second philosophy.
>
> Alexandre Dumas, *The Count of Monte Cristo* [110]

In this chapter we conclude the research of this thesis. We begin with a brief summary and concluding remarks, and follow with suggestions for future research directions.

7.1 Summary and conclusions.

The research in this thesis was motivated by the lack of distributed sampled-data controllers incorporating tools from modern optimal sampled-data control such as generalized D/A (hold) and A/D (sampler) devices. The obvious reason is that such tools tend to naturally incorporate local emulation of the entire system [69], [74], [111], thus scale with the dimension of the system. In multi-agent systems scalability is paramount since the dimension of the overall system increases with the number of agents, reducing the attractivity of such tools. Hence, we asked the simple yet practical question: is there an interplay between spatial (distributed) and temporal (sampling) constraints in multi-agent systems that can be exploited.

In an attempt to answer this question, we considered the simplest of multi-agent system control problems, consensus of integrator agents, with the constraint that the communication is sampled. This resulted in the first insight to be exploited, hidden in the word *communication*. Unlike lumped systems, multi-agent systems can naturally have at least two time-scales (or sampling rates) – one for communicated and one for locally measured information. A second insight into how the spatial and temporal constraints interplay is that asynchronous sampling is equivalent to a switching

spatial topology, as discussed in Example 2.1.3. We then combined these two insight with an additional assumption, motivated by practical considerations through Moore's law [112]: computing power tend to becomes exponentially smaller, cheaper, and more powerful over time. Hence, assuming that each agent has significant computing power locally available can be reasonable. Combined, these three principles led us to down the road of the emulation architecture in Chapter 2, which employs fast local measurements and complex local emulators. While deriving the main results, we inadvertently discovered another fundamental insight, as detailed in Subsection 2.3.1. In short, to reach agreement one does not actually require the states (or outputs) of the entire group, but rather only their *centroid*. Since the centroid always has the same dimension as a single agent, emulating it directly renders the controllers invariant to the number of agents and thus scalable. These ideas were then further exploited in the various generalizations described in Chapter 3 and Chapter 4, where scalable, distributed sampled-data controllers were designed for increasingly complicated agreement problems with intermittent and asynchronous sampling.

While the first part of this thesis is clearly motivated by the original research question, the second part diverges slightly. An important property exploited within the emulation architecture is the importance of the centroid and its orthogonality to the disagreement. This is an intrinsic property of the diffusive structure, regardless of sampling. Interestingly, the ideal analog closed-loop derived in Subsection 3.2.1 shows that in general the centroid is unstable, regardless of the feedback gain used on the disagreements. This hinted that there may be some hidden instability even in classical MAS controllers, a notion supported by anecdotal evidence. Since problems in the emulated analog loop will naturally propagate to the sampled-data version, this prompted the investigation of general analog diffusive coupling in Chapter 5. The system-theoretic investigation in Chapter 5 revealed a fundamental issue common to all diffusive controllers: it cannot internally stabilize agents with a common unstable pole and direction. This in particular covers all variations of agreement with homogenous agents, since such a pole is necessary for agreement [44]. Moreover, this instability manifests in complete lack of feedback in the agreement direction, just like the centroid dynamics in the ideal case. Consequently, even if this instability is disregarded, diffusive controllers cannot attenuate the effects of disturbances and measurement noise on the agreement trajectory. Robustness to such signals is the bread and butter of classical control, making diffusive coupling an unsuitable architecture in practical applications.

Over the years diffusive coupling and the consensus protocol became almost synonymous with solving $\boldsymbol{\mathcal{P}_1}$, these architectures are usually assumed a priori when variations of $\boldsymbol{\mathcal{P}_1}$ are considered. The instability results of Chapter 5 imply that it may be advantageous to sever this tie, and reexamine the consensus protocol and $\boldsymbol{\mathcal{P}_1}$ independently. This preliminary research is the focus of Chapter 6. By conducting an input-output analysis of the consensus protocol, we were able to both highlight its shortcomings and isolate its strengths. Namely, the lack of feedback in the direction of $\mathbb{1}$ naturally

reduces agreement to a simple stabilization problem assuming that the agents P_i share an internal model of the agreement trajectory. However, it also makes it impossible to change the agreement trajectory or attenuate the effects of external inputs in this direction. Furthermore, as shown in (6.3′), locally the consensus protocol is equivalent to a proportional controller acting on the error between the local measurement and neighboring centroid. With external inputs, this is essentially a servo-regulation problem with the reference generated internally. This parallel to classical servo problems motivated the two degrees-of-freedom approach derived in Section 6.2 as an alternative to consensus. We have then shown that not only does this architecture retain the feedback path in the $\mathbb{1}$ direction, it almost completely decouples the local loops. Unlike standard consensus, reaching agreement is no longer a simple stabilization problem, but also requires some interpolation constraints to be satisfied. This complexity is offset by the ability to attenuate network noise and even completely reject certain disturbances as shown in Section 6.3, hinting that that this architecture may have significant upside.

7.2 Future research directions

A reoccurring theme in this research, is that small modifications can have significant ramifications. Consequently, there are numerous additional modifications and extensions that can be made to the difference results in this thesis. Some of these proposed directions are outlined in the three subsections below, each dedicated to one of the three main components of this work: i) the emulation scheme, ii) internal stability, and iii) the 2DOF structure.

7.2.1 Extending the emulation scheme

Heterogeneous agents. In Chapter 3 we considered homogeneous agents under either full-state or output only measurements. Despite these changes, the underlying structure was the same as it was for simple integrators: track the emulated centroid and update via discrete consensus-like update scheme. It is well-known that a necessary and sufficient condition for synchronization in continuous time is the existence of a common internal model which the agents must track [44], [46]. This result fits in nicely with the emulation scheme, where locally the agents had to stabilize the disagreements and just emulate the required trajectory. It seems reasonable that given agents $(A_i, B_i, C_i, 0)$, controllers inspired by [44, Eq. 10] such as

$$\begin{cases} \dot{\hat{x}}_i(t) = A_i\hat{x}_i(t) + B_iu_i(t) + L_i(C\hat{x}_i(t) - y_i(t)) \\ \dot{\bar{\mu}}_i(t) = A_0\bar{\mu}_i(t) \\ \bar{\mu}_i(s_k^+) = \bar{\mu}_i(s_k) - \dfrac{1}{\nu}\displaystyle\sum_{l\in\mathcal{N}_i[k]}(\bar{\mu}_i(s_k) - \bar{\mu}_l(s_k)) \\ u_i(t) = K_i(\hat{x}_i(t) - \Pi_i\bar{\mu}_i) + \Gamma_i\bar{\mu}_i(t) \end{cases},$$

where K_i and L_i are stabilizing feedback and observer gains, and Π_i satisfy the regulator equations Γ_i

$$A_i\Pi_i + B_i\Gamma_i = \Pi_i A_0$$
$$C_i\Pi_i = I$$

will solve $\boldsymbol{\mathcal{P}_2}$. There are technical difficulties, of course, due to the heterogeneity of the agents, but not a conceptual one. Note that replacing the above with $C_i\Pi_i = R$ for some common R would change $\boldsymbol{\mathcal{P}_2}$ from state to putpur synchronization.

Modifying the update map. In Chapter 4, we have seen that the convergence rate of the discrete dynamics lower bound the convergence rate of the overall system. In addition, since the update map operates a consensus-like protocol, it is sensitive to measurement noise as described in Chapter 6. The fact the discrete dynamics follow those of first-order consensus regardless of the actual dynamics can be exploited to improve the convergence rate and reduce the noise sensitivity. Following Proposition 4.3.1, we know it is possible, under certain conditions, to tune edge weights without changing the validity of the results. These weights can also be time-varying as long as they satisfy the aforementioned conditions. This implies that we can directly implement known results about optimal edge weights to the existing scheme, cf. [40], [83], [85]–[88], assuming the conditions of these works hold.

Additional analog dynamics. Our underlying assumption has been that the agents can continuously measure local information. An interesting and non-trivial extension is to assume that the agents not only have continuous local information, but continuous information of some *subgraph*. This could represent, for example, situations where the agents start in some clusters which are sufficiently close to allow communication. This is not a trivial extension, since it is not obvious what the analog part should be doing. On the one hand, trying to use it to stabilize the agents would require designing structured stabilizing controllers which are known to be NP-hard [9]. On the other hand, using it to keep clusters "close" can result in multiconsensus [66], resulting in two different agreement problems - one in the intersample and one in on sampling instances.

Nonlinear agents and controllers. Although conceptually similar, i.e. locally stabilizing and discretely agreeing, this poses significant technical hurdles. It was conjectured in discussions with Prof. WPMH Heemels and Prof. Erik Steur, that a sufficient condition on the local loops would be strict incremental stability. Alas, this direction has not been pursued thus far.

Applications. Consensus and synchronization are important not only for themselves, but as building blocks for more realistic objectives. Two interesting applications are rigidity-based formation control and distributed algorithms. In rigidity-based formation control the controllers used involve objects called rigidity matrices, which can be

thought of as Laplacians with state-dependent nonlinear weights [113]–[115]. It would be interesting to try and adapt the emulation scheme for formation control, or at least for its consensus part. The second application is more straightforward. Consensus-like steps are a well-known feature of distributed algorithms [17]–[19], [116], hence implementing the emulation scheme seems like a natural step. Possible hurdles include accommodating for additional dynamics, possible nonlinearity, and modifications for discrete-time dynamics.

7.2.2 Doubling down on internal stability

Nonlinear controllers. It is a non-trivial question whether or not nonlinear controllers can stabilize diffusively coupled systems. Clearly the methods used in proving Theorem 5.1 relied heavily on the linear nature of the controllers through the coprime factorization, which does not easily translate to the nonlinear case. However, it seems intuitive that similar results would hold even in the nonlinear case. Even in the nonlinear case, the controller can be represented as

$$K \coloneqq (E \otimes I_m) K_{\mathrm{e}} (E^\top \otimes I_p).$$

for some nonlinear operator K_{e}. Controllers are diffusive if and only if they follow the structure from Section 5.2, namely the process relative outputs and sum the along the edges. The pre and post multiplication by the incidence matrix constraint the measurement to be relative, and output to sum over the edges. Hence, the deficient rank would still be there even for nonlinear edge controllers. Following the discussion in Chapter 6, this still implies that there are certain inputs which are in a sense uncontrollable by the diffusive structure. Proofs might be more technically involved, but the results seem plausible. One possible way to analyze this system is by considering some nonlinear projection onto the agreement and disagreement spaces, and analyzing whether the feedback path disappears in the agreement space.

Nonlinear agents. As before, it is reasonable to believe that the instability result would hold even for nonlinear agents. Some evidence in this direction is due to Lyapunov's indirect method, i.e., analyzing the linearized model. Since for any linearized model the system satisfies the assumptions of Theorem 5.1, it will be locally unstable around every point.

