



NUI MAYNOOTH

Ollscoil na hÉireann Má Nuad

Topics in Cooperative Control

A dissertation submitted for the degree of
Doctor of Philosophy by

Florian Knorn

Maynooth, June 2011

Adviser: Prof. Robert N. Shorten

Examiners: Prof. Abraham Berman (external)
Prof. Seán McLoone (internal)

Hamilton Institute
National University of Ireland, Maynooth
Ollscoil na hÉireann, Má Nuad
Co. Kildare, Ireland

Für Steffi, Sophie

und

meiner Mutter, meinem Vater und dem Rest meiner Familie

Table of contents

Abstract	ix
Preface	xi
Acknowledgements	xi
Publications during doctorate	xii
1 Introduction	1
1.1 Overview and structure	1
1.2 Motivating examples	2
1.2.1 Stability of a wireless network power control algorithm	2
1.2.2 Emissions control in traffic networks	3
1.2.3 Topology control in wireless sensor networks	4
2 Literature Review	5
2.1 Introduction	5
2.2 Switched Systems and Positive Systems	6
2.2.1 Switched systems	7
2.2.2 Positive Systems	13
2.2.3 Switched positive systems	15
2.3 Large-Scale Systems and Decentralised Control	16
2.3.1 Decomposition	18
2.3.2 Aggregation	20
2.3.3 Basic concepts of Decentralised Control	24
2.4 Cooperation and consensus	28
2.4.1 History	28
2.4.2 Networked dynamic systems	29
3 Switching	33
3.1 Introduction	33
3.2 Preliminaries	34
3.2.1 Notation	34
3.2.2 Definitions	35
3.3 State-dependent switching case	36
3.3.1 Main result	36

3.3.2	Numerical test based on a linear program	38
3.4	Arbitrary switching case	39
3.4.1	Main result	39
3.4.2	Remarks	43
3.4.3	Insights from Hurwitz condition	43
3.5	Discrete-time switched positive systems	45
3.6	Examples of usage	45
3.6.1	Numerical example	45
3.6.2	Switched positive systems with multiplicative noise	46
3.6.3	Robustness of switched positive systems with channel dependent multiplicative noise	46
3.7	Conclusion	47
4	Switching and Feedback	51
4.1	Introduction	51
4.2	Preliminaries	53
4.2.1	Overall setting and problem statement	53
4.2.2	Notation	55
4.2.3	Growth conditions	56
4.2.4	Feasibility and existence of unique solution	56
4.3	Algorithm 1: Complete knowledge of system	57
4.3.1	Main result	59
4.3.2	Simulations	61
4.4	Algorithm 2: System only partially known	62
4.4.1	Main result	62
4.4.2	Simulations of Algorithm 2	66
4.4.3	Extension when access to the global term is limited	67
4.5	Algorithm 3: Dynamics and controllers	69
4.5.1	Main result	70
4.5.2	Simulations of Algorithm 3	72
4.6	Extension to asynchronous state updates	75
4.6.1	Asynchronous version of Algorithm 1	75
4.6.2	Asynchronous versions of Algorithms 2 and 3	76
4.6.3	Simulations of the extensions of Algorithm 2	77
4.7	Conclusion	78
4.A	Chapter appendix	80
4.A.1	An expression for the global term	80
4.A.2	Global- and utility functions used in our simulations	81
5	Switching, Feedback and Estimation	83
5.1	Introduction	83
5.2	Preliminaries	86
5.2.1	Basic idea	86
5.2.2	General setting	88

5.3	Decentralised estimation of the second eigenvalue	89
5.3.1	Estimate_real()	90
5.3.2	RLS_real()	92
5.3.3	Estimate_complex()	92
5.3.4	Remarks	93
5.3.5	Demonstration of estimation	95
5.4	Decentralised connectivity control	97
5.4.1	Consensus with input	97
5.4.2	Application to decentralised connectivity control	99
5.4.3	Conditions for convergence of the decentralised control law	100
5.5	Simulation results	107
5.5.1	Example 1: Controller stability bounds	108
5.5.2	Example 2: Combining Control and Estimation	110
5.5.3	Further Examples of control	110
5.5.4	Validation of control results	112
5.5.5	Examples of other control objectives	113
5.6	Conclusion	114
6	Applications	117
6.1	Stability of the Foschini-Miljanic algorithm	117
6.1.1	Introduction	117
6.1.2	Mathematical preliminaries	119
6.1.3	Wireless communications	121
6.1.4	Main results	123
6.1.5	Example	124
6.2	Emissions control in a fleet of Hybrid Vehicles	126
6.2.1	Introduction	126
6.2.2	Background	126
6.2.3	Controlling emissions, maximising driving distance	128
6.2.4	Implementation	129
6.2.5	Simulations	131
6.3	Real-world implementation of cooperative control	135
6.3.1	Introduction	135
6.3.2	Overall setting	136
6.3.3	Practical issues	140
6.3.4	Results from experiments	141
7	Conclusion	145
7.1	Summary	145
7.2	Future directions	147
	Notation	153
	Bibliography	155

Abstract

The main themes of this thesis are networked dynamic systems and related cooperative control problems. We shall contribute a number of technical results to the stability theory of switched positive systems, and present a new cooperative control paradigm that leads to several cooperative control schemes which allow multi-agent systems to achieve a common goal while, at the same time, satisfying certain local constraints. In this context, we also discuss a number of practical applications for our results.

On a very abstract level, we first investigate the stability of an unforced dynamic system or network that switches between different configurations. Next, a control input is included to regulate the aggregate behaviour of the network. Lastly, looking at a particular instance of this problem setting, an estimation component is added to the mix.

To be more specific, we first derive a number of necessary and sufficient, easily verifiable conditions for the existence of common co-positive linear Lyapunov functions for switched positive linear systems. This is particularly useful given the classic result that, roughly, existence of such functions is sufficient for exponential stability of the switched system under arbitrary switching. Such switched systems may represent a networked dynamic system that switches between different configurations.

Next, we develop several cooperative control schemes for networked, dynamic multi-agent systems. Several decentralised algorithms are devised that allow the network to achieve what may be called implicit, constrained consensus: Constrained in the sense that the aggregate behaviour of the network (assumed to be a function of the totality of its states) should assume a prescribed value; implicit in the sense that the consensus is not to be reached on the states directly, but on values that are a function of the states. This can be used to assure inter-agent fairness in some sense, which makes this result relevant to a large class of real-world problems. Initially, three algorithms will be given that work in a variety of settings, including non-linear and uncertain settings, time-changing and asymmetric network topologies, as well as asynchronous state updates. For these results, the general assumption is that the aggregate behaviour of the network is made accessible to each node so that it can be incorporated into the control algorithm.

Then, a somewhat more specific application is addressed, namely (algebraic) connectivity control in wireless networks. This is a setting where the aggregate behaviour (the network's connectivity level, roughly an algebraic measure of how well information can flow through the network) has to be estimated first before it can be regulated. To that end, a fully decentralised scheme is developed that allows the connectivity level to be estimated locally in each node. This estimate is then used to inform a decentralised scheme to adjust the nodes' interconnections in order to drive the network to the desired connectivity level.

Finally, three further real-world applications are discussed that rely on the results presented in this thesis.

Preface

*God is love.
Whoever lives in love lives in God,
and God in them.*

1 John 4:16

Acknowledgements

Due to the “longevity” and somewhat permanent nature of a printed Ph.D. thesis, this acknowledgement section is the perfect place to express my deepest gratitude to all the wonderful people without whom this thesis would never have been possible.

First and foremost, I could never overstate my gratitude towards Prof. R. Shorten, for his support, ideas, guidance and encouragement, both in academic and personal matters. In my eyes, you are the best *Doktorvater* on the planet, and a person I will always look up to, in all aspects of life.

I am also greatly indebted to my other co-authors and collaborators, Prof. M. Corless, Dr. O. Mason, Dr. R. Stanojević, Mr. J. Kinsella, Prof. D. Leith, Dr. T. Charalambous, and Dr. A. Zappavigna. It was only through the joint work with them that this thesis came together.

It was an honour to have spent the last four years here at the brilliant *Hamilton Institute*, probably the best research environment I came to know so far. It is such a high-profile yet friendly, open minded environment and its almost natural open door policy is so incredibly helpful for any kind of academic work. In particular, I am most thankful to Prof. S. Kirkland and Prof. R. Middleton for their generous and patient help and advice, but there are also many more colleagues that have made important contributions in some form or another to this thesis.

On a more private level, I would like to thank all the amazing people and fellow students — far too numerous to name them all — that I have met over the years here in Maynooth. You have made it all too easy to maintain a healthy (probably too healthy) balance between “fun” and “work”, even if the line between both was never too well defined in the first place. This also includes the lovely *Maynooth Community Church*.

Finally, I am forever indebted to my parents and family for their continued love and support over the past almost 30 years. Words are inadequate to express the joy and love I feel for my daughter Sophie and my wife. Steffi, thank you for being the amazing person you are, for your love and never ending support and patience with me and my many quirks. Few men can function properly (if at all) without the right partner on their side, and even fewer can consider themselves as lucky as I am.

Publications during doctorate

The following journal publications were prepared in the course of this doctorate.