7.2.3 Maturing the two-degrees-of-freedom protocol

Conditions for graph independence. A staple of diffusive coupling is that if certain local conditions are satisfied, the controllers will drive the agents to agreement for any connected graph. This is still a missing piece in the 2DOF protocol, as Theorem 6.1 requires a robust stability condition with respect to the eigenvalues of $A^\star_{\mathcal{G}}$.

Example 6.2.3 and Example 6.2.4 showed certain families of network filters that can ensure consensus or simple harmonic synchronization for arbitrary graphs, but the architecture still lacks more general results. Preliminary results based on the small-gain theorem and the generalized Nyquist criterion hint that such conditions may exist for SISO agents. Unfortunately, the results were not finalized in time and are therefore omitted from this thesis.

Dealing with uncertainty. Here there are two types of uncertainties: i) local uncertainties, and ii) dynamic uncertainties on the edges. The difficulties with the first kind are due to the component C_{OL} in Figure 6.2, which acts as

$$C_{\mathrm{OL}} = P^{-1}T_{\mathrm{r}}.$$

If P is uncertain, then the plant inversion is not perfect and (6.15a) does not hold. Since the "reference" signal in this case is a filtered version of the real output, the resulting dynamics are harder to analyze and predict. In contrast, the second type of uncertainty might be easier to analyze. This type would affect $A^{\star}_{\mathcal{G}}$, turning it into a dynamics variant $A^{\star}_{\mathcal{G}}(s)$. For some interesting cases, such as when each edge of $A^{\star}_{\mathcal{G}}(s)$ is afflicted by a different delay, certain important properties are unaltered. For example, note that in the delay case $A^{\star}_{\mathcal{G}}(0) = A^{\star}_{\mathcal{G}}$. Since static consensus is essentially a DC objective, it implies that it is still achievable. This is supported both by simulations and preliminary theoretical results saying that if certain small-gain conditions hold for T_{r}, then consensus can be achieved for arbitrary edge delays.

Designing good 2DOF agreement controllers. As established in Theorem 6.1, reaching agreement requires that $\hat{S}_1$ has prescribed unstable poles, and that $\hat{S}_i$ are stable for all $i > 1$. One possible approach to solving this problem is by first parameterizing all filters T_{r} ensuring the prescribed poles, and then solving a stabilization problem with uncertain gain. The resulting problem, however, is not a trivial one for two reasons. First, potentially there can be some i such that $\alpha_i = 0$, hence $\hat{S}_i = 1$. Assuming nonzero noise, this implies that T_{r} itself must be stable, rendering this a strong stabilization problem. Second, strong stabilization problems with uncertain gain has been solved in the past, cf. [117], but under the assumption that the gain is strictly positive, which is not the case here. Consequently, even the mere stabilization problem cannot be generally solved using standard tools.

Moreover, T_{r} and α_i completely determine the noise response, thus ideally we would like to solve a robust H_2 problem for $\hat{S}_i$ and not simply stabilize it. If the graph is perfectly known, then this might be slightly simpler as (6.17), for example, could be designed directly. Note that even directly minimizing (6.17) would only minimize the variance of the disagreements, and that the agreement mode must be treated separately. Since $\hat{S}_1$ is by definition unstable, regular optimal control methods are also not directly

applicable to design T_{r}.

Revisiting classical tools. The 2DOF protocol could serve as a catalyst for a new variation in classical control - designing "good" controllers while enforcing certain unstable closed-loop poles. This is somewhat of a dual problem to classic internal model principle [109], where the goal was to place specific blocking *zeros* in the closed-loop. Such problem would have been considered nonsense in classical settings, but arises naturally in agreement problems. For example, the Youla-Kucera parametrization is a powerful tool that parametrizes all internally stabilizing controllers for a given system. It is then commonly used in designing both optimal [104, Ch. 7] and robust [118, Ch. 6 – 7] controllers. As established, agreement problems require specific unstable closed-loop poles. This precludes internal stability and by extension the use of the Youla-Kucera parametrization. However, if we could derive a similar parametrization that relaxes the internal stability requirement, it can bridge the gap between classical H_∞ tools and agreement protocols. In particular we would like to be able to enforce only prescribed unstable closed-loop poles. Assuming this can be done, it could also be used to derive a parametrization of all possible 2DOF protocols using coprime factorizations as in the servo case [55], [119].

A 2DOF sampled-data emulation approach. Finally, an obvious direction would be to revisit the emulation scheme from Chapter 3 and apply it in conjunction with the 2DOF approach. This idea has two immediate variations. One, is to directly emulate the 2DOF scheme instead of the heuristic diffusive controllers from Chapter 3. The second variation is to implement the scheme as it is, locally emulate $\tilde{y}$, and update it at sampling instances. This potentially would not increase the overall dimension of the controller, and still keep only the communicated information as sampled-data.

Appendix A

Introduction to Graph Theory

Graphs are widely studied mathematical objects [23], [120] and are one of the main tools in modeling multi-agent systems [22]. Therefore, it is important to introduce some common terminology, notation, and results.

A mathematical graph $\mathcal{G}$ is a pair of sets $(\mathcal{V}, \mathcal{E})$ where $\mathcal{V}$ is the set of vertices, and $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ is the set of connecting edges. Given an edge (v_i, v_j), the node v_i is called an *in-neighbor* of v_j, while node v_j is called an *out-neighbor* of v_i. The in (out) *degree* of a vertex, denoted $\deg_i(v_i)$ (resp. $\deg_o(v_i)$) is the number of in (out) neighbors of v_i. The *in-neighborhood* of vertex j is defined as all vertices v_i such that

$$v_i \in \mathcal{N}_j^{\text{in}} \iff (v_i, v_j) \in \mathcal{E}.$$

Similarly, the *out-neighborhood* of vertex i is defined as all vertices v_j such that

$$v_j \in \mathcal{N}_i^{\text{out}} \iff (v_i, v_j) \in \mathcal{E}.$$

A graph is called *undirected* or symmetric if $(v_i, v_j) \in \mathcal{E} \implies (v_j, v_i) \in \mathcal{E}$, and *directed* (digraph) if the ordering of the vertex pair matter. If the graph is undirected then $\mathcal{N}_i^{\text{out}} = \mathcal{N}_i^{\text{in}} = \mathcal{N}_i$ for all i. A graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}')$ is called a *subgraph* of $\mathcal{G}$ if $\mathcal{V}' \subseteq \mathcal{V}$ and $\mathcal{E}' \subseteq \mathcal{E}$. A subgraph is called *spanning* if $\mathcal{V}' = \mathcal{V}$.

A *directed path* is an ordered sequence of nodes such that any pair of consecutive nodes in the sequence is a directed edge of the graph. A *cycle* is a directed path that starts and ends at the same node, a graph with no cycles is called *acyclic*. A *directed tree* (sometimes rooted tree) is an acyclic graph with the property that there exists a *root* node such that there is one and only one directed path from the root to every other node. Given a graph $\mathcal{G}$ and subgraph $\mathcal{G}'$, $\mathcal{G}'$ is called a *directed spanning tree* of $\mathcal{G}$ if it is a spanning subgraph that is a directed tree. An undirected graph is called *connected* if there is some sequence of adjacent edges between every such pair of nodes. A directed graph is said to be *weakly connected* if its undirected version is connected. A graph is said to be *strongly connected* if there exists a directed path from any node to any other node. An undirected graph is said to be *complete* if every pair of distinct

vertices is connected by an edge. These concepts are illustrated in Figure A.1.

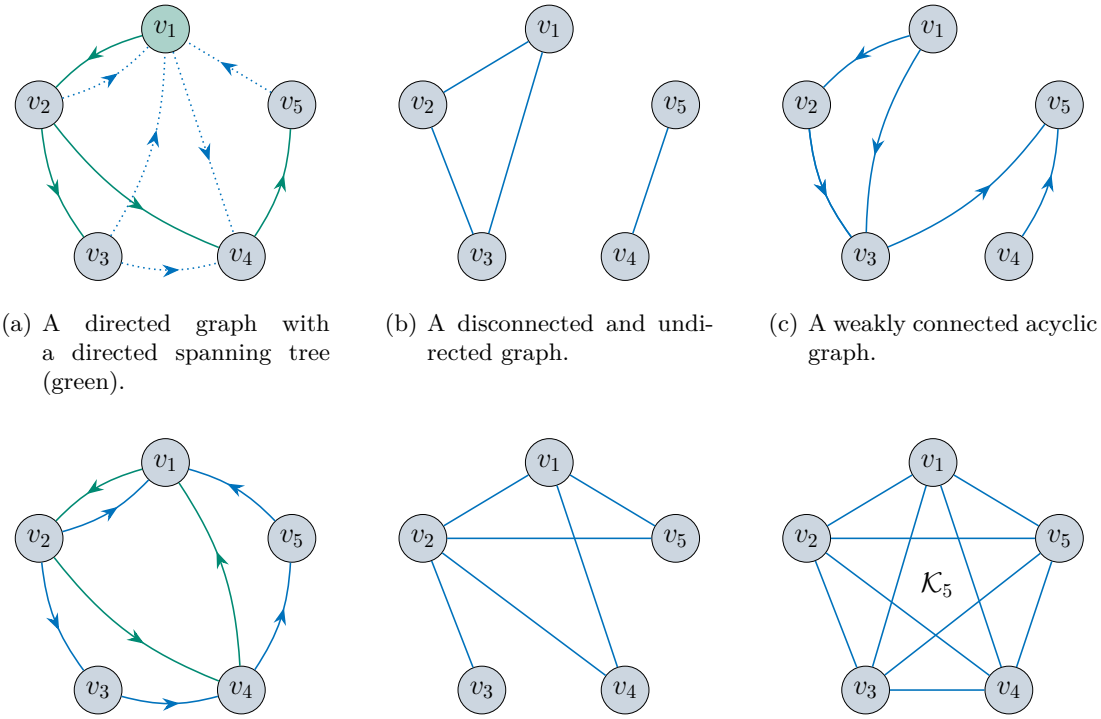

(a) A directed graph with a directed spanning tree (green).

(b) A disconnected and undirected graph.

(c) A weakly connected acyclic graph.

(d) A cyclic (one cycle in green) strongly connected graph.

(e) A connected undirected graph.

(f) The complete graph on 5 vertices.

Figure A.1: Illustrations of basic graph concepts.

We shall now define several important matrix representations of graph $\mathcal{G}$. The *adjacency matrix* of a graph $\mathcal{G}$ is denoted as $A_{\mathcal{G}}$ and defined as:

$$[A_{\mathcal{G}}]_{ij} = \begin{cases} 1, & \text{if } (v_j, v_i) \in \mathcal{E} \\ 0, & \text{else} \end{cases}, \tag{A.1}$$

which is a square binary matrix, and it is symmetric if $\mathcal{G}$ is undirected. Note that $[A_{\mathcal{G}}]_{ij}$ depends on (v_j, v_i); this is consistent with [21], but [72] defines the transpose.