Florian Knorn, Oliver Mason, and Robert N. Shorten. On linear co-positive Lyapunov functions for sets of linear positive systems. *Automatica*, 45(8):1943–1947, Aug. 2009a. DOI : 10.1016/j.automatica.2009.04.013

Florian Knorn, Rade Stanojević, Martin J. Corless, and Robert N. Shorten. A framework for decentralised feedback connectivity control with application to sensor networks. *International Journal of Control*, 82(11):2095–2114, Nov. 2009d. DOI : 10.1080/00207170902912056

Florian Knorn, Martin J. Corless, and Robert N. Shorten. Results in cooperative control and implicit consensus. *International Journal of Control*, 84(3):476–495, Mar. 2011a. DOI : 10.1080/00207179.2011.561442

Annalisa Zappavigna, Themistoklis Charalambous, and Florian Knorn. Unconditional stability of the Foschini-Miljanic algorithm. To appear in *Automatica*, Mar. 2011. <http://goo.gl/o8rKm>

Additionally, the following conference publications provide further applications and extensions to the work reported in the journal papers.

Florian Knorn and Douglas J. Leith. Adaptive Kalman filtering for anomaly detection in software appliances. In *INFOCOM-ANM'08: Proceedings of the 1st IEEE Workshop on Automated Network Management*, pages 1–6, Phoenix, AZ, USA, Apr. 2008. DOI : 10.1109/INFOCOM.2008.4544581

Florian Knorn, Oliver Mason, and Robert N. Shorten. Applications of linear co-positive Lyapunov functions for switched linear positive systems. In *POSTA'09*, pages 331–338. DOI : 10.1007/978-3-642-02894-6_32

Florian Knorn, Rade Stanojević, Martin Corless, and Robert N. Shorten. A problem in positive systems stability arising in topology control. In *POSTA'09*, pages 339–347. DOI : 10.1007/978-3-642-02894-6_33

Florian Knorn, Martin J. Corless, and Robert N. Shorten. Results and application of cooperative control using implicit and constrained consensus. Submitted to *CDC-ECC '11: Joint 50th IEEE Conference on Decision and Control and the 2011 European Control Conference*, Mar. 2011b

Introduction

In this first chapter we briefly establish the context for the work developed in this thesis and give an overview of its structure. We also provide a number of motivating examples to set the stage for some of the main results derived in subsequent chapters.

Chapter contents

- 1.1 Overview and structure
 - 1.2 Motivating examples
-

1.1 Overview and structure

With man’s innate desire and drive to expand, conquer, progress, improve and optimise, the technological tools created in the process also never cease to grow. This growth may happen both in terms of sheer size and in complexity. In the past century in particular, two new key ingredients were added to the development: miniaturisation and communication. On the one hand, systems increased in functionality but at the same time decreased in size (computers are just one of the many examples for this trend). On the other, systems also became more and more connected thanks to more efficient, faster, capable and reliable communication means (think of the banking and stock trading systems, governments, or indeed the Internet). Both trends combined lead to large systems composed of many “small” but interconnected components rather than of one large, monolithic block. The advantages of that are evident — due to the distributed nature of the system it would be more robust to disturbances than a centralised system with its single point of failure, and it could also better adapt to locally changing environments. However, it is also clear that many individuals need to “cooperate” to achieve a common task.

Cooperation is typically defined as the process of working together toward the same end, and cooperation is clearly paramount between the elements in such networked settings as lack thereof would certainly not lead to the desired common goal. This may explain the growing interest in recent decades in enabling large systems to exhibit such needed cooperative behaviour.

In this context, the main theme of this Ph.D. thesis is networked dynamic systems related to which three problems are studied. First, we will be looking at switched positive systems which, in some sense, may be interpreted as networks of scalar systems that switch between different topologies. Here, we shall make several contributions to the relatively young research area of *switched positive systems* by providing a number of necessary and sufficient stability conditions for switched positive linear systems. Second, we will investigate networks of systems with switching topologies that have some form of global control input in order to regulate the network's aggregate behaviour. In particular, we will derive a number of decentralised algorithms that enable multi-agent systems to cooperatively achieve a common goal while additionally fulfilling certain localised constraints. Third, an extension of this problem is studied where an estimation component needs to be added to the network in order to first estimate the aggregate network behaviour before it can be cooperatively regulated.

Stability of switched positive systems can be seen as a sub-problem of general systems theory and switched systems in particular. Cooperative control, in turn, is a relatively novel concept that is closely related to several “traditional” control approaches, in particular *large-scale systems*, *decentralised control*, and more recently *multi-agent systems*. These relevant fields of research will first be discussed in detail in the literature review in the next chapter. We shall then present our main results in Chapters 3, 4 and 5. The applications chapter, Chapter 6, will complement the theoretical contributions by providing several applications where those results could be of use. Finally, we draw some conclusions from our work and suggest future directions.

Before moving on to the literature review, let us give a few motivating examples for the work carried out in this thesis.

1.2 Motivating examples

1.2.1 Stability of a wireless network power control algorithm

Various radio communication technologies rely on the so-called *Code Division Multiple Access (CDMA)* method to select and use radio channels for broadcast and reception, Schulze and Lüders (2005).¹ It is based on the general idea in data communications that several transceivers should simultaneously utilise a single communication channel to transmit and receive information in order to maximise spacial and temporal use of the spectrum. This concept is known as the *multiple access* concept. However, with multiple sources broadcasting at the same time, the broadcast power needs to be carefully adjusted and controlled as each transmission between one pair of nodes interferes with the communication between other nearby nodes in the network. Thus, a compromise needs

¹ To name two of the most high-profile applications, the *Global Positioning System (GPS)* as well as mobile phone standards *cdmaOne* and *CDMA2000* are based on this method.

to be found for each communication pair — on the one hand power output should be minimised to limit the interference with other nodes' communications, but on the other it must be large enough to guarantee a stable communication link (i.e. the signal needs to be by a factor larger than the local interference level in order to be correctly picked up by the receiver).

A seminal power control algorithm for wireless networks is the *Foschini-Miljanic (FM)* algorithm, Foschini and Miljanic (1993), which works in a fully decentralised way. It adjusts and minimises each node's power output all while observing certain quality of service requirements. This algorithm has been proved to be stable to various kinds of perturbations and adverse conditions. However, only recently has it been shown that it is stable even in the presence of time-varying time-delays.

At the heart of this result (presented in Section 6.1) is a delay-independent stability property of switched positive systems that ultimately relies on the existence of certain types of Lyapunov functions. Necessary and sufficient conditions and checks for their existence, as derived in the third chapter, are thus relevant to a large class of real world problems.

1.2.2 Emissions control in traffic networks

A second example would be a network of cars driving around in a city, where the city council is trying to implement some form of CO₂ emissions control. Assume the overall objective would be that the aggregate emissions of all cars participating in the scheme do not exceed a prescribed level. Fairness dictates that no car should be allowed to pollute more than others, thus the cars should adjust their behaviour so that they all produce the same CO₂ emissions (in other words, reach a consensus on the emissions). But assuming that the emissions are a direct function of the cars' speed (and that different cars have different efficiency levels, depending on their weight, engine, *etc.*) some cars may be able to drive faster than others for a given level of permissible emissions.

In order to implement the emissions control scheme the council may place a number of monitoring units around the city to measure the overall emissions level and broadcast that (global) information to all the cars in the network, along with the value of the desired or allowable emission level. Clearly, the cars need to *cooperate* in order to achieve the desired emission level since the city-wide (traffic related) emissions are just the sum of the individual contributions.

To make such cooperation possible we assume that the cars are able to broadcast their own emission level to vehicles in their vicinity. The so-established communication network can then be used to reach an "agreement" among the cars on a common emission level. Additionally, incorporating the information from the city-wide emissions broadcast, the cars should now be able to conjointly adjust their speed so that the resulting emissions match those of other cars in the network, and also so that the overall emissions produced throughout the city reach the admissible level.

Highlighting some of the particularities of this setting we note that the topology of the resulting communication network would be constantly changing as the cars drive around and move in and out of range from each other; the communication network will not necessarily be symmetric — some cars may not be able to broadcast as far as others, or some of the transmissions may be lost; and the dependence of emissions on the driving speed is usually non-linear.

Problems of this type will be considered in Chapter 3 and a real-life application along these lines is discussed in the fifth chapter.

1.2.3 Topology control in wireless sensor networks

Lastly, consider a different type of wireless network, this time one that interconnects small sensor units or *motes*. Assume that a large number of such battery powered motes are dropped roughly uniformly distributed over a defined area. The (usually identically built) motes would be equipped with a battery, a transceiver, one or more sensors and some kind of processing unit. Networks of this type are very common and widely used, Akyıldız *et al.* (2002). The radio in the motes is used to form a network between all the nodes, and one objective here could be to adjust the broadcast power of their radios so that this network reaches a prescribed level of (algebraic) connectivity. However, one may additionally require that all nodes should last equally long in terms of battery power. The first objective would be important for certain types of algorithms whose rate of convergence depends on the level of connectedness of the graph they evolve on, and the second objective guarantees maximum life-time of the network without node failures (due to power shortage).

Clearly, the power used by the radio directly influences the time-to-live (TTL) of a node. However, the overall power consumption may vary among nodes depending on their individual workload, and the batteries may also have slightly varying initial charges. Assuming that the radio is the biggest power consumer in each mote, they will be able to influence their TTL by varying the power setting of their radios. But now, depending on the power used, each node can broadcast information to more or fewer nodes in its vicinity. As different nodes will use different power settings, the resulting topology of the communication graph will generally be asymmetric, and changing over time. In this setting, we would like to find a decentralised algorithm that adjusts the node's power setting so that on the one hand all nodes eventually have equal TTLs, but on the other hand also guaranteeing a certain guaranteed level of connectedness of the resulting communication network. This means that again the objective is a combination of local and global constraints, with additionally an identification component involved.

This problem setting will later be studied in detail in the fourth chapter.

With these motivating examples in mind, let us now move on to the literature review.

Literature Review

This second chapter reviews related work reported in the literature and puts the thesis into the context of existing research. In particular, we discuss the areas of switched positive systems, large-scale systems, decentralised control, and cooperation in networked multi-agent systems.

Chapter contents

- 2.1 Introduction
 - 2.2 Switched Systems and Positive Systems
 - 2.3 Large-Scale Systems and Decentralised Control
 - 2.4 Cooperation and consensus
-

2.1 Introduction

As we mentioned in the introduction, three areas of research are particularly relevant to this thesis. Before going into the details, let us briefly state their key objectives:

- *Switched positive systems* focus on systems whose overall dynamics switch over time between a number of distinct constituent behaviours or dynamics, and whose states are only defined in and thus confined to the non-negative orthant.
- *Large-scale systems and decentralised control* theory aims at developing a theoretical framework particularly suited for the analysis and control of large systems, and typically attempts to find or design constituent system dynamics with the property that, when connected together, the resulting closed-loop system will be stable. In particular, the implemented control laws should be decentralised, so that there is no single, centralised entity that regulates the system.
- *Networked multi-agent systems*, and in particular consensus and cooperation therein: Attempts are made to develop consensus algorithms or protocols that pose an interaction rule specifying the information exchange between agents and usage of communicated information to update the agents' states so that the system reaches an "agreement" of sorts, and that the system achieves a certain goal "cooperatively".

Leaving the first research area aside for a moment, the last two fields generally deal with systems that are not “monolithically” large, but large in the sense that they are composed of a great number of interconnected, more granular subsystems that have both some amount of “self-interaction” as well as some interaction with neighbouring subsystems, but not *every* other subsystem. Put differently, the graph describing interactions among subsystems is assumed to be sparsely connected.

Such a setting naturally lends itself to be treated by decomposing the system into its “parts” rather than investigating everything as a whole. Similarly, with our growing desire for even larger, even more complex systems, it may not be attractive to use a single large, central computer to control the system — be it for economic, reliability or pure technical feasibility reasons. This becomes evident by considering the many, diverse real world applications such as power networks, communications networks, large chemical plants and oil refineries, ecological systems, traffic networks, economic and financial systems, or finite element discretisations, just to name a few.

In the following literature review, by no means encyclopedic in nature, we begin by discussing switched positive systems as they are particularly relevant to the third chapter of this thesis. We then approach the area of large-scale systems and decentralised control, reviewing some of the most common results used to analyse and stabilise large systems. Finally, we visit the more recent notion of achieving an aggregate behaviour “cooperatively” as well as the idea of consensus and agreement in dynamical systems. Cooperative control may be considered as a separate field from the more traditional decentralised control theory in that it typically deals with even larger, but more homogeneous systems formed by a network of interconnected, but all in all similar entities.

2.2 Switched Systems and Positive Systems

The class of *switched positive systems* refers to dynamical systems that have two important qualities: They are *positive*, which means their states are only defined for non-negative values and that they remain in the closed positive orthant throughout time. Additionally, they are of *switched* nature, that is their evolution is not governed by a single but several, different dynamic system formulations between which the system switches over time, and which represent different, distinct system behaviours.

Both types of systems play a crucial role in many real world applications: For many physical variables only positive values are meaningful (for instance, masses, liquid concentrations, temperatures, volumes, *etc.*; but also quantities of objects or probabilities), and while switched behaviour can be observed in a number of natural sciences, it is most prevalent in man-made applications (for instance, consider robotic systems switching between different operating modes, transmission boxes in vehicles, networked systems with changing communication topologies, event-driven systems, *etc.*).

In the following, we begin by giving an overview of switched systems. This is followed by a discussion of positive systems where restriction of the state to the closed positive orthant allows for much more comprehensive stability results than are available in the general case. In the last subsection, we finally present a number of results from the relatively young field of switched positive systems that unites both fields.

2.2.1 Switched systems

It is generally understood that a switched system consists of a number of dynamic systems called *constituent systems*, *subsystems* or *modes* (representing different “behaviours”) together with a *switching rule* or *switching signal* that orchestrates the switching among them. Switched systems are thus closely related to and can be seen as a sub-class of *hybrid* systems since they constitute a mix of both dynamic elements (the state evolution governed by differential or difference equations) and discrete time elements (the piecewise constant switching function).

A great deal of attention has been given to switched systems for a number of reasons. First, this framework allows a much more natural modelling of many real-world phenomena which exhibit switching between different, distinct behaviours (common in biological networks for instance, de Jong *et al.*, 2004). Then, it is also of particular use in the context of intelligent control systems which attempt to improve overall performance by switching between different, tailor-made controllers that are more appropriate for different (local) operating regimes, Ge and Sun (2005). Furthermore, switching between even the simplest, linear systems can produce very complex behaviours including chaos and multiple limit cycles, Yang and Chen (2008). Another interesting fact is that even if given two planar, linear, time-invariant systems that are exponentially stable, stability under arbitrary switching among the two vector fields associated with these LTI systems is not, in general, guaranteed to be stable. In other words, it may well be possible to construct a switching sequence that results in an unstable overall behaviour, Liberzon and Morse (1999). To illustrate this point, an example of a destabilising switching sequence applied to a system consisting of two (individually exponentially stable) LTI systems is given in Figure 2.1 on the following page.

While work on the more general problem of differential equations with time varying parameters has been ongoing since the early 1900s (Perron, 1930; Maizel’, 1954; Sell, 1963; Conti, 1967; Coppel, 1978), a new body of literature focusing in particular on switched systems (where system parameters vary *abruptly* with time) has been growing since the 1990s. For a more in-depth treatment of the wealth of results (the vast majority of which only applies to linear systems) refer to the books by Liberzon (2003); Murray-Smith and Shorten (2003); Li *et al.* (2005); Ge and Sun (2005); Boukas (2006); Mahmoud (2010) or the survey articles by Liberzon and Morse (1999); Michel (1999); Decarlo *et al.* (2000); Hespanha (2004a); Lin and Antsaklis (2009) and in particular Shorten *et al.* (2007) on




Figure 2.1: Trajectory resulting from a destabilising switching sequence in a switched system of two Hurwitz stable second order LTI systems.

which parts of this section are based. Most of the research can typically be attributed to two fundamental questions: Is a switched system stable under arbitrary switching, or (if not) is it stable when certain restrictions are placed on the switching signal?

In the following, we shall discuss some of the literature that dealt with these questions.

Stability under arbitrary switching

Since Lyapunov theory plays a key role in the stability analysis of dynamic systems, it should come as no surprise that most of the results concerning switched systems also rely on such ideas. It is easy to see that if a Lyapunov function exists for a switched system under arbitrary switching — which also includes constant “switching” signals — then this function must be valid for each constituent system in isolation as well. In other words, such function would have to be a *common* Lyapunov function for all subsystems. Indeed, a classic result for linear (continuous-time) switched systems shows that existence of a common Lyapunov function is equivalent to *uniform exponential stability* or *absolute stability*, see Molchanov and Pyatnitskii (1989); Dayawansa and Martin (1999); Liberzon and Morse (1999); Fornasini and Valcher (2011) for more details and the precise definition of these terms. A similar result for discrete time systems can be derived from Brayton and Tong (1979); Barabanov (1988). In that context, most of the literature appears to focus on finding common *quadratic* Lyapunov functions, but other types such as linear or piecewise quadratic / linear have also received attention.

Converse Lyapunov theorems Loosely speaking, these results guarantee existence of Lyapunov functions given stability. They have been established for different types of switched systems, including linear systems (Molchanov and Pyatnitskii, 1989; Blanchini, 1995), non-

linear systems (Dayawansa and Martin, 1999; Mancilla-Aguilar and García, 2000), uncertain systems (Lin and Antsaklis, 2005a), systems with dwell-time¹ (Wirth, 2005b), or input-to-state stable systems (Mancilla-Aguilar and García, 2001). But while it is useful to know such correspondence between stability and Lyapunov function existence, finding tests that *guarantee* the existence of a common Lyapunov function (and thus stability) is probably most relevant for practical applications. In the linear case, this basically means: What conditions must the system matrices of the constituent systems fulfil in order for the overall system to be stable under arbitrary switching? Such existence questions can be approached numerically and algebraically.

Numerical tests The advantage of focusing on common quadratic Lyapunov functions is that their existence problem can be formulated as a set of *linear matrix inequalities*. If the resulting system of inequalities is *feasible*, that is if a solution exists, then the switched system will be exponentially stable, Boyd *et al.* (1994); Ghaoui and Niculescu (2000); Liberzon and Tempo (2004); Ibrir (2008). A different technique involving periodic switching signals was derived in Margaliot and Yfoulis (2006). Techniques for the systematic construction of common piecewise linear Lyapunov functions (which were considered as early as the 1960s in the context of *Lur'e systems*, Rosenbrock, 1963; Weissenberger, 1969) and common polyhedral Lyapunov functions have been studied in Brayton and Tong (1979, 1980); Barabanov (1989); Polański (1995, 1997); Johansson and Rantzer (1998); Polański (2000); Yfoulis and Shorten (2004); Christophersen and Morari (2007). Unfortunately, all these approaches only provide sufficient conditions for stability, and even if they can answer the stability question (provided the original problem is not too large), they usually provide little insight as to *why* a system is stable or not.