The (oriented) incidence matrix of $\mathcal{G}$ is denoted by $E_{\mathcal{G}}$ or simply E when the association with a concrete graph is clear. It is a $|\mathcal{V}| \times |\mathcal{E}|$ matrix, whose (i, j) entry is

$$[E_{\mathcal{G}}]_{ij} = \begin{cases} 1 & \text{if vertex } i \text{ is the head of edge } j \\ -1 & \text{if vertex } i \text{ is the tail of edge } j \\ 0 & \text{if vertex } i \text{ does not belong to edge } j \end{cases}. \tag{A.2}$$

Note that the construction of the incidence matrix implies that $\mathbb{1}' E_{\mathcal{G}} = 0$ for every $\mathcal{G}$.

The *degree matrix* of a graph is a diagonal matrix denoted as $D_{\mathcal{G}}$ and defined as:

$$[D_{\mathcal{G}}]_{ij} = \begin{cases} \deg(v_i), & \text{if } i = j \\ 0, & \text{else} \end{cases}. \tag{A.3}$$

When required we shall use the notation $D^i_{\mathcal{G}}$ and $D^o_{\mathcal{G}}$ to differentiate between the in and out degree matrices, and for undirected graphs $D^i_{\mathcal{G}} = D^o_{\mathcal{G}} = D_{\mathcal{G}}$. The degree matrices can be equivalently defined via the Adjacency matrix through

$$D^i_{\mathcal{G}} = \operatorname{diag}\{A_{\mathcal{G}}\mathbb{1}\}, \quad D^o_{\mathcal{G}} = \operatorname{diag}\{\mathbb{1}'A_{\mathcal{G}}\}.$$

Assuming that $D_{\mathcal{G}}$ is invertible (i.e., no isolated nodes), we can also define the *normalized adjacency matrix* $A^\star_{\mathcal{G}} \coloneqq D^{-1}_{\mathcal{G}} A_{\mathcal{G}}$.

The *Laplacian matrix* of an undirected graph $\mathcal{G}$ is denoted as $L_{\mathcal{G}}$ and defined as:

$$[L_{\mathcal{G}}]_{ij} = \begin{cases} \deg(v_i), & \text{if } i = j \\ -1, & \text{if } i \neq j \text{ and } (v_i, v_j) \in \mathcal{E} \\ 0, & \text{else} \end{cases} \implies \quad L_{\mathcal{G}} = D_{\mathcal{G}} - A_{\mathcal{G}}. \tag{A.4}$$

We can similarly define the in-degree and out-degree Laplacians of a directed graph as $L^i_{\mathcal{G}} = D^i_{\mathcal{G}} - A_{\mathcal{G}}$ and $L^o_{\mathcal{G}} = D^o_{\mathcal{G}} - A_{\mathcal{G}}$ respectively.

These matrices, and the Laplacian in particular, play a pivotal role in the control of multi-agent systems and agreement problems. The following proposition lists several known properties for undirected graphs that will be used throughout this thesis.

Proposition A.0.1. *Let $\mathcal{G}$ be an undirected graph.*

1. *[23, Lemma. 13.1.1] The Laplacian, $L_{\mathcal{G}}$, is symmetric and positive semi-definite thus its real spectrum can be ordered as*

$$0 = \lambda_1(L_{\mathcal{G}}) \leq \lambda_2(L_{\mathcal{G}}) \leq \ldots \leq \lambda_{|\mathcal{V}|}(L_{\mathcal{G}}).$$

2. *[23, Lemma. 13.1.1] The multiplicity of the zero eigenvalue of the graph Laplacian is equal to the number of connected components of the graph.*

3. *[23, Thm. 8.3.2] The Laplacian can be written as*

$$L_{\mathcal{G}} = EE'$$

where E is the incidence matrix with any arbitrary orientation.

4. *For any graph $\mathcal{G}$, $\operatorname{span}\{\mathbb{1}\} \subseteq \ker L_{\mathcal{G}}$. Furthermore, $\operatorname{span}\{\mathbb{1}\} = \ker L_{\mathcal{G}}$ if and only if $\mathcal{G}$ is connected.*

5. *[12, Prop. 1–4] If the graph is connected, then $A^\star_\mathcal{G}$ is well defined, its eigenvalues are real and satisfy*

$$\operatorname{spec} A^\star_\mathcal{G} \in [-1, 1].$$

Moreover, it always has a unique simple eigenvalue at 1 with corresponding eigenvectors

$$A^\star_\mathcal{G}\mathbb{1} = \mathbb{1} \quad \textit{and} \quad (D_\mathcal{G}\mathbb{1})' A^\star_\mathcal{G} = (D_\mathcal{G}\mathbb{1})'.$$

Proof We only have to prove item 4, which is an immediate consequence of items 2 and 3. ∎

From the above we may note that the second smallest eigenvalue of $L_\mathcal{G}$ is non-zero if and only if the graph is connected. Hence, it is used as a measure of the connectivity of the graph, and called the *algebraic connectivity* [121].

For directed graphs, we can similarly define the following proposition.

Proposition A.0.2. *Let $\mathcal{G}$ be a directed graph with in-degree and out-degree Laplacian matrices $L^i_\mathcal{G}$ and $L^o_\mathcal{G}$ respectively.*

1. *For any graph $\operatorname{span}\{\mathbb{1}\} \subseteq \ker L^i_\mathcal{G}$ and $\operatorname{span}\{\mathbb{1}\} \subseteq \ker(L^o_\mathcal{G})'$. Equivalently, the in and out directed Laplacians have zero row and column sums, respectively.*
2. *[22, Prop. 3.10] The non-zero eigenvalues of $L^i_\mathcal{G}$ and $L^o_\mathcal{G}$ have strictly positive real parts.*
3. *[22, Prop. 3.8] $L^i_\mathcal{G}$ has a unique and simple zero eigenvalue if and only if $\mathcal{G}$ has a directed spanning tree.*
4. *[12, Prop. 1–3] If $D^i_\mathcal{G}$ (respectively, $D^o_\mathcal{G}$) is invertible, then the directed counterpart of $A^\star_\mathcal{G}$ is well defined and $\operatorname{spec} A^\star_\mathcal{G} \in \mathcal{B}_1$, where $\mathcal{B}_1$ is unit ball in the complex plane.*
5. *Given a graph with ν nodes the adjacency matrix satisfies $\left\|A_\mathcal{G}\right\|_2 \leq \nu - 1$, where $\|\cdot\|_2$ is the spectral norm, i.e. the induced Euclidean norm.*
6. *Given a graph with ν nodes, both directed Laplacians satisfy $\left\|L_\mathcal{G}\right\|_2 \leq \nu(\nu - 1)$.*
7. *If $L_\mathcal{G}$ is the in-degree Laplacian of a graph with ν nodes, then the matrix*

$$M = I - \frac{1}{\nu} L_\mathcal{G}$$

is row-stochastic (see Definition B.2.1).

Proof The first property is immediate by construction, hence we only have to prove the last three properties.

For the adjacency matrix, note that $A_\mathcal{G}$ is a binary matrix with $a_{ii} = 0$ for all k. By Gershgorin's circle theorem [122, Thm. 6.1.1], this means that all of its eigenvalues are located in a circle centered at the origin with maximal radius of $\nu - 1$ (the maximal

number of non-zero elements in each row). By [123, Thm.1.1], the spectral norm of any oriented adjacency matrix is upper bounded by the spectral radius of its unoriented counterpart. Combining the two results we have $\left\|A_{\mathcal{G}}\right\| \leq \nu - 1$ for all k.

Now recall that the maximal possible element of $L_{\mathcal{G}}$ is equal to $\nu - 1$. Through the equivalence of norms [122] we know that for any $n \times m$ matrix A

$$\|A\|_2 \leq \sqrt{nm}\, \|A\|_{\max},$$

hence for the Laplacian we obtain

$$\left\|L_{\mathcal{G}}\right\|_2 \leq \nu(\nu - 1).$$

For the last property, by definition the Laplacian has non-negative diagonal entries and either 0 on -1 on the off diagonal entries. As established, the diagonal entries are bounded by $\nu - 1$ thus

$$0 \leq [M]_{ij} \leq 1$$

for all i, j and any graph and is always non-negative. Moreover, since $L_{\mathcal{G}}\mathbb{1} = 0$, we have $M\mathbb{1} = \mathbb{1}$ and thus a row sum of 1, making it row-stochastic. ■

Appendix B

Elements of Matrix Theory and Linear Algebra

In this thesis, we make use of some notions and constructions from matrix theory which go beyond the scope of basic undergraduate courses. This appendix is meant to briefly introduce some relevant definitions and results to ensure common ground.

B.1 The Kronecker product

The definitions and results here are adapted from [65, Ch. 13].

Definition B.1.1. The Kronecker product of $A \in \mathbb{R}^{n\times m}$ and $B \in \mathbb{R}^{q\times r}$ is the $nq \times mr$ matrix

$$A \otimes B = \begin{bmatrix} a_{11}B & \cdots & a_{1m}B \\ \vdots & \ddots & \vdots \\ a_{n1}B & \ddots & a_{nm}B \end{bmatrix}.$$

Proposition B.1.2 (Properties of the Kronecker product)**.** *For appropriate matrices A, B, C, D and scalars $\alpha, \beta, \gamma, \delta$, the Kronecker product satisfies the following properties.*

Bi-linearity

$$(\alpha A + \beta B) \otimes (\gamma C + \delta D) = \alpha\gamma A \otimes C + \alpha\delta A \otimes D + \beta\gamma B \otimes C + \beta\delta B \otimes D. \tag{B.1a}$$

Associativity

$$(A \otimes B) \otimes C = A \otimes (B \otimes C) \tag{B.1b}$$

Transpose

$$(A \otimes B)' = A' \otimes B' \tag{B.1c}$$

The mixed-product

$$(A \otimes B)(C \otimes D) = (AC \otimes BD) \tag{B.1d}$$

Eigenpair

$$Av = \lambda v,\ Bw = \mu w \quad \Longrightarrow \quad (A \otimes B)(v \otimes w) = \lambda\mu\,(v \otimes w) \tag{B.1e}$$

Spectrum

$$\operatorname{spec}(A \otimes B) = \left\{\lambda\mu \;\middle|\; \lambda \in \operatorname{spec}(A),\ \mu \in \operatorname{spec}(B)\right\} \tag{B.1f}$$

Inverse *For square A, B, the product $A \otimes B$ is invertible if and only if A and B are invertible, and is given by*

$$(A \otimes B)^{-1} = A^{-1} \otimes B^{-1} \tag{B.1g}$$

B.2 Stochastic and non-negative matrices

Definition B.2.1. A square matrix $A \in \mathbb{R}^{n\times n}$ is said to be

(i) non-negative (positive) if $a_{ij} \geq 0$ for all i, j.