Algebraic conditions These tend to provide more meaningful answers to the stability question and shine more light on the dynamical properties of switched systems. However, the general problem of proving common Lyapunov function existence for linear systems is yet to be solved. There are nonetheless a number of useful results for specific types of linear systems (all, of course, under the assumption that each of the constituent systems is stable). For instance, if the system matrices are symmetric or normal, then the resulting system will be stable under arbitrary switching, Zhai and Lin (2004); Zhai *et al.* (2006). Triangular systems also always have a common (quadratic) Lyapunov function Mori *et al.* (1997); Shorten and Narendra (1998). In fact, for such systems, exponential stability of the constituent systems is equivalent to uniform exponential stability under arbitrary switching. This is particularly useful since even certain non-triangular systems can be brought into triangular form: For instance, it is well known that if system matrices commute with each other, then there exists a unitary matrix which can be used to transform each system matrix into upper triangular form, Horn and Johnson (1985); Narendra and Balakrishnan

¹ As we shall see later, these are systems which cannot switch arbitrarily fast, but have a uniform upper bound on the switching rate.

(1994). If the system matrices do not commute, but if the *Lie-Algebra* generated by them is solvable, then it is Lie's theorem (Humphreys, 1972) that guarantees that the system is simultaneously triangularisable.

Extensions to these ideas have been reported in Shorten and Cairbre (2001a,b, 2002); Solmaz *et al.* (2007), attempting to relax the somewhat restricting requirement of *simultaneous* triangularisability to pairwise triangularisability. Further necessary and sufficient stability results for special classes of systems concern pairs of: planar systems (Shorten and Narendra, 2000, 2002), third-order systems (King and Shorten, 2004, 2006), and systems with rank one difference (Shorten and Narendra, 2003; King and Nathanson, 2006). A necessary and sufficient condition for the *robust* existence of a common quadratic Lyapunov function (hence implying exponential stability) with respect to certain types of perturbations is discussed in Hinrichsen and Pritchard (1989); Shorten *et al.* (2007) where the concept of *stability radii* is used. Sufficient conditions based on Lyapunov operators were developed in Ooba and Funahashi (1997a,b,c, 1999). Lastly, necessary and sufficient asymptotic stability conditions for general switched linear systems were reported for the discrete-time case in Lin and Antsaklis (2005b); Bhaya and das Mota Chagas (1994) and for the continuous-time case in Bhaya and das Mota Chagas (1994); Lin and Antsaklis (2009).

While all these results are promising they are generally hard or computationally expensive to check for systems of larger dimensions and/or with many constituent systems. Also, not all applications require stability under *arbitrary* switching, as we shall see next.

Stability under restricted switching

Many real-world system cannot switch instantaneously or have a natural upper bound on the switching rate (consider gear changes in a car for instance); in other cases the system may not be able to switch from any one mode to any other mode, but must adhere to a prescribed switching sequence/order (for example, it would be rather unlikely that an automatic gearbox would change directly from fifth to first gear). Given such *a priori* knowledge of time domain or state space restrictions on the switching signal, it is possible to find less conservative stability results. Also, another interesting question concerns whether it is possible to restrict switching to result in a stable overall behaviour for systems that contain a number of unstable modes.

Slow switching On an abstract level, it is easy to understand how restrictions on the switching rate can contribute to stability: Assume a switched system is composed of stable subsystems with the property that, when a subsystem is activated, it exhibits a short intermittent increase in energy. Since the subsystems are stable, they would absorb the initial energy increase quickly. But if one switches “too quickly” between the systems, this increase may build up quicker than it can be absorbed — with the result that the switched system would not be stable. If, however, the switching rate was restricted and

each subsystem is given enough time to absorb the temporary increase, then the switched system would be stable. Recall Figure 2.1 on page 8 which showed a somewhat “fast” switching sequence — if the same system is switched just a little bit slower, the solution will actually converge, see Figure 2.2 below.




Figure 2.2: Trajectory of the same system and same initial condition as used in Figure 2.1, but this time using a slower switching sequence.

Such ideas of constraining the switching rate have been studied extensively over the past decades, initially in the context of systems with slowly varying parameters, see for instance Desoer (1969); Ilchmann *et al.* (1987); Guo and Rugh (1995). In the switched systems literature, the term *dwell-time* captures this concept, Hespanha (2004b); Hespanha and Morse (1999); Morse (1996); Zhai *et al.* (2001). It defines the (uniform) lower bound on the time intervals between consecutive switching instants. A classical result then confirms the intuition, Morse (1996): If the dwell-time is sufficiently large, a switched system based on Hurwitz stable subsystems is asymptotically stable for any switching system respecting the dwell-time constraint. However, it is also intuitive that, occasionally, the dwell-time constraint may be violated without compromising stability, provided this does not happen too frequently. This led to the introduction of the more forgiving *average dwell-time* concept (Hespanha and Morse, 1999), for which a similar result exists — but since the required average dwell-time may be smaller than the fixed one it will allow for a broader class of switching signals. Similar concepts for the discrete time case exist as well, Zhai *et al.* (2002). Unfortunately, it appears that most of the existing results only provide rather conservative bounds on the dwell-time — tight conditions on the truly required minimum dwell-time are still a topic of research, Shorten *et al.* (2007). Converse Lyapunov theorems for the dwell-time case are reported in Wirth (2005a); De Santis *et al.* (2004); Pola *et al.* (2004).

Apart from defining a minimum time between switches, it may also be required, in cases, to introduce an *upper bound* on the time the system is allowed to stay in a mode. Switching signals obeying such upper bound then may allow a switched system with unstable modes to be overall stable — as the system is not allowed to spend too much time in the unstable mode. Work investigating such situations includes Lin *et al.* (2003); Zhai *et al.* (2001, 2002); Yedavalli and Sparks (2001).

State-dependent switching As mentioned earlier, the switching may also be constrained by rules that depend on the state vector of the system. This can come in two flavours — either the switching is directly a function of the state value (switching is entirely dictated by the state vector alone), or it is arbitrary but subject to certain *constraints* that depend on the state. The latter (more general) set-up is considered in the next chapter. The former, more common set-up assumes that the state space is partitioned *a priori* into closed (but possibly unbounded) regions or “cells” whose interiors are pairwise disjoint but whose union covers the entire state space (such regions are usually denoted by Ω in the literature), and each of these Ω -regions has a particular subsystem associated to it so that the system automatically switches into that mode whenever its state enters that region. In other words, it is assumed that there are a number of hyper-surfaces that completely determine all the system’s mode switches. Such a situation is illustrated in Figure 2.3 below. Since the switching can no longer be arbitrary it may be unduly restrictive to require the existence of a *common* Lyapunov function. In fact, there may not be such a function altogether — but the system may still be asymptotically stable. A common approach is then to look for a family of (local) Lyapunov functions — usually one Lyapunov function for each region — which are then “pieced together” to create an overall function which then provides for asymptotic stability.




Figure 2.3: Illustration of the positive orthant being divided into three pairwise disjoint and conic Ω -regions. In each of these Ω_i -regions, only mode i can be activated.

This idea was applied in Johansson and Rantzer (1998) to switched affine systems by adopting a numerical technique called the *S-procedure* (Aizerman and Gantmakher, 1965; Uhlig, 1979). It allows the systematic construction of piecewise quadratic Lyapunov functions which, combined, then guarantee stability under the state-dependent switching rule. Further results for different system types using this type of Lyapunov function being mostly based on Linear Matrix Inequalities have been reported in Pettersson and Lennartson (1996, 1997); Hassibi and Boyd (1998); Johansson *et al.* (1999); Feng (2002); Pettersson and Lennartson (2002); Lee (2008); Yong *et al.* (2008). An attempt to generalise the piecewise quadratic Lyapunov function approach to more general functions of polynomial form were given in Prajna and Papachristodoulou (2003); Papachristodoulou and Prajna (2009).

Multiple Lyapunov functions This framework is another way of deriving restrictions on the switching rate (but in some formulations also the switching sequence) in order to guarantee stability. It sits somewhere in the middle between time space and state space based restrictions. As the name suggests, the rough idea is to use not just one but combine multiple non-traditional Lyapunov-like functions (usually one for each subsystem) to construct another non-traditional overall Lyapunov function — non-traditional in the sense that it may have discontinuities and may not be decreasing everywhere. This Lyapunov-like function then dictates the restrictions on the switching sequence. There are several versions of this concept, but the simplest is to constrain the switching in such a way as to guarantee that if the system is to switch into one particular mode i then the associated Lyapunov-like function must 1) be strictly decreasing at that point and 2) its value must be less than what it was when the system last left that mode. Ideas initially due to Peleties and DeCarlo (1991, 1992) motivated a number of useful results in that direction, see for instance as Branicky (1994, 1998); Ye *et al.* (1998); Geromel and Colaneri (2006a,b); Zhang *et al.* (2009a,b). Unfortunately, as in classical Lyapunov theory, it is not straightforward to choose the candidate Lyapunov functions, in particular those that would minimise the resulting times.

Before moving on to switched positive system, it should be noted that a third fundamental question relating to constraining the switching may be asked: Namely, whether it is possible (and if so and how) to construct a *stabilising* switching sequence when one or more subsystems are unstable. To limit the scope of this chapter, this shall not be dealt with here, but the interested reader is referred to the survey papers mentioned earlier (in particular Lin and Antsaklis, 2009).

2.2.2 Positive Systems

A somewhat different restriction arises when studying so-called *positive systems* (sometimes, if non-linear, they are also referred to as *monotone systems* with the assumption that the origin is stable, Rüffer *et al.*, 2010). As we mentioned earlier, these are systems

where the states only “make sense” for non-negative values — hence, the dynamics must be such that the system never leaves the closed positive orthant.

Systems with such constraints on the state space have been the subject of many recent studies in the control engineering and mathematics literature, see for instance Berman and Plemmons (1979); Berman *et al.* (1989); Johnson *et al.* (1993); Farina and Rinaldi (2000); Kaczorek (2002); Virnik (2008); Haddad *et al.* (2010) or the proceedings of the series international symposia on *Positive Systems: Theory and Applications* (POSTA’03; POSTA’06; POSTA’09). The interest in such systems is hardly surprising since they are encountered in as diverse areas as economics (Johnson, 1974; Meyn, 2008), biology (Godfrey, 1983; Jacquez and Simon, 1993, 2002; Arcak and Sontag, 2006), electronics (Benvenuti and Farina, 1996), social sciences (Bartholomew *et al.*, 1991; de Kerchove and Van Dooren, 2006), communication networks Zander (1992); Foschini and Miljanic (1993); Shorten *et al.* (2006), decentralised control Šiljak (1978), or indeed mathematics (probabilities are positive quantities) just to name a few. While both nonlinear and linear positive systems have been studied, much recent attention has focused on both time-varying (in particular switched) and time-invariant positive linear systems, and on the Metzler matrices that characterise the properties of such systems. A classical result states that a continuous-time linear time-invariant (LTI) system starting in the positive orthant will remain positive if and only if the system matrix is a Metzler matrix (that is, it has non-negative off-diagonal elements); in the discrete time case, it must be a non-negative matrix, Farina and Rinaldi (2000). Note that this property is independent of stability. For discussions on reachability and controllability in positive systems, which are out of the scope of this literature review, please refer to Caccetta and Rumchev (2000); Fornasini and Valcher (2005); Valcher and Santesso (2010).

As for general systems, any type of Lyapunov function may of course be used to study the stability properties of positive systems. For a general LTI system, the existence of a quadratic Lyapunov function (which is based on general but positive definite matrices) is necessary and sufficient asymptotic stability. In the case of positive LTI systems, however, this matrix has a simpler structure: Here the existence of a strictly positive *diagonal* matrix is necessary and sufficient for asymptotic stability, Farina and Rinaldi (2000). Furthermore, thanks to the positivity property of these systems, *co-positive* Lyapunov functions may also be employed to study stability — and as noted in Çamlıbel and Schumacher (2004), these may be less conservative as they take into account that the states only evolve in the positive orthant. For instance, with linear co-positive Lyapunov functions one searches for a strictly positive *vector*, which is even more attractive due to the even simpler structure. In the LTI case the existence of such a linear function is also equivalent to the system matrix being Hurwitz, see for instance Mason *et al.* (2009); Horn and Johnson (1991). Stability properties of positive non-linear systems were recently studied in Mason and Verwoerd (2009) and Rüffer *et al.* (2010); positive descriptor systems were considered in Virnik

(2008). But while positive linear time-invariant system are now completely understood, time-varying results appear to be scarce.

2.2.3 Switched positive systems

In this last subsection, we now turn our attention to the combination of both system types. When studying the stability of a switched system that switches between positive LTI systems, the types of Lyapunov functions mentioned above (i.e. quadratic and linear co-positive) would naturally suggest themselves. Clearly, since switched positive systems are a subclass of switched systems, all results mentioned in the previous section on general switched systems hold. However, since they do not take into account the positivity constraint on the state, attempts have been made to find less conservative stability results that are tailor suited to this system type. Let us conclude this section by reviewing a number of recent results first for the continuous time, and then the discrete time case.

Continuous-time switched positive systems

Common quadratic Lyapunov functions Necessary and sufficient conditions for existence of common quadratic Lyapunov functions for arbitrary switching between two continuous-time positive 2D systems were discussed in Gurvits *et al.* (2007). An eigenvalue condition on the product of the system matrices was derived that is equivalent to uniform asymptotic stability. Attempts to generalise these results and the general problem of finding necessary and sufficient conditions for systems with higher dimensions so far only include the 3D case in Fainshil *et al.* (2009). Common diagonal Lyapunov functions in particular were investigated in Mason and Shorten (2004). A very recent publication (Alonso and Rocha, 2010) presented general (but only sufficient) existence conditions for common quadratic Lyapunov functions in both the continuous- and discrete time case for systems of arbitrary size (both in terms of dimension and number of subsystems) using multidimensional systems analysis. Their condition relies on a certain test-matrix (which is constructed based on the constituent system matrices) being Schur-stable.

Common linear co-positive Lyapunov functions Necessary and sufficient conditions for existence of common linear co-positive Lyapunov functions were initially studied in Mason and Shorten (2007). A result was presented for switching between two constituent systems of arbitrary dimensions involving the convex hull of the system matrices being Hurwitz stable. This work was later extended in Knorn *et al.* (2009a) to arbitrarily many systems, which is the content of the next chapter of this thesis.

Common quadratic co-positive Lyapunov functions Additional equivalent conditions to the previous result were given in the Fornasini and Valcher (2010), including the fact that such common linear co-positive Lyapunov function may be used directly to construct

common quadratic co-positive Lyapunov functions (although they are of rank one). Necessary and sufficient existence conditions were also studied in Bundfuss and Dür (2009) and formulated amounting to feasibility checks of suitably defined linear inequalities, in an attempt to answer some of the general problems posed in Çamlıbel and Schumacher (2004). The work by Gurvits *et al.* (2007) also includes equivalent conditions for the existence of such functions for the 2D case with two modes studied.

A different approach involving “most unstable switching laws” was applied to the case of arbitrary dimensions in Margaliot and Branicky (2009).

Discrete-time switched positive systems

The results for linear co-positive Lyapunov functions find straightforward extensions to the discrete time case, Fornasini and Valcher (2011). In fact, in said paper it is shown that if a common linear co-positive Lyapunov function exists, then a common quadratic Lyapunov function can be found, which in turn implies that a common quadratic co-positive Lyapunov function must also exist. Switched linear co-positive Lyapunov functions were discussed in Liu (2009), where existence of such functions can be equivalently formulated as linear programming problems as well as linear matrix inequality problems.

In some sense, the types of systems encountered so far typically do not involve thousands of states and are usually of “dense” nature (in the linear case for instance it is never assumed that the system matrices are sparse). This contrasts with the next class of systems that we turn our attention to, where the opposite is assumed — “many” states, but overall “sparse” systems.

2.3 Large-Scale Systems and Decentralised Control

While research in the area of large-scale systems and control therein started in the second half of the 20th century, they continue to be of interest to this day as shown, for instance, by the ongoing series of IFAC symposia “Large Scale Systems: Theory and Applications”, (IFAC TC 5.4, 2010). Although the term “large” is of rather relative nature, we shall simply assume that it refers to systems that are large enough so that “traditional” analysis and control techniques start to reach their limits, and where a partitioned interpretation is of benefit either conceptually or computationally. Many classical approaches pre-suppose some form of “centrality” — be it centrality of *a priori* information (system model, parameters, *etc.*), centrality of measurements or centrality of actuation. However, as systems grow larger, complexity also grows rapidly: if not exponentially, it still grows faster than the system size. This implies that typically sooner rather than later centralised design, analysis or control approaches cannot be used due to the sheer size of the problem. For instance, in principle Lyapunov’s Method (Khalil, 1992; Miller and Michel, 2007) can be applied to large, complex multidimensional systems, but in practice, apart from the fact

that there is no obvious choice for a suitable Lyapunov function candidate, one would also quickly encounter computational problems.




Figure 2.4: Visual representation of a “small” protein–protein interaction network based on data by Uetz et al. (2000). In these networks, proteins form the nodes, and they are linked together if they interact in some way or other, resulting in an undirected graph. The graph shown here “only” contains about 500 nodes; other publicly available data sets contain significantly larger networks, but these are difficult to visualise.

As stated by Sezer and Šiljak (1996), one can usually identify three basic reasons why it is often necessary to move beyond classic “one-shot” approaches: i) dimensionality, ii) information structure constraints, iii) uncertainty.

“Decentralising” or decomposing the task at hand (be it modelling, analysis, or indeed control of a large-scale system), that is breaking the problem down into smaller but interconnected sub-problems, oftentimes is not only the only chance at regaining tractability, but in many cases also allows for much more meaningful insights into the problem, especially if it is of distributed nature in the first place. Presumably, these sub-problems could initially be treated independently by analysing their stability properties in isolation, to then be re-combined again (taking into account the nature of their interconnections) to give insights into the original, large system. In addition to control theoretic aspects, questions of interconnection- and communication structure and related stability issues then become relevant.

An intuitive way of creating large-scale systems is to take a large number of individual systems and interconnect them. This “bottom up” approach is typically referred to as *synthesis*. Alternatively, at “top down” approach is taken in the *decomposition-aggregation*

procedure. Here, the overall system first needs to be broken down into somewhat “independent” sub-groups (decomposition) to be then studied in isolation and finally “put together” again (aggregation) to derive properties for the overall system. Both approaches are illustrated in Figure 2.5 on the facing page.