(ii) row-stochastic if it is non-negative and $A\mathbb{1} = \mathbb{1}$.

(iii) column-stochastic if it is non-negative and $A'\mathbb{1} = \mathbb{1}$.

(iv) doubly-stochastic if it is both row and column stochastic.

Proposition B.2.2. *Below are several properties of stochastic matrices.*

1. *Let A be a $\nu \times \nu$ row-stochastic matrix, then*

$$\|A\|_2 \leq \sqrt{\|A\|_1} \leq \sqrt{\nu}.$$

2. *The product of row stochastic matrices is again row-stochastic.*

Proof 1. The first inequality is an immediate consequence of the known inequality

$$\|A\|_2 \leq \sqrt{\|A\|_\infty \|A\|_1}$$

and the fact that the row sum of a row-stochastic matrix is always 1. The second inequality is also immediate since if each row sums to 1, the maximal possible column sum is $\sum_{i=1}^{\nu} 1 = \nu$.

2. Let A, B be two row-stochastic matrices, i.e. square, non-negative and with row-sum 1. Note that the row-sum condition is equivalent to $A\mathbb{1} = \mathbb{1}$. Consider the AB, clearly AB is square and non-negative since it is the product of square and non-negative matrices. Since B is row-stochastic $B\mathbb{1} = \mathbb{1}$ and $A\mathbb{1} = \mathbb{1}$, thus $AB\mathbb{1} = \mathbb{1}$. Thus the product AB is also row-stochastic. ∎

Appendix C

Background on Dynamical Systems

C.1 Linear time-invariant systems

A system $G : u \mapsto y$ is an operator mapping some input signal $u \in \mathbb{F}^m$ to output signal $y \in \mathbb{F}^p$. The signals u and y are often described as functions of an independent variable t which represent time, and often take values in $\mathbb{R}$. We say a system is SISO if $m = p = 1$, and MIMO otherwise.

Definition C.1.1. A system G is called *linear* if it satsifies the property of superposition:

$$G(\alpha_1 u_1 + \alpha_2 u_2) = \alpha_1 (G u_1) + \alpha_2 (G u_2), \quad \text{for all admissible signals } u_1, u_2 \text{ and scalars } \alpha_1, \alpha_2.$$

A system is called *time-invariant* if any constant time shift of its input results in the same time shift of its output.

Note that Gu represents the action of system (operator) G on signal u. For an LTI system, this action is given by the following *convolution integral*

$$y(t) = \int_{\mathbb{R}} g(t-s)u(s)ds \tag{C.1}$$

where $g(t)$ is called the *impulse response* of system G. An LTI system is said to be *causal* if and only if $g(t) = 0$ whenever $t < 0$.

Definition C.1.2 (Stability)**.** Define

$$L_2^n(\mathbb{R}) \coloneqq \left\{ x \ : \ \mathbb{R} \to \mathbb{R}^n \ \middle| \ \|x\|_2 \coloneqq \left(\int_{\mathbb{R}} \|x(t)\|^2 \, dt \right)^{1/2} < \infty \right\}.$$

A system is said to be L_2 *stable* if for any $u \in L_2$ we have $y \in L_2$.

C.1.1 LTI systems and their transfer functions

It is often useful to analyze LTI systems in transformed domains. In particular, applying the *Laplace Transform* to (C.1) simplifies the integral to

$$Y(s) = G(s)U(s), \tag{C.2}$$

where $G(s), U(s)$, and $Y(s)$ are the Laplace transforms of $g(t), u(t)$, and $y(t)$ respectively. We then call $G(s)$ the *transfer function* of G, and it is a complex function of $s = \sigma + \mathrm{j}\omega$. An LTI system is said to be *finite-dimensional* if its transfer function is real-rational.

In this domain, we can directly characterize the stability of LTI and causal systems via their transfer functions [95, §A.6.3].

Lemma C.1.3 (Stability via Transfer Functions [124])**.** *An LTI and causal system is stable if and only if its transfer function $G(s)$ is holomorphic and bounded in $\mathbb{C}_0$, i.e. iff $G \in H_\infty^{p\times m}$, where*

$$H_\infty^{p\times m} \coloneqq \left\{ G \;:\; \mathbb{C}_0 \to \mathbb{C}^{p\times m} \;\middle|\; G(s) \text{ is holomorphic in } \mathbb{C}_0 \text{ and } \|G\|_\infty \coloneqq \sup_{s\in\mathbb{C}_0} \|G(s)\| < \infty \right\} \tag{C.3}$$

and $\|G(s)\|$ is the matrix spectral norm. We write H_∞ when the dimensions are clear.

Definition C.1.4 (Properness)**.** A $p \times m$ transfer function $G(s)$ is said to be *proper* if

$$\exists \alpha \geq 0 \text{ such that } \sup_{s\in\mathbb{C}_\alpha} \|G(s)\| < \infty,$$

and *strictly proper* if there exists $\alpha \geq 0$ such that

$$\lim_{|s|\to\infty, s\in\mathbb{C}_\alpha} \|G(s)\| = 0.$$

Functions $M \in H_\infty^{m\times m}$ and $N \in H_\infty^{p\times m}$ are said to be *right coprime* if there are $X \in H_\infty^{m\times m}$ and $Y \in H_\infty^{m\times p}$ (Bézout coefficients) such that

$$XM + YN = I_m. \tag{C.4a}$$

Functions $\tilde{M} \in H_\infty^{p\times p}$ and $\tilde{N} \in H_\infty^{p\times m}$ are said to be *left coprime* if there are $\tilde{X} \in H_\infty^{p\times p}$ and $\tilde{Y} \in H_\infty^{m\times p}$ such that

$$\tilde{M}\tilde{X} + \tilde{N}\tilde{Y} = I_p. \tag{C.4b}$$

A transfer function $G(s)$ is said to have coprime factorizations over H_∞ if there are right coprime $M_G, N_G \in H_\infty$ and left coprime $\tilde{M}_G, \tilde{N}_G \in H_\infty$, known as right and left coprime factors of G, respectively, such that

$$G = N_G M_G^{-1} = \tilde{M}_G^{-1}\tilde{N}_G. \tag{C.5}$$

Coprime factors are unique up to post- or pre-multiplication by bi-stable transfer functions for right and left factors, respectively.

Lemma C.1.5. *If $G(s)$ has coprime factorizations, then*

$$G \in H_\infty \iff M_G^{-1} \in H_\infty \iff \tilde{M}_G^{-1} \in H_\infty.$$

Proof The "if" part of the first equivalence relation is immediate from (C.5). Its "only if" part follows from rewriting the Bézout equality (C.4a) as $M_G^{-1} = X_G + Y_G G$. The second relation follows by similar arguments. ∎

Lemma C.1.6. *Let $G(s)$ have coprime factorizations. If $\lambda \in \bar{\mathbb{C}}_0$ is a pole of $G(s)$, then $M_G(\lambda)$ and $\tilde{M}_G(\lambda)$ are singular.*

Proof Because $\lambda \in \bar{\mathbb{C}}_0$, the singularity of $M_G(\lambda)$ or $\tilde{M}_G(\lambda)$ does not depend on concrete factorizations taken. If $M_G(\lambda)$ is nonsingular, then $N_G(\lambda)M_G(\lambda)^{-1}$ is bounded, which implies that λ cannot be a pole of $G(s)$. The proof for $\tilde{M}_G$ is similar. ∎

A comprehensive exposition of the subject can be found in [119].

C.1.2 State-space realizations and transfer Functions

Let G be a finite-dimensional LTI system having a proper transfer function $G(s)$. The system G has a state-space realization

$$G(s) = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array}\right] := D + C(sI - A)^{-1}B. \tag{C.6}$$

The eigenvalues of A are known as *poles* of the realization (C.6). A realization is called *minimal* if and only if the set of all realization poles, multiplicities counted, coincides with that of the poles of the transfer function $G(s)$.

A matrix pair $(A, B) \in \mathbb{R}^{n\times n} \times \mathbb{R}^{n\times m}$ is called *controllable* if the eigenvalues of $A + BK$ can be freely assigned by a suitable choice of K, and called *uncontrollable* otherwise. Eigenvalues of A that cannot be freely assigned are called *uncontrollable modes*. A pair is called *stabilizable* if all of its uncontrollable modes are in $\mathbb{C} \setminus \bar{\mathbb{C}}_0$.There are numerous criteria for analyzing controllability, some common ways are presented below, and proven in [104, Thm. 4.1].

Lemma C.1.7. *The following statements are equivalent:*

1. *The pair (A, B) is controllable.*
2. *The matrix $\left[\begin{array}{cc} A - sI & B \end{array}\right]$ has full rank $\forall s \in \mathbb{C}$ (the PBH test).*
3. *The controllability matrix*

$$M_c := \left[\begin{array}{cccc} B & AB & \cdots & A^{n-1}B \end{array}\right]$$

has full rank.

Similarly, a pair $(C, A) \in \mathbb{R}^{p\times n} \times \mathbb{R}^{n\times n}$ is called *observable* if the eigenvalues of $A + LC$ can be freely assigned by a suitable choice of L. Eigenvalues of A that cannot be freely assigned are called *unobservable modes*. A pair is called *detectable* if all of its unobservable modes are in $\mathbb{C} \setminus \bar{\mathbb{C}}_0$. Like in the controllability case, there are numerous tools to analyze observability which are proven in [104, Prop. 4.6].

Lemma C.1.8. *The following statements are equivalent:*

1. *The pair (C, A) is observable.*
2. *The matrix $\begin{bmatrix} A - sI \\ C \end{bmatrix}$ has full rank $\forall s \in \mathbb{C}$ (the PBH test).*
3. *The observability matrix*
$$M_o \coloneqq \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{bmatrix}$$
has full rank.
4. *The pair (A', C') is controllable.*

A realization is minimal if and only if it is controllable and observable.

Invariant zeros of the realization (C.6) are defined as the points $\lambda \in \mathbb{C}$ at which

$$\operatorname{rank} \begin{bmatrix} A - \lambda I & B \\ C & D \end{bmatrix} < \operatorname{nrank} \begin{bmatrix} A - sI & B \\ C & D \end{bmatrix}$$

(the matrix polynomial of s in the right-hand side is dubbed the Rosenbrock system matrix). The set of all invariant zeros comprises transmission zeros of the transfer function $G(s)$ and hidden modes of realization (C.6).