The latter procedure is prominent in and probably originated from the economics literature, see for instance Theil (1954); Green (1964). It has been described by Simon and Ando (1961) as:²

- (i) *We can somehow classify all the variables in the economy into a small number of groups;*
- (ii) *we can study the interactions within the groups as though the interactions among groups did not exist;*
- (iii) *we can define indices representing groups and study the interaction among these indices without regard to the interactions within each group*

In the context of large-scale systems, this three-step process takes the following form, see Sandell *et al.* (1978):

- Step 1: The system is supposed to consist of interconnected subsystems. It is assumed that this decomposition or tearing has already been specified, and that a description of each subsystem and a description of the interconnection is available.*
- Step 2: It is assumed that each subsystem, when considered in isolation, is stable [or has been stabilised]. Furthermore, some quantitative measure of this stability (e. g., a lower bound on the rate of decrease of a Lyapunov function) is available.*
- Step 3: A condition is now specified in terms of this quantitative measure and some quantitative measure of the magnitude of the interconnection, and it is shown that the interconnected system is stable if the condition holds.*

Let us first review how this procedure applies to large-scale systems, starting with decomposition techniques followed by ways of aggregating the stability properties of the subsystems to derive stability of the overall system. We then discuss how such systems may be stabilised.

2.3.1 Decomposition

As we mentioned earlier, decomposition of a given large-scale system is in many cases the only option one has to analyse the system, even with ever more powerful computing equipment and increasingly sophisticated numerical tools. While work on how to best decompose complex systems started in the second half of the last century by the seminal work of Kron (1963) on electrical networks, it is reported in Himmelblau (1973) that as early as 1830 and 1843 C. F. Gauss and his student C. L. Gerling successfully solved

² Emphasis added.




Figure 2.5: Illustration of the “bottom up” synthesis and the “top down” decomposition-aggregation approach in the analysis of large-scale systems; the overall large-scale system is shown on the top, whereas the individual subsystems in isolation are shown on the bottom.

systems of equations by exploiting diagonal structures. Relevant monographs in the area include Himmelblau (1973); Sage (1977); Jamshidi (1983); Chen *et al.* (2004); Antoulas (2005).

Two basic approaches can be distinguished: Tearing along *physical* or *mathematical* lines. In the former case, the system is broken down according to physical considerations and the subsystems have a physical coherence usually representing distinct, natural structures. In the latter case, the system is decomposed by a purely mathematical algorithm — hence without any consideration for physical meaning — together with, possibly, some coordinate transformations before and after the decomposition. As the physical decomposition is strongly application dependent but usually intuitive to perform (given enough insight into the problem at hand) it shall not be discussed here.

Mathematical decomposition in itself can be of exact or approximative nature. That is, either they produce equivalent models with identical behaviour, or reduced models (*via model-reduction*) that are a *simplification* of the original system, thus introducing approximation errors. In the exact case, the objective is to yield subsystems that are as independent as possible, as then the remaining, hopefully small interactions among subsystems can be regarded as perturbations to otherwise isolated systems — which facilitates their study significantly. In the approximative case, however, one aims to significantly reduce the size of the system (that is, approximate the overall system with a low-dimensional one) while preserving key properties such as stability, passivity or steady-state response, so that then traditional analysis methods can be applied.

Exact decomposition

While the influential work by Kron (1963) investigated decomposition along physical lines, it was Steward (1962, 1965) that first introduced information flow-based algorithms for identifying sparsity in large systems of equations in order to produce weakly coupled subsystems. Further partitioning / tearing methods were developed in Sargent and Westerberg (1964); Ledet and Himmelblau (1970); Young (1971); Himmelblau (1973). Unfortunately, most decomposition techniques have been developed for systems of algebraic equations only; it appears that the systematic decomposition of dynamic equations is still unresolved. Therefore, decomposition is usually performed based on the physical or structural characteristics of the system.

Model reduction

Classical model reduction techniques for dynamic systems (typically in state-space formulation, both continuous- or discrete-time) are numerous, and basically fall into three categories:

- (i) *Singular value decomposition (SVD) based methods*
- (ii) *Krylov (or moment matching) based methods*
- (iii) *Iterative methods that combine aspects of both.*

As only exact analysis methods are considered in this thesis, we shall not describe these techniques in detail. The interested reader is invited to refer to the excellent tutorial papers by Antoulas *et al.* (1999); Antoulas and Sorensen (2001) and the numerous references therein.

Nonetheless, model reduction techniques can significantly reduce the size of a system to a point where traditional analysis techniques become feasible again. However, in the case of exact decomposition of the system, or where the model is already available in decomposed form, stability of the overall system cannot be readily determined unless the stability properties of the subsystems are aggregated by observing original interconnection structure. This will be discussed in the following subsection.

2.3.2 Aggregation

A natural question to ask is whether stability of an interconnected system can be readily deduced or derived from stability properties of its individual subsystems. To answer this question, it is natural to attempt to somehow “aggregate” the stability properties of the individual systems to determine overall stability. General references discussing the key results in this area include Šiljak (1978); Michel and Miller (1977); Vidyasagar (1981); Michel (1983); Grujić *et al.* (1987); Lakshmikantham *et al.* (1991). It appears that work in

this area has followed two strands: To derive stability with Lyapunov methods, and with input-output methods.

For both approaches, two different assumptions are imaginable, see Šiljak (1978): Either the constituent systems are assumed to be stable in isolation, or they cannot function properly (are unstable) when on their own. This leads to the somewhat philosophical question whether the increase in complexity by interconnecting the systems will lead to an improvement in stability and reliability of the aggregate system, or not. Intuitively, in the second case where the systems are not self-sufficient, interconnection *may* lead to certain cooperative effects that could potentially produce overall stability — contrary to the first case where interconnection may actually produce an *unstable* system, say for instance due to unstable feedback loops being introduced by certain connections.

A key property of large-scale systems is uncertainty in the interconnection structure. Whether this is due to inexact models or time-changing interconnections from structural perturbations, subsystems generally may connect or disconnect from each other during operation, and this behaviour needs to be included in any stability analysis of such systems. To take this into consideration, the concept of *connective stability* was introduced in Šiljak (1972): A system is connectively stable if and only if it remains stable (in the sense of Lyapunov) for all possible interconnection topologies, in other words under *any* structural perturbation. Since this includes in particular the case where all subsystems are completely isolated from each other, one generally assumes that all subsystems are stable on their own, Sandell *et al.* (1978).³

Lyapunov methods

Indeed, the initial work by Bailey (1965) and the flood of subsequent papers followed this path by assuming that a Lyapunov function exists for each subsystem in isolation.⁴ Then, the individual Lyapunov functions can either be cast into another *scalar* Lyapunov function for the aggregate system by forming a weighted sum of the original functions, or they can be combined into what is called a *Vector Lyapunov function* (Bellman, 1962; Matrosov, 1972, 1973). In both cases, the interconnection structure plays an important role: In order to derive stability, certain constraints must be placed on the nature and magnitude of the interactions between the free subsystems.

In the context of large-scale systems, vector Lyapunov functions were first used in the seminal work by Bailey (1965). Subsequent results — both for linear and non-linear systems — were obtained by Piontkovskii and Rutkovskaya (1967); Matrosov (1972, 1973);

³ Exceptions to this assumption however are commented on in the section dedicated to Input-Output based methods, see below.

⁴ Roughly speaking, a Lyapunov function is a norm-like, positive-definite function that decreases along all system trajectories — if one such function can be found, then the system can be shown to be stable, Lyapunov and Fuller (1992). The advantage of using such functions in general is that knowledge of actual solutions of the dynamic system are not required for the stability analysis, and they do not assume linearity of the original system.

Grujić and Šiljak (1973); Šiljak (1983); Lunze (1989); Nersesov and Haddad (2006), most of which rely on the *comparison principle* (Müller, 1926; Lakshmikantham and Leela, 1969; Miller and Michel, 2007) to ultimately show stability of the original problem. References for the scalar Lyapunov function approach include Thompson (1970); Araki *et al.* (1971); Araki and Kondo (1972); Michel and Porter (1972); Michel *et al.* (1982); Liu and Lewis (1992), and some argue that this approach leads to less conservative stability results than in the vector case. In fact, it can be shown that many applications of the vector Lyapunov function approach can be reduced to the scalar approach, Michel (1977).

As mentioned earlier, the nature of the interconnections between the subsystems play an important role. Both procedures require the construction of certain test-matrices, and in many cases the required interconnection properties will cause those test-matrices to be *M-matrices* (which will be discussed in detail in the next chapter). The special properties of this class of matrices plays a key role in the technical proofs of the relevant results; additionally, they elegantly allow to show connective stability, Šiljak (1972).

Generalisations Both methods were generalised in a number of ways, Michel and Miller (1977). To name a few, *matrix Lyapunov functions* were used in Drici (1994); Martyniuk (1998, 2002) to further extend the above techniques to systems with overlapping decompositions (that is systems, where states may be “shared” among subsystems) as well as to find more efficient and less conservative stability tests. For decomposition techniques based on graph theoretic considerations, which can be of great advantage if the connected system is composed of multiple strongly connected components, refer to Michel *et al.* (1978); Tang *et al.* (1980). Discrete time versions of the above results were presented in Araki *et al.* (1971); Grujić and Šiljak (1973); Araki (1975); Martyniuk *et al.* (1996). Modifications of both Lyapunov approaches required for dealing with infinite dimensional systems were considered in Matrosov (1973); Rasmussen and Michel (1976b); Michel and Miller (1978). This allowed to apply these results to systems with delay (Anderson, 1979; Mori *et al.*, 1981; Chang, 1985; Xu, 1995), functional and partial differential equations (Ohta, 1981), Volterra integro-differential equations (Wang *et al.*, 1992) or hybrid systems (Michel and Miller, 1977). Stochastic systems were considered in Michel (1975a); Ladde and Šiljak (1975); Michel (1975b); Rasmussen and Michel (1976a); Socha (1986) and discontinuous systems in Michel and Porter (1971); Ruan (1991); Stipanović and Šiljak (2001).

While one can safely say that the stability theory for large-scale systems based on Lyapunov methods has reached a relatively mature level, Michel (1983), it has one major drawback: Lyapunov stability only applies to the equilibria of *unforced* systems.

Input-Output based methods

While this restriction on the system structure is not only removed by input-output based methods, they also typically give even less conservative results, are more easy to apply in practise as crucial test parameters (the gains) are more readily related to actual design

parameters in the overall system, and the equilibrium of the interconnected system does not need to be known *a priori*, Sandell *et al.* (1978). *Input-output stability* ignores the internal system description and only focuses on the stability of how the system's inputs are mapped to its outputs. In other words, it considers a system to be stable if its outputs will be bounded for every input signal that is also bounded (in some sense), that is, loosely speaking, the system cannot be destabilised by the input.⁵

Literature in this area can again be classified into two main categories, namely deriving methods involving *finite gains*, and methods using notions of *dissipativity* / *passivity*. Both approaches of input-output stability (Sandberg, 1964; Zames, 1966; Desoer and Vidyasagar, 2009) have then been applied to arbitrary interconnections of a large number of *multi-input multi-output* (MIMO) feedback systems. While such interconnections could be viewed as one large MIMO-system in itself, as before, it is often preferable to take advantage of its decomposed form.

Finite gains Initial results that fall in the first category were given by Tokumaru *et al.* (1973); Porter and Michel (1974); Cook (1974); Araki (1976); Lasley and Michel (1976). They followed the typical steps of first assuming that the MIMO subsystems are given in a particular (but very general) standard formulation (often referred to as *input-output feedback system*), then requiring the operators used in these formulations to have *small gains* and the non-linear elements in it to be *sector bounded*, and finally showing stability of the overall system by placing further conditions on the gains of the operators that reflect the interconnection structure. Using such general operator based input-output descriptions allows the theory to also cover non-linear, time-varying systems both in continuous- and discrete-time, Callier *et al.* (1978). The gain condition on the subsystems is required for their input-output stability (*via the small gain theorem*, Zames, 1966). The interconnection gains are usually used to construct a test matrix whose leading principal minors are required to be all positive. Somewhat similar to the Lyapunov-based approach discussed in the previous section, M -matrices again play a key role as they fulfil this property, Lasley and Michel (1976); Moylan (1977); Araki (1978), and also elegantly provide for connective stability. Placing more restrictions on the isolated subsystems and their interconnection structure, a number of additional results are possible such as obtaining *circle criterion* based (Araki, 1978) or *Popov-type* (Lasley and Michel, 1976) stability conditions, or using results from positive operator theory, Sundareshan and Vidyasagar (1977). Graph theoretic decomposition techniques were developed by in Callier *et al.* (1976, 1978) to derive simpler stability tests; this work also helped Vidyasagar (1980) to derive conditions for the well-posedness of large-scale interconnected systems. Input-output stability results for interconnections of stochastic systems were studied in Gutmann and Michel (1979a,b).

⁵ The general input-output approach for linear systems has also received some criticism however as the truncation operator required in most proofs introduces a non-linearity and unwanted harmonics in the frequency domain that make the approach only applicable to certain types of systems, namely *small gain* and *dissipative* systems.

Dissipativity Another way of approaching input-output stability can be found for interconnections of *dissipative* or *passive* systems, Willems (1972); Hill and Moylan (1976); Moylan and Hill (1978); Hill and Moylan (1980). Roughly speaking, the concept of dissipativity is a natural generalisation of Lyapunov theory to open systems (that is systems with inputs and outputs). In the context of dynamical systems it refers to systems that cannot produce energy on their own and cannot store all the energy that is given to them, in other words they “absorb” supplied energy in some way.⁶ The study of such systems often involves construction of an internal function called the *storage function*. For stability analysis, this function can be seen as (or used to derive) a Lyapunov function for the system; in thermodynamics, it can be related to the internal energy and entropy of the system. A classical result (Willems, 1972, 73) shows that any neutral interconnection of dissipative systems forms itself a dissipative system (which is thus input-output stable as well); by “neutral” it is meant that the interconnections must be lossless, i. e. not introduce additional supply or dissipation. This was extended to more general interconnections in Vidyasagar (1977); Moylan and Hill (1978); Sandberg (1978); Vidyasagar (1979) where conditions are presented that require certain test matrices reflecting the interconnection structure to be positive definite. Extensions to discrete-time systems can be found in Haddad *et al.* (2004).

Before moving on to the area of decentralised control we note that attempts have been made to compare and draw parallels between the Lyapunov and input-output stability based approaches, Araki (1978); Moylan and Hill (1978).

2.3.3 Basic concepts of Decentralised Control

Closely related to the stability analysis of large-scale system is the area of decentralised control. Its concepts are somewhat complementary to large-scale systems analysis and, over the last four decades, it has been concerned with developing control techniques that are particularly suited for these types of systems. The decomposition and analysis techniques presented earlier also give answers to the fundamental question of how to break down a given large-scale *control* problem into manageable and only weakly coupled sub-problems, which can then be solved in isolation with relative ease. The implementation of such solutions will be greatly simplified if only locally available information (system states and outputs) are used, and the reduced communication overhead will certainly have reliability and economic benefits as well. Furthermore, delays in the information availability and exchange generally have a detrimental effect on control systems. Thus, if the control stations only use local information that is presumably more readily and quickly available, then this approach poses another advantage over centralised solutions.

⁶ A simple example would be passive components in electrical circuits, such as resistors or capacitors; a transistor in turn is not dissipative as it is an “active” component.

There is a large number of excellent books and survey papers covering this vast topic (including both theory and applications). To name a few, the monographs by Šiljak (1978); Jamshidi (1983); Tamura and Yoshikawa (1990); Šiljak (1991); Lunze (1992); Zečević and Šiljak (2010); Davison and Aghdam (2011) cover the topic more broadly whereas the review papers by Sezer and Šiljak (1996); Sandell *et al.* (1978); Ikeda (1989); Chae and Bien (1991); Šiljak (1996); Šiljak and Zečević (1999); Jiang (2004); Šiljak and Zečević (2005); Bakule (2008); Perutka (2010) are also good starting points to explore the field.

In the following, we briefly give an overview of the typical methodologies encountered, the necessary presumptions to guarantee feasibility of the control problem, and some of the most common design approaches for both weakly and strongly coupled systems.

Methodologies

When attempting to design suitable controllers given the complexity of large-scale systems, three basic methodologies can be identified: i) decentralisation, ii) decomposition, iii) robustness and model simplification, Bakule (2008).

The first one, *decentralisation*, concerns the *structure* of the information to be used in the control system. As stated above, the objective is to only use locally available information in each subsystem, leading to a more or less independent implementation of the control stations. Šiljak (1991) and Lunze (1992) suggest two different scenarios — decentralised controller design for strongly or weakly coupled subsystems. In the first case there is a strong interdependence between subsystems, hence the controller design for each subsystem must take into account at least an approximate model of the neighbouring subsystems, whereas such coupling effects can be neglected in the second case. Clearly, due to the increased complexity of the resulting closed loop system in the first case, weakly coupled systems are preferable for controller design.

The *decomposition* methodology, which was already extensively discussed in the previous section, serves as a tool to analyse and synthesise large-scale systems, with the main goal of reducing the computational complexity of the task. *Robustness analysis and model simplification* attempt to exploit the nature of the uncertainties or the model in order to further reduce the complexity of the control design task.

Reachability and decentrally stabilisable systems

As in classic control theory, *controllability* and *reachability* requirements need to be satisfied for any feedback controller design to succeed. By its very nature, the idea of feedback control consists of regulating a system by some form of action applied to its *inputs*, where this action depends on and is a response to the system's behaviour as reflected by its *outputs*. Clearly, in order for the control action to be successful, it must be able to influence or “reach” the system's states, and the system's states need to be sufficiently “represented” (or at least “observable”) in its outputs for the controller to react appropriately. These two

fundamental concepts are defined as *input-* and *output reachability* (Šiljak, 1978). Inspired by the work of Lin (1974) on structural controllability, analysis of such system properties is formalised by graph-theoretic concepts. To apply this powerful machinery, the state-space model of the system is described as a directed graph (whose vertices are the states, inputs and outputs, and whose arcs represent interactions among them). Structural conditions guaranteeing that systems can indeed be stabilised by a decentralised control action include the so-called *matching conditions* and *non-matching conditions* (Ikeda, 1989; Leitmann, 1993; Šiljak, 1991).