Poles and zeros have (spatial) directions for MIMO systems. Assume through the rest of this appendix that the realization in (C.6) is minimal. By *input* and *output* directions of a realization pole λ of (C.6), we understand the subspaces

$$\operatorname{pdir}_{\mathrm{i}}(G, \lambda) \coloneqq B^{\top} \ker(\lambda I - A)^{\top} \subset \mathbb{C}^m \tag{C.7a}$$

and

$$\operatorname{pdir}_{\mathrm{o}}(G, \lambda) \coloneqq C \ker(\lambda I - A) \subset \mathbb{C}^p, \tag{C.7b}$$

respectively. If λ is not a pole of $G(s)$, then both definitions in (C.7) result in the trivial subspace $\{0\}$.

Lemma C.1.9. *If $\lambda \in \mathbb{C}$ is a pole of $G(s)$, then*

i) λ is a pole of $G(s)v$ whenever $0 \neq v \in \mathrm{pdir_i}(G, \lambda)$,

ii) λ is a pole of $v^\top G(s)$ whenever $0 \neq v \in \mathrm{pdir_o}(G, \lambda)$.

Proof Bring in a minimal realization of G as in (C.6). If (A, Bv) is controllable, then every eigenvalue of A is a pole of $G(s)v$, by the observability of (C, A). If (A, Bv) is uncontrollable, without loss of generality we may assume that

$$(A, B) = \left(\begin{bmatrix} A_{\mathrm{c}} & A_{12} \\ 0 & A_{\bar{\mathrm{c}}} \end{bmatrix}, \begin{bmatrix} B_{\mathrm{c}} \\ B_{\bar{\mathrm{c}}} \end{bmatrix} \right)$$

with controllable $(A_{\mathrm{c}}, B_{\mathrm{c}}v)$ and $B_{\bar{\mathrm{c}}}v = 0$. In this case λ is not a pole of $G(s)v$ iff $\lambda \notin \mathrm{spec}(A_{\mathrm{c}})$. So assume that $\lambda \notin \mathrm{spec}(A_{\mathrm{c}})$, which implies that $\lambda \in \mathrm{spec}(A_{\bar{\mathrm{c}}})$ and that

$$B^\top \ker(\lambda I - A)^\top \subset \begin{bmatrix} B_{\mathrm{c}}^\top & B_{\bar{\mathrm{c}}}^\top \end{bmatrix} \mathrm{Im} \begin{bmatrix} 0 \\ I \end{bmatrix} = \mathrm{Im}\, B_{\bar{\mathrm{c}}}^\top .$$

But then $v \in \mathrm{pdir_i}(G, \lambda) \implies v \in \mathrm{Im}\, B_{\bar{\mathrm{c}}}^\top = (\ker B_{\bar{\mathrm{c}}})^\perp$, which contradicts the condition $B_{\bar{\mathrm{c}}}v = 0$. Hence, λ must be a pole of $G(s)v$. The second item follows by similar arguments. ∎

Input and output directions of an invariant zero λ are defined as

$$\mathrm{zdir_i}(G, \lambda) := \begin{bmatrix} 0 & I_m \end{bmatrix} \ker \begin{bmatrix} A - \lambda I & B \\ C & D \end{bmatrix} \subset \mathbb{C}^m \tag{C.8a}$$

and

$$\mathrm{zdir_o}(G, \lambda) := \begin{bmatrix} 0 & I_p \end{bmatrix} \ker \begin{bmatrix} A - \lambda I & B \\ C & D \end{bmatrix}^\top \subset \mathbb{C}^p, \tag{C.8b}$$

respectively. With some abuse of notation we use the definitions in (C.8) also if λ is not an invariant zero of (C.6), but the normal rank of $G(s)$ is deficient. For example, in our notation

$$\mathrm{zdir_i}\left(\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}, \lambda \right) = \mathrm{zdir_o}\left(\begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix}, \lambda \right) = \mathrm{Im}\, \mathbb{1}_2$$

for all $\lambda \in \mathbb{C}$. In such situations directions are understood as normal null spaces.

Lemma C.1.10. *If $\lambda \notin \mathrm{spec}(A)$, then it is an invariant zero of G iff* $\mathrm{rank}\, G(\lambda) < \mathrm{nrank}\, G(s)$ *and*

$$\mathrm{zdir_i}(G, \lambda) = \ker G(\lambda) \quad \textit{and} \quad \mathrm{zdir_o}(G, \lambda) = \ker [G(\lambda)]^\top .$$

Proof Follows from the relations

$$\begin{bmatrix} A-\lambda I & B \\ C & D \end{bmatrix} = \begin{bmatrix} A-\lambda I & 0 \\ C & G(\lambda) \end{bmatrix} \begin{bmatrix} I & (A-\lambda I)^{-1}B \\ 0 & I \end{bmatrix}$$
$$= \begin{bmatrix} I & 0 \\ C(A-\lambda I)^{-1} & I \end{bmatrix} \begin{bmatrix} A-\lambda I & B \\ 0 & G(\lambda) \end{bmatrix}$$

and the assumed invertibility of $A - \lambda I$. ■

Lemma C.1.11. *If $\lambda \in \bar{\mathbb{C}}_0$, then it is a pole of $G(s)$ if and only if it is a zero of the denominators $M_G(s)$ and $\tilde{M}_G(s)$ of its coprime factorizations. Moreover,*

$$\mathrm{pdir_i}(G,\lambda) = \mathrm{zdir_o}(M_G,\lambda) \quad \textit{and} \quad \mathrm{pdir_o}(G,\lambda) = \mathrm{zdir_i}(\tilde{M}_G,\lambda)$$

in this case.

Proof Follows by [104, Prop. 4.16] and the fact that a pole of $G(s)$ in $\bar{\mathbb{C}}_0$ is a zero of all possible denominators. ■

More details can be found in [93], although we use slightly different definitions of directions (subspaces, rather than vectors), in line with [104].

Bibliography

[1] J. Tolkien, *The Fellowship of the Ring*. London, England: HarperCollins, 2020.

[2] D. D. Šiljak, *Large-Scale Dynamic Systems: Stability and Structure*. NY: North-Holland, 1978.

[3] E. J. Davison, A. G. Aghdam, and D. E. Miller, *Decentralized Control of Large-Scale Systems*. New York, NY: Springer-Verlag, 2020.

[4] N. R. Sandell Jr., P. Varaiya, M. Athans, and M. G. Safonov, "Survey of decentralized control methods for large scale systems," *IEEE Trans. Automat. Control*, vol. 23, no. 2, pp. 108–128, 1978.

[5] C. Papadimitriou and J. Tsitsiklis, "Intractable problems in control theory," *SIAM J. Control Optim.*, vol. 24, pp. 639–654, 1986.

[6] V. Blondel and J. Tsitsiklis, "A survey of computational complexity results in systems and control," *Automatica*, vol. 36, no. 9, pp. 1249–1274, 2000.

[7] 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.

[8] A. Rantzer and M. Valcher, "Scalable control of positive systems," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 4, no. 1, pp. 319–341, 2021.

[9] M. Rotkowitz and S. Lall, "A characterization of convex problems in decentralized control," *IEEE Trans. Automat. Control*, vol. 51, no. 2, pp. 274–286, 2006.

[10] A. Mahajan, N. C. Martins, M. C. Rotkowitz, and S. Yüksel, "Information structures in optimal decentralized control," in *Proc. 51st IEEE Conf. Decision and Control*, Maui, HW, 2012, pp. 1291–1306.

[11] M. di Bernardo, A. Salvi, and S. Santini, "Distributed consensus strategy for platooning of vehicles in the presence of time-varying heterogeneous communication delays," *IEEE Transactions on Intelligent Transportation Systems*, vol. 16, no. 1, pp. 102–112, 2015.

[12] J. A. Fax and R. M. Murray, "Information flow and cooperative control of vehicle formations," *IEEE Trans. Automat. Control*, vol. 49, no. 9, pp. 1465–1476, 2004.

[13] S. Zhao and D. Zelazo, "Translational and scaling formation maneuver control via a bearing-based approach," *IEEE Transactions on Control of Network Systems*, vol. 4, no. 3, pp. 429–438, 2017.

[14] I. F. Akyildiz, Weilian Su, Y. Sankarasubramaniam, and E. Cayirci, "A survey on sensor networks," *IEEE Communications Mag.*, vol. 40, no. 8, pp. 102–114, 2002.

[15] R. Olfati-Saber, "Distributed kalman filter with embedded consensus filters," in *Proceedings of the 44th IEEE Conference on Decision and Control*, 2005, pp. 8179–8184.

[16] M. DeGroot, "Reaching a consensus," *Journal of the American Statistical Association*, vol. 69, no. 345, pp. 118–121, 1974.

[17] J. Tsitsiklis, D. Bertsekas, and M. Athans, "Distributed asynchronous deterministic and stochastic gradient optimization algorithms," *IEEE Trans. Automat. Control*, vol. 31, no. 9, pp. 803–812, 1986.

[18] D. Bertsekas and J. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1989.

[19] A. Agarwal, M. Wainwright, and J. Duchi, "Distributed dual averaging in networks," in *Advances in Neural Information Processing Systems (NIPS)*, vol. 23, 2010.

[20] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, "Novel type of phase transition in a system of self-driven particles," *Phys. Rev. Lett.*, vol. 75, pp. 1226–1229, 6 1995.

[21] W. Ren and R. W. Beard, *Distributed Consensus in Multi-vehicle Cooperative Control: Theory and Applications*. London: Springer-Verlag, 2008.

[22] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*. Princeton: Princeton University Press, 2010.

[23] C. D. Godsil and G. F. Royle, *Algebraic Graph Theory*. Springer, 2001.

[24] R. Olfati-Saber, A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proc. IEEE*, vol. 95, no. 1, pp. 215–233, 2007.

[25] R. S. Smith and F. Y. Hadaegh, "Control of deep-space formation-flying spacecraft; relative sensing and switched information," *Journal of Guidance, Control, and Dynamics*, vol. 28, no. 1, pp. 106–114, 2005.

[26] U. A. Khan, S. Kar, and J. M. F. Moura, "Distributed sensor localization in random environments using minimal number of anchor nodes," *IEEE Transactions on Signal Processing*, vol. 57, no. 5, pp. 2000–2016, 2009.

[27] D. Zelazo and M. Mesbahi, "Graph-theoretic analysis and synthesis of relative sensing networks," *IEEE Transactions on Automatic Control*, vol. 56, no. 5, pp. 971–982, 2011.

[28] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Automat. Control*, vol. 48, no. 6, pp. 988–1001, 2003.

[29] R. Olfati-Saber and R. M. Murray, "Consensus problems in networks of agents with switching topology and time-delays," *IEEE Trans. Automat. Control*, vol. 49, no. 9, pp. 1520–1533, 2004.

[30] Z. Lin, M. Broucke, and B. Francis, "Local control strategies for groups of mobile autonomous agents," *IEEE Trans. Automat. Control*, vol. 49, no. 4, pp. 622–629, 2004.

[31] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Transactions on Automatic Control*, vol. 50, no. 2, pp. 169–182, 2005.

[32] M. Arcak, "Passivity as a design tool for group coordination," *IEEE Trans. Automat. Control*, vol. 52, no. 8, pp. 1380–1390, 2007.

[33] D. Zelazo and M. Mesbahi, "Edge agreement: Graph-theoretic performance bounds and passivity analysis," *IEEE Transactions on Automatic Control*, vol. 56, no. 3, pp. 544–555, 2011.

[34] A. Rahmani, M. Ji, M. Mesbahi, and M. Egerstedt, "Controllability of multiagent systems from a graph-theoretic perspective," *SIAM J. Control Optim.*, vol. 48, no. 1, pp. 162–186, 2009.

[35] X. Dong, B. Yu, Z. Shi, and Y. Zhong, "Time-varying formation control for unmanned aerial vehicles: Theories and applications," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 1, pp. 340–348, 2015.

[36] J. Seo, H. Shim, and J. Back, "Consensus of high-order linear systems using dynamic output feedback compensator: Low gain approach," *Automatica*, vol. 45, no. 11, pp. 2659–2664, 2009.

[37] A. Seuret, D. Dimarogonas, and K. Johansson, "Consensus under communication delays," in *Proc. 47th IEEE Conf. Decision and Control*, 2008, pp. 4922–4927.

[38] W. Yu, G. Chen, M. Cao, and W. Ren, "Delay-induced consensus and quasi-consensus in multi-agent dynamical systems," *IEEE Trans. Circuits Syst. I*, vol. 60, no. 10, pp. 2679–2687, 2013.

[39] W. Yu, G. Chen, and M. Cao, "Some necessary and sufficient conditions for second-order consensus in multi-agent dynamical systems," *Automatica*, vol. 46, no. 6, pp. 1089–1095, 2010.

[40] T. Li and J. Zhang, "Mean square average-consensus under measurement noises and fixed topologies: Necessary and sufficient conditions," *Automatica*, vol. 45, no. 8, pp. 1929–1936, 2009.

[41] T. Li, F. Wu, and J. Zhang, "Multi-agent consensus with relative-state-dependent measurement noises," *IEEE Trans. Automat. Control*, vol. 59, no. 9, pp. 2463–2468, 2014.

[42] L. Scardovi and R. Sepulchre, "Synchronization in networks of identical linear systems," *Automatica*, vol. 45, no. 11, pp. 2557–2562, 2009.

[43] Z. Li, Z. Duan, G. Chen, and L. Huang, "Consensus of multiagent systems and synchronization of complex networks: A unified viewpoint," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 57, no. 1, pp. 213–224, 2010.

[44] P. Wieland, R. Sepulchre, and F. Allgöwer, "An internal model principle is necessary and sufficient for linear output synchronization," *Automatica*, vol. 47, no. 5, pp. 1068–1074, 2011.

[45] H. Knobloch and A. Isidori, *Topics in control theory* (Oberwolfach Seminars). Basel, Switzerland: Birkhäuser, 1993.

[46] A. Isidori, L. Marconi, and G. Casadei, "Robust output synchronization of a network of heterogeneous nonlinear agents via nonlinear regulation theory," *IEEE Trans. Automat. Control*, vol. 59, no. 10, pp. 2680–2691, 2014.

[47] H. Kim, H. Shim, and J.-H. Seo, "Output consensus of heterogeneous uncertain linear multi-agent systems," *IEEE Trans. Automat. Control*, vol. 56, no. 1, pp. 200–206, 2011.

[48] X. Ge, Q.-L. Han, D. Ding, X.-M. Zhang, and B. Ning, "A survey on recent advances in distributed sampled-data cooperative control of multi-agent systems," *Neurocomputing*, vol. 275, pp. 1684–1701, 2018.

[49] M. Xing, F. Deng, and Z. Hu, "Sampled-data consensus for multiagent systems with time delays and packet losses," *IEEE Trans. Syst. Man, and Cybernetics: Systems*, vol. 50, no. 1, pp. 203–210, 2020.

[50] Z. Ding, "Consensus disturbance rejection with disturbance observers," *IEEE Transactions on Industrial Electronics*, vol. 62, no. 9, pp. 5829–5837, 2015.

[51] S. Mou, M.-A. Belabbas, S. M. Z. Sun, and B. Anderson, "Undirected rigid formations are problematic," *IEEE Trans. Automat. Control*, vol. 61, no. 10, pp. 2821–2836, 2016.

[52] L. Cheng, Z. Hou, M. Tan, and X. Wang, "Necessary and sufficient conditions for consensus of double-integrator multi-agent systems with measurement noises," *IEEE Trans. Automat. Control*, vol. 56, no. 8, pp. 1958–1963, 2011.

[53] M. Bürger and C. De Persis, "Dynamic coupling design for nonlinear output agreement and time-varying flow control," *Automatica*, vol. 51, pp. 210–222, 2015.

[54] G. Lang and J. M. Ham, "Conditional feedback systems—a new approach to feedback control," *Trans. American Inst. Electrical Engin., Part II: Applications and Industry*, vol. 74, no. 3, pp. 152–161, 1955.

[55] D. C. Youla and J. J. Bongiorno Jr., "A feedback theory of two-degree-of-freedom optimal Wiener-Hopf design," *IEEE Trans. Automat. Control*, vol. 30, no. 7, pp. 652–665, 1985.

[56] O. Smith, "Closer control of loops with dead time," *Chem. Eng. Progress*, vol. 53, no. 5, pp. 217–219, 1957.

[57] L. Mirkin, "On the extraction of dead-time controllers and estimators from delay-free parametrizations," *IEEE Trans. Automat. Control*, vol. 48, no. 4, pp. 543–553, 2003.

[58] G. Tadmor, "H_∞ optimal sampled-data control in continuous time systems," *Int. J. Control*, vol. 56, no. 1, pp. 99–141, 1992.

[59] G. Barkai, L. Mirkin, and D. Zelazo, "On sampled-data consensus: Divide and concur," *IEEE Control Syst. Lett.*, vol. 6, pp. 343–348, 2022.

[60] G. Barkai, L. Mirkin, and D. Zelazo, "An emulation approach to sampled-data synchronization," in *Proc. 62nd IEEE Conf. Decision and Control*, Singapore, 2023, pp. 6449–6454.

[61] G. Barkai, L. Mirkin, and D. Zelazo, "An emulation approach to output-feedback sampled-data synchronization," in *Proc. 22nd European Control Conf.*, 2024.

[62] G. Barkai, L. Mirkin, and D. Zelazo, "Asynchronous sampled-data synchronization with small communications delays," in *Proc. 63rd IEEE Conf. Decision and Control.*, 2024.

[63] G. Barkai, L. Mirkin, and D. Zelazo, "On the internal stability of diffusively coupled multi-agent systems and the dangers of cancel culture," *Automatica*, vol. 155, p. 111 158, 2023.

[64] L. Carroll, *Alice in Wonderland* (Oxford Bookworms, Green S.). London, England: Oxford University Press, 1994.

[65] A. Laub, *Matrix Analysis for Scientists & Engineers*. Philadelphia: SIAM, 2005.

[66] M. Mattioni, "On multiconsensus of multi-agent systems under aperiodic and asynchronous sampling," *IEEE Control Syst. Lett.*, vol. 4, no. 4, pp. 839–844, 2020.

[67] W. Sun, K. M. Nagpal, and P. P. Khargonekar, "$\mathcal{H}_\infty$ control and filtering for sampled-data systems," *IEEE Trans. Automat. Control*, vol. 38, no. 8, pp. 1162–1174, 1993.

[68] L. Mirkin, H. Rotstein, and Z. J. Palmor, "H^2 and H^∞ design of sampled-data systems using lifting. Part I: General framework and solutions," *SIAM J. Control Optim.*, vol. 38, no. 1, pp. 175–196, 1999.

[69] L. Mirkin, "Intermittent redesign of analog controllers via the Youla parameter," *IEEE Trans. Automat. Control*, vol. 62, no. 4, pp. 1838–1851, 2017.

[70] L. A. Montestruque and P. J. Antsaklis, "On the model-based control of networked systems," *Automatica*, vol. 39, no. 10, pp. 1837–1843, 2003.

[71] J. Lunze and D. Lehmann, "A state-feedback approach to event-based control," *Automatica*, vol. 46, no. 1, pp. 211–215, 2010.

[72] F. Bullo, *Lectures on Network Systems*, 1.6. Kindle Direct Publishing, 2024. [Online]. Available: `https://fbullo.github.io/lns`.

[73] A. Weir, *Project hail Mary*. New York, NY: Ballantine Books, May 2021.

[74] E. Garcia, P. J. Antsaklis, and L. A. Montestruque, *Model-Based Control of Networked Systems*. Boston: Birkhäuser, 2014.

[75] Y. Wu, H. Su, P. Shi, Z. Shu, and Z. Wu, "Consensus of multiagent systems using aperiodic sampled-data control," *IEEE Transactions on Cybernetics*, vol. 46, no. 9, pp. 2132–2143, 2016.

[76] J. Sun, Z. Wang, and N. Rong, "Sampled-data consensus of multiagent systems with switching jointly connected topologies via time-varying lyapunov function approach," *Int. J. Robust and Nonlinear Control*, vol. 30, no. 14, pp. 5369–5385, 2020.

[77] H. Abou-Kandil, G. Freiling, V. Ionescu, and G. Jank, *Matrix Riccati Equations in Control and Systems Theory*. Birkhäuser, 2003.

[78] J. Fu, G. Wen, W. Yu, T. Huang, and J. Cao, "Exponential consensus of multiagent systems with Lipschitz nonlinearities using sampled-data information," *IEEE Trans. Circuits Syst. I*, vol. 65, no. 12, pp. 4363–4375, 2018.

[79] B. Wang, Y.-P. Tian, and Z. Han, "Convergence to zero of quadratic Lyapunov functions for multi-agent systems in time-varying directed networks," *IEEE Trans. Automat. Control*, vol. 68, no. 12, pp. 8178–8184, 2023.

[80] T. Ménard, E. Moulay, P. Coirault, and M. Defoort, "Observer-based consensus for second-order multi-agent systems with arbitrary asynchronous and aperiodic sampling periods," *Automatica*, vol. 99, pp. 237–245, 2019.

[81] A. Goldenshluger and L. Mirkin, "H_2 control under intermittent sampling and small communication delays," *IEEE Control Syst. Lett.*, vol. 3, no. 3, pp. 583–588, 2019.

[82] M. E. J. Newman, "The structure and function of complex networks," *SIAM Review*, vol. 45, no. 2, pp. 167–256, 2003.

[83] L. Xiao and S. Boyd, "Fast linear iterations for distributed averaging," *Syst. Control Lett.*, vol. 53, no. 1, pp. 65–78, 2004.

[84] Y. Shafi, M. Arcak, and L. E. Ghaoui, "Graph weight allocation to meet laplacian spectral constraints," *IEEE Trans. Automat. Control*, vol. 57, no. 7, pp. 1872–1877, 2012.

[85] L. Xiao, S. Boyd, and S. Kim, "Distributed average consensus with least-mean-square deviation," *Journal of Parallel and Distributed Computing*, vol. 67, no. 1, pp. 33–46, 2007.

[86] D. Jakovetić, J. Xavier, and J. Moura, "Weight optimization for consensus algorithms with correlated switching topology," *IEEE Trans. Signal Processing*, vol. 58, no. 7, pp. 3788–3801, 2010.

[87] S. van Dijk, P. Chanfreut, and W. Heemels, "Decentralized design for lq consensus in multi-agent systems," in *Proc. 62nd IEEE Conf. Decision and Control*, 2023, pp. 7427–7432.

[88] M. E. Chamie and T. Başar, "Optimal strategies for dynamic weight selection in consensus protocols in the presence of an adversary," in *Proc. 53rd IEEE Conf. Decision and Control*, 2014, pp. 735–740.

[89] W. Shakespeare, *The Tragedy of Hamlet, Prince of Denmark* (Penguin Shakespeare). Harlow, England: Penguin Books, May 1970.

[90] T. Yucelen and M. Egerstedt, "Control of multiagent systems under persistent disturbances," in *"Proc. 2012 American Control Conf.,*, 2012, pp. 5264–5269.

[91] L. Mo and S. Guo, "Consensus of linear multi-agent systems with persistent disturbances via distributed output feedback," *Journal of Systems Science and Complexity*, vol. 32, no. 3, pp. 835–845, 2019.

[92] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control*. Englewood Cliffs, NJ: Prentice-Hall, 1996.

[93] S. Skogestad and I. Postlethwaite, *Multivariable Feedback Control: Analysis and Design*, 2nd. Chichester: John Wiley & Sons, 2005.

[94] T. T. Georgiou and M. C. Smith, "Graphs, causality and stabilizability: Linear, shift-invariant systems on $\mathcal{L}_2[0,\infty)$," *Math. Control, Signals and Systems*, vol. 6, pp. 195–223, 1993.

[95] R. F. Curtain and H. Zwart, *Introduction to Infinite-Dimensional Systems Theory: A State-Space Approach*. New York, NY: Springer-Verlag, 2020.

[96] M. C. Smith, "On stabilization and the existence of coprime factorizations," *IEEE Trans. Automat. Control*, vol. 34, no. 9, pp. 1005–1007, 1989.

[97] P. A. Fuhrmann, "On the corona theorem and its application to spectral problems in Hilbert space," *Trans. Amer. Math. Soc.*, vol. 132, no. 1, pp. 55–66, 1968.

[98] W. Rudin, *Real and Complex Analysis*, 3rd. New York, NY: McGraw-Hill, 1987.

[99] M. A. Belabbas, X. Chen, and D. Zelazo, "On structural rank and resilience of sparsity patterns," *IEEE Trans. Automat. Control*, vol. 68, no. 8, pp. 4783–4795, 2023.

[100] S. Sundaram and C. N. Hadjicostis, "Distributed function calculation and consensus using linear iterative strategies," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 4, pp. 650–660, 2008.

[101] J. R. Partington and C. Bonnet, "H_∞ and BIBO stabilization of delay systems of neutral type," *Syst. Control Lett.*, vol. 52, no. 8, pp. 283–288, 2004.

[102] M. S. Verma, "Coprime fractional representations and stability of non-linear feedback systems," *Int. J. Control*, vol. 48, pp. 897–918, 1988.

[103] M. Sharf and D. Zelazo, "A network optimization approach to cooperative control synthesis," *IEEE Control Syst. Lett.*, vol. 1, no. 1, pp. 86–91, 2017.

[104] L. Mirkin, *Linear Control Systems*, course notes, Faculty of Mechanical Eng., Technion—IIT, 2024. [Online]. Available: `http://leo.technion.ac.il/Courses/LCS/LCSnotes.pdf`.

[105] B. D. O. Anderson and J. B. Moore, "Time-varying feedback laws for decentralized control," *IEEE Trans. Automat. Control*, vol. 26, no. 5, pp. 1133–1139, 1981.

[106] B. Sanderson, *Oathbringer* (Stormlight Archive). Tor Books, 2017.

[107] H. Kwakernaak and R. Sivan, *Linear Optimal Control Systems*. New York, NY: John Wiley & Sons, 1972.

[108] L. Qiu and K. Zhou, *Introduction to Feedback Control*. Upper Saddle River, NJ: Prentice-Hall, 2010.

[109] B. A. Francis and W. M. Wonham, "The internal model principle of control theory," *Automatica*, vol. 12, no. 5, pp. 457–465, 1976.

[110] A. Dumas, *The Count of Monte Cristo* (Everyman's Library Classics & Contemporary Classics). New York, NY: Random House, Jun. 2009.

[111] L. Mirkin, H. Rotstein, and Z. J. Palmor, "H^2 and H^∞ design of sampled-data systems using lifting. Part II: Properties of systems in the lifted domain," *SIAM J. Control Optim.*, vol. 38, no. 1, pp. 197–218, 1999.

[112] R. Schaller, "Moore's law: Past, present and future," *IEEE Spectrum*, vol. 34, no. 6, pp. 52–59, 1997.

[113] B. Anderson, C. Yu, B. Fidan, and J. Hendrickx, "Rigid graph control architectures for autonomous formations," *IEEE Control Syst. Mag.*, vol. 28, no. 6, pp. 48–63, 2008.

[114] K.-K. Oh and H.-S. Ahn, "Distance-based undirected formations of single-integrator and double-integrator modeled agents in n-dimensional space," *Int. J. Robust and Nonlinear Control*, vol. 24, no. 12, pp. 1809–1820, 2014.

[115] S. Zhao and D. Zelazo, "Bearing rigidity and almost global bearing-only formation stabilization," *IEEE Trans. Automat. Control*, vol. 61, no. 5, pp. 1255–1268, 2016.

[116] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[117] A. Tannenbaum, "Feedback stabilization of linear dynamical plants with uncertainty in the gain factor," *Int. J. Control*, vol. 32, no. 1, pp. 1–16, 1980.

[118] M. Vidyasagar, *Control System Synthesis: A Factorization Approach, Part II* (Synthesis lectures on control and mechatronics). Morgan & Claypool Publishers, 2011.

[119] M. Vidyasagar, *Control System Synthesis: A Factorization Approach*. Cambridge, MA: The MIT Press, 1985.

[120] J. Bondy and U. Murty, *Graph Theory*. Springer Publishing Company, Incorporated, 2008.

[121] M. Fiedler, "Algebraic connectivity of graphs," *Czechoslovak Mathematical Journal*, vol. 23, no. 2, pp. 298–305, 1973.

[122] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 2nd ed. Cambridge University Press, 2012.

[123] C. Hoppen, J. Monsalve, and V. Trevisan, "Spectral norm of oriented graphs," *Linear Algebra and its Applications*, vol. 574, pp. 167–181, 2019.

[124] G. Weiss, "Representation of shift-invariant operators on L_2 by H_∞ transfer functions: An elementary proof, a generalization to L_p, and a counterexample for L_∞," *Mathematics of Control, Signals and Systems*, vol. 4, no. 2, pp. 193–203, 1991.

סוכנים המבוקרות בעזרת בקרים דיפוסיביים, דוגמת פרוטוקול הקונצנזוס, תחת תנאים מסוימים. בפרט, בהינתן פירוק לגורמים זרים יציבים של הסוכנים, אם כל המכנים חולקים גרעין לא טריוויאלי עבור נקודות כלשהן בחצי המישור הימני אזי המערכת אינה ניתנת לייצוב בעזרת בקר דיפוסיבי. עבור סוכנים מממד סופי, הדבר מתבטא בקוטב לא יציב וכיוון משותפים המצטמצמים בחיבור הטורי, אך מופיעים באחת מן התמסורות של החוג הסגור. זו דוגמה לתופעה ייחודית של מערכות מרובות־קלט מרובות־פלט: צמצומי קטבים הנובעים ממחסור בדרגה הנורמלית של הבקר. תוצאה זו מסבירה את השבריריות של מערכות שכאלו תחת הפרעות ורעשי מדידה אשר צויינה במספר מקומות בספרות ללא הסבר. תוצאה זו גם מסבירה מדוע בקרים המנסים להביא להסכמה תחת הפרעות ורעשים חיצוניים נוטים לכלול אלמנט מבוסס קונצנזוס ועוד אלמנט מקומי לייצוב המערכת.

מגבלה יסודית זו מעידה על צורך לחפש ארכיטקטורת בקרה חלופית למערכות מבוזרות. בעבודה זו מוצגת שיטת בקרה חלופית שכזו, השואבת השראה מבקרת שתי דרגות חופש מבעיות עקיבה קלאסיות, המייצרת הפרדה בין יציבות בחוג המקומי לבין התיאום בין הסוכנים דרך הרשת. תחת חוק בקרה זה, תגובת החוג הסגור לרעשי מדידה תלויה במסנן רשת אחיד ובגרף התקשורת, אך לא בדינמיקת הסוכנים והחוגים המקומיים. לעומת זאת, התגובה להפרעות מקומיות תלויה הן בדינמיקה המקומית והגלובלית, אך בצורה פשוטה יותר מאשר זו המתקבלת משימוש בבקרים מבוססי קונצנזוס. תכונה מעניינת של הארכיטקטורה החדשה היא שתחת ההנחה שקיים לפחות סוכן אחד "מוגן", תכנון נכון של הבקרים המקומיים יכול לדחות באופן מושלם הפרעות מקומיות. שתי תכונות אלו חסרות בפרוטוקולי הסכמה דיפוסיבים, ונובעות מן העובדה כי בארכיטקטורה זו עדיין קיים ענף משובי בכיוון משתנה ההסכמה, דבר שמאפשר הן לדחות הפרעות והן להנחית את השפעת רעשי המדידה על משתנה זה.

הניתוח והתכן בעבודה זו משלבים כלים מתורת הגרפים הספטקרלית, בקרה דגומה, ובקרה קלאסית, ומייצרים תוצאות תאורטיות חדשות, כמו גם שיטות תכן מעשיות המתאימות למגוון רחב של בעיות בקרה עבור סוכנים לינארי קבועים-בזמן. הפרק האחרון בעבודה זו מוקדש לסיכום ולכיווני מחקר עתידיים, דוגמת הרחבת שיטת האמולציה לסוכנים הטרוגניים ולארכיטקטורת שתי דרגות החופש, תכנון מסנני רשת אופטימליים באופן סיסטמטי ושאלות נוספות הנוגעות לביצועים תחת אילוצי תקשורת וחישה.

תקציר

עבודת דוקטורט זו עוסקת בבעיות סנכרון והסכמה במערכות מרובות־סוכנים בזמן רציף, הפועלות תחת אילוצי דגימה בנוסף לאילוצי תקשורת מרחביים. העבודה מאתגרת את הגישה הרווחת הנשענת על פרוטוקולי קונצנזוס ובקרים דיפוסיביים, בטענה כי הפשטות המבנית שלהם עלולה להסוות מגבלות יסודיות—במיוחד בנוכחות הפרעות חיצוניות, רעש מדידה ודינמיקה בלתי יציבה.

המחקר נפתח בבעיית הקונצנזוס עבור סוכנים אינטגרליים מסדר ראשון, המחליפים מידע בצורה לא מחזורית ובזמני דגימה א־סינכרוניים. בעייה זו נפתרה בעבר בעזרת דיסקרטיזציה של פרוטוקול הקונצנזוס לזמן רציף בעזרת מחזיק מסדר אפס ודוגם אידאלי. צורת בקרה זו מוסיפה באופן טבעי שמרנות למערכת, לדוגמא בזמן רציף הגבר הפרוטוקול אינו מוגבל, בעוד שבגירסה הדגומה קיים חסם עליון שהוא פונקציה של זמן הדגימה וגרף התקשורת. היות וככל שהגבר הבקר גבוה קצב ההתכנסות להסכמה גדל, שימוש בבקר זה מחייב חסם קשיח על הביצועים האפשריים. על מנת לפתור בעיה זו, מוצע פרוטוקול חדש הלוקח השראה מעקרונות קלאסיים בבקרה דגומה אופטימלית. בפרט, כל סוכן מבצע אמולציה מקומית של דינמיקה אנלוגית גלובלית ומשתמש בה בכדי לייצר את אות הבקרה בין רגעי הדגימה ברגעי הדגימה, כל סוכן משדר את מרכז-המסה של האמולטור שלו לשכניו, במקום לשדר את המדידה האמיתית בלבד. תחת חוק בקרה זה, הדינמיקה בחוג הסגור מתפרקת באופן טבעי לשני חלקים נפרדים: מרכזי-מסה ואי-הסכמות. בפרט, עבור מרכזי-המסה מתקבלת דינמיקת קונצנזוס בדיד תחת גרף המשתנה בזמן, אך ללא תלות בזמני או מרווחי הדגימה. דינמיקה זו מתכנסת להסכמה תחת הנחות סטנדרטיות על קשירות גרף האיחוד וזמני דגימה חסומים ללא תלות בסדרת הדגימות עצמה או בהגבר הבקר. יתר על כן, עבור בחירה מסוימת של דינמיקת האמולטורים, כל סוכן יכול לממש את חוק הבקרה בעזרת בקר היברידי מסדר ראשון ללא תלות במספר הסוכנים הכולל.

לאחר מכן, אותה גישת בקרה מבוססת אמולציה מורחבת למקרים כלליים יותר, בפרט, עבור סוכנים זהים בעלי דינמיקה לינארית קבועה-בזמן המסנכרנים לאות דינמי תחת משוב מצב מלא או משוב יציאה בלבד. טכניקות ההוכחה במקרים אלו מורכבות יותר עקב הדינמיקה הנוספת, אך, באופן מפתיע, תנאי הקשירות בין הסוכנים נותר זהה גם במקרה הכללי. עבור המקרה הכללי נעשתה גם אנליזת ביצועים, ונמצא תנאי מספיק להתכנסות מעריכית להסכמה על אף הדגימה הלא סדירה, תוצאה שאינה סטנדרטית בארכיטקטורות בקרה אחרות. לבסוף מוצגות שתי הרחבות של שיטת הבקרה המוצעת. הראשונה מוסיפה חזאי על מנת לפצות על השהיות קטנות אך לא ידועות בחילופי המידע, והשנייה כוללת משקלים דינמיים עבור היתוך המידע הנשלח בין הסוכנים, משקלים שעשויים לשפר את ביצועי הבקר.

תרומה תאורטית מרכזית של המחקר בעבודה זו היא זיהוי אי יציבות פנימית במערכות מרובות

תודות

ראשית, ברצוני להודות לשני מנחי, פרופ' לאוניד מירקין ופרופ' דניאל זלזו, על ההנחיה, העידוד, החוכמה, ובעיקר – על הסבלנות. נדיר למצוא אפילו מנחה אחד שמקדיש עצמו להתפתחותו של תלמידו כחוקר ולרווחתו כאדם, ולא רק לקידום הקריירה האישית שלו. בדרך פלא, כאשר גילגלו את הקוביות עבורי, זכיתי בשניים כאלה. בסופו של יום אתם עיצבתם אותי כחוקר,בין אם תרצו ובין אם לא, אז מקווה שאתם מרוצים מן התוצאה.

דרכי לכאן לא הייתה שגרתית, ואני מבקש להודות גם לכמה אנשים נוספים, שחלקם בלי לדעת או בלי כוונה, סללו עבורי את הדרך. למורב למתמטיקה שלי בתיכון שאמרה לי "מקצועות ריאלים זה לא בשבילך, נסה משהו אחר", ושיבצה אותי לשלוש יחידות מתמטיקה בלי יכולת לערער. הכרחת אותי להתמודד עם תסמונת המתחזה הרבה לפני רוב החוקרים, אז אפילו מזה יצאה קצת תועלת. לאייל הרשקו, המורה שלי לרובוטיקה בתיכון, שמייד התעלם מדעתה והצית בי את התשוקה להנדסה. אתה הראית לי הלכה למעשה את ההבדל בין מורה למנטור, וכמה שניהם חיוניים להצלחה -- תודה שהייתה ועודך מנטור עבורי. לפרופ' אלון וולף, שאמר לי להפסיק לבזבז זמן בעבודות מזדמנות, לקח את ידי, והוליך אותי פיזית אל המרכז הקדם־אקדמי של הטכניון. אותה שיחה פשוטה — שסביר להניח שאינך זוכר — גרמה לי להפסיק להתמהמה, להתמודד עם הפחד מכישלון ולהירשם. לחברי ליאור ימאי ויותם קנת, שלא רק שהם חברים טובים, אלא גם עזרו לי לצלוח את הסערה שהיא השנה הראשונה בטכניון. לגב' דניאלה אבידן, שהטרדתי ללא הרף בשאלות על הוכחות חסרות והכללות אחרי כל הרצאה במד"ר — עד שגררה אותי למחלקה למתמטיקה והציעה שאירשם. אמנם לא סיימתי תואר במתמטיקה, אך עידודך גרם לי לקחת עוד קורסים, ללמוד דברים חדשים, ובעיקר להבין שהמחסום העיקרי ללמידת דברים חדשים הוא הפחד לנסות. לפרופ' רודולף ספולקרה ולפרופ' אימן שאמס, שהזמינו אותי לקבוצות שלהם רק כדי שאלמד דברים חדשים ואראה זוויות נוספות.

ולבסוף, מעל הכול — למשפחתי. להורי, שאיכשהו הצליחו לתמוך בי ולעודד אותי לא משנה מה' ובד בבד לא ללחוץ עלי כאשר חיפשתי את עצמי. אני יודע שאתם גאים בי, ואני מקווה שגם ברור לכם שהדרך שבה גידלתם אותי — בהבנה מתי לתת לי מרחב לטעות — הייתה חיונית להצלחה הזו. לאחותי הקטנה והחכמה ממני לאין שיעור, שחר, שדיברה איתי כשווה גם כשאני למדתי מתמטיקה בתיכון והיא עסקה בפיזיקה של מצב מוצק. אני גם מצטער שביקשתי ממך לפתור משוואות בלי פתרון ממשי כשהיית בת שמונה, בלי לגלות לך. ולבסוף, לרעייתי המופלאה מאיה. את העוגן שלי, נקודת המשען שלי, ואהבת חיי. הבאת לעולם את ילדינו הנפלאים, ודחית את הקריירה שלך כדי שאוכל לרדוף אחר החלום הזה של קריירה כחוקר. יש לי הוכחה נהדרת לכך שאת מושלמת עבורי, אבל העמוד קצר מלהכיל אותה. כל מה שעשיתי ואעשה אפשרי בזכותך, ושווה כקליפת השום אם לא תהיי שם לצידי.

מסע לפני יעד.

אני מודה לטכניון על התמיכה הכספית הנדיבה בהשתלמותי.

המחקר בוצע בהנחייתם של פרופסור לאוניד מירקין ופרופסור דניאל זלזו, בפקולטה להנדסת מכונות.

חלק מן התוצאות בחיבור זה פורסמו כמאמרים מאת המחבר ושותפיו למחקר בכנסים ובכתבי-עת במהלך תקופת מחקר הדוקטורט של המחבר, אשר גרסאותיהם העדכניות ביותר הינן:

G. Barkai, L. Mirkin, and D. Zelazo, "On sampled-data consensus: Divide and concur," *IEEE Control Systems Letters*, vol. 6, pp. 343–348, 2022.

G. Barkai, L. Mirkin, and D. Zelazo, "On the internal stability of diffusively coupled multi-agent systems and the dangers of cancel culture," *Automatica*, vol. 155, p. 111 158, 2023.

G. Barkai, L. Mirkin, and D. Zelazo, "An emulation approach to sampled-data synchronization," in *Proc. 62nd IEEE Conf. Decision and Control.*, 2023, pp. 6449–6454.

G. Barkai, L. Mirkin, and D. Zelazo, "An emulation approach to output-feedback sampled-data synchronization," in *Proc. 22nd European Control Conf.*, 2024.

G. Barkai, L. Mirkin, and D. Zelazo, "Asynchronous sampled-data synchronization with small communications delays," in *Proc. 63rd IEEE Conf. Decision and Control.*, 2024.

המחקר שבבסיס חיבור זה נערך כולו ביושר, ולפי אמות המידה האתיות המקובלות באקדמיה. בפרט אמורים הדברים בפעילויות איסוף הנתונים, עיבודם והצגתם, התייחסות והשוואה למחקרים קודמים וכו', ככל שהיוו חלק מן המחקר. כמו כן, הדיווח על המחקר ותוצאותיו בחיבור זה הוא מלא וישר, לפי אותן אמות מידה.

ניצול יחסי הגומלין בין אילוצים במרחב ובזמן בבקרת מערכות מרובות סוכנים

חיבור על מחקר

לשם מילוי חלקי של הדרישות לקבלת התואר
דוקטור לפילוסופיה

גל ברקאי

הוגש לסנט הטכניון — מכון טכנולוגי לישראל
אב התשפ"ה חיפה אוגוסט 2025

ניצול יחסי הגומלין בין אילוצים במרחב ובזמן בבקרת מערכות מרובות סוכנים

גל ברקאי