Weakly coupled systems

Loosely speaking, systems where the interaction between different subsystems are only “weak” are referred to as *weakly coupled systems*. In such systems the control design can be performed independently and based on the individual subsystem models only. This allows the wealth of classical control techniques to be employed to achieve suitable stability properties (to name a few, such techniques include for instance pole placement by state feedback, root-locus or parameter plane methods, Šiljak, 1978; Lunze, 1992; Chen *et al.*, 2004; Lunze, 2008). After stabilisation of the isolated systems, an aggregate model of the system is built to derive stability of the interconnected system, taking into account the nature and magnitude of the interactions.

Unfortunately, the more basic control design techniques tend to lead to high-gain feedback solutions which may be prohibitive in practical applications — or even infeasible if the strength of the interconnections is not known *a priori*. This led to shift of attention towards adaptive control solutions where the gains are automatically adjusted as needed for overall stability. An extensive overview of these methods in the context of large-scale systems is given in Šiljak (1996); Perutka (2010).

Nonetheless, weak coupling between subsystems is a desirable property, and the next section discusses a number of techniques from the decentralised control literature that allow decomposition of a given system into weakly coupled systems.

Decomposition techniques for decentralised control

The “decomposition principle” stands for a loose collection of methods surrounding the common goal of breaking down a given large-scale system into a set of lower dimensional subsystems that are weakly coupled. As we mentioned earlier, such decomposition is often done based on physical or structural characteristics of the system, provided of course that the subsystems are sufficiently disjoint in nature. But while tearing along the boundaries of physical subsystems may yield useful insights into the overall system behaviour and interplay of its components, it may not necessarily lead to the most efficient decomposition. Since universal decomposition techniques do not depend on particular *a priori* engineering

knowledge about the system, they can usually be applied to larger classes of problems and additionally lead to computationally more efficient results.

A common decomposition technique, the *nested ϵ -decompositions* (Sezer and Šiljak, 1986, 1991; Zečević and Šiljak, 1994; Amano *et al.*, 1996), consists in its basic form of graph-theoretic algorithm that clusters system states together through symmetric row and column permutations of the matrices of the state-space representation. It yields a weakly coupled collection of subsystems where the strength of the coupling (which impacts the size and number of the subsystems) can be adjusted by varying the ϵ parameter. This basic approach was extended in many directions to cope with time-delays in the interconnections, nonlinear and uncertain interactions, stochastic systems or descriptor systems, to name a few. An extensive list of references for these extensions can be found in Bakule (2008).

Further composition algorithms like the *Lower Block Triangular (LBT) compositions* (Sezer and Šiljak, 1996) or *input and/or output reachable acyclic decompositions* (Šiljak, 1991) yield hierarchical interconnection patterns between the subsystems. These structures offer significant computational advantages when standard feedback controller design or observer design techniques are used.

Another class of decomposition techniques are the so-called *overlapping decompositions* (Šiljak, 1991, 1996). When systems are strongly coupled and *overlap*, they share common parts and inputs, which means that control needs to conform with these information structure constraints. This also means that the overall system will have no effective ϵ -decompositions in its original form. To deal with these situations, one often-used approach consists of expanding the original problem (with its strongly coupled subsystems) into a higher dimensional system where the subsystems then appear weakly coupled and permit a suitable ϵ -decomposition — an *overlapping ϵ -decomposition*. A general framework for this concept and surrounding ideas is given by the *inclusion principle*, see Ikeda *et al.* (1981); Bakule (1985); Šiljak (1991); Chu and Šiljak (2005).




Figure 2.6: Different matrix structures after decomposition, c.f. Šiljak and Zečević (2005).

A related class of decompositions for strongly coupled systems are *BBD decompositions* (Šiljak, 1996; Bakule, 2008; Zečević and Šiljak, 2005b, 2010). Whereas in disjoint systems the feedback gain matrices (relating the system outputs to the inputs) can be transformed into block-diagonal (BD) forms, this is not possible in overlapping systems,

and only *block tri-diagonal* (BTD) or *bordered block-diagonal* (BBD) forms can be achieved (see Figure 2.6 on the previous page for an illustration of these structures). Nonetheless, these formulations have still the advantage that they allow controllers gains particularly in *very* large and sparse systems to be computed in an efficient way, in particular allowing those computations to be performed on massively parallelised architectures with minimal inter-processor communication overheads.

Many of the existing techniques for overlapping systems, see Šiljak and Zečević (2005); Bakule (2008) for a comprehensive overview, involve linear matrix inequalities (Boyd *et al.*, 1994) for which efficient solvers exist, Šiljak and Stipanović (2000); Šiljak and Zečević (2005); Zečević and Šiljak (2005a); Swarnakar *et al.* (2007).

These remarks conclude this section on large-scale systems and decentralised control. The idea of cooperatively controlling a large system's behaviour is closely related to the area of decentralised control, but has been treated somewhat separately in the literature. Decentralised control is typically concerned with an overall system that is to exhibit a certain behaviour, and to achieve this, a global control goal is decentralised. In cooperative control, a somewhat different angle of attack appears to be taken, presenting more of a bottom-up approach: A large number of individual, largely similar and mostly autonomous entities is joined up to form an aggregate, networked system that is then to exhibit a certain behaviour.

2.4 Cooperation and consensus

As mentioned above, consensus and cooperation in networked multi-agent systems has recently attracted much attention in the research community. For a great introduction into the field and examples of its many, diverse applications see for instance the surveys by Ren *et al.* (2005), Olfati-Saber *et al.* (2007) and Murray (2007), as well as the collection of references at Reynolds (2001).

2.4.1 History

Consensus and agreement problems were studied systematically as early as the 1960 in the context of management science and statistics, Eisenberg and Gale (1959); Norvig (1967); Winkler (1968); DeGroot (1974). Later, those ideas were picked up in different contexts, such as fusion of sensor data (Luo and Kay 1989; Benediktsson and Swain 1992; Estrin *et al.* 2001; Olfati-Saber and Shamma 2005; or see the proceedings of the IEEE conferences on Multisensor Fusion and Integration for Intelligent Systems), medicine (Weller and Mann, 1997), decentralised estimation (Levy *et al.*, 1983; Mutambara, 1998; Gupta, 2006; Olfati-Saber, 2007), clock synchronisation (Schenato and Gamba, 2007; Carli *et al.*, 2008), or simulation of flocking behaviour (Reynolds, 1987; Vicsek *et al.*, 1995; see also Figure 2.7 on the facing page for an example) just to name a few.




Figure 2.7: Illustration of a flock of birds where, in grossly simplified terms, each bird adjust its speed and heading relative that of nearby flockmates, which leads to the coordinated group behaviour often observed in nature (such as in bird flocks, fish schools, herds, etc.)

2.4.2 Networked dynamic systems

Particularly in the last decade the general problem of consensus finding in networked dynamic systems has been focused on intensely. It typically comes in many “flavours” depending on the application. These variations include whether the topology of the graph representing the inter-agent communications remains fixed or changes over time; it is undirected or directed; the agents can manipulate the state on which to reach consensus instantly or only with certain dynamics; if each node’s state is scalar or multidimensional; whether there are delays in the information exchange; or if all nodes update their states in a synchronous fashion or on their own pace. While the initial work by Borkar and Varaiya (1982); Tsitsiklis (1984); Tsitsiklis *et al.* (1986); Reynolds (1987); Vicsek *et al.* (1995) on consensus and coordination was based on bi-directional information exchange between neighbouring nodes (leading to undirected communication graphs) with rigorous convergence proofs given in Jadbabaie *et al.* (2003), this has been extended to include directed communication graphs for instance in Beard and Stepanyan (2003); Olfati-Saber and Murray (2004); Moreau (2005); Ren and Beard (2005); Fang *et al.* (2005). Another generalisation allowed asynchronous consensus protocols so that not all nodes had to perform state updates at the same instant, Olfati-Saber and Murray (2004); Hatano and Mesbahi (2005); Blondel *et al.* (2005); Fang *et al.* (2005); Cao *et al.* (2006). Closely related was the work that also considered changing graph topologies, Jadbabaie *et al.* (2003); Tanner *et al.* (2003b); Beard and Stepanyan (2003); Ren and Beard (2005); Olfati-Saber (2006). Further generalisations of the problem allowed the inclusion of agent dynamics (typically linear, second order systems) in the consensus problem, Tanner *et al.* (2003a,b); Olfati-Saber and Murray (2003); Olfati-Saber (2006), which play an important role in networks of mobile agents that move with finite dynamics. In some situations the consensus variable may not be directly altered by the nodes, but only implicitly. Such a situation is dealt with in Stanojević and Shorten (2008, 2009b).

However, most of these papers only focus on so-called *unconstrained* consensus applications. When the consensus, that the system is to reach, should fulfil external conditions (such as a common heading of a flock of agents, but in a particular direction), three approaches are usually taken, see Beard *et al.* (2001); Lawton *et al.* (2003); Ren and Beard (2004) and citations therein: leader-following (Wang, 1991; Mesbahi and Hadaegh, 1999; Singh *et al.*, 2000; Fax and Murray, 2004; Ji *et al.*, 2006), virtual structure based (Lewis and Tan, 1997; Beard *et al.*, 2000; Shi *et al.*, 2006) or behaviour based (Balch and Arkin, 1998; Anderson and Robbins, 1998; Lawton *et al.*, 2003; Parker, 1998; Chen and Luh, 1994; Veloso *et al.*, 2000) approaches.

Leader-following

The first concept presents a common technique used typically to make formations of autonomous mobile agents follow desired trajectories. The idea is that all agents in the are programmed to follow a designated “leader” node, as sketched in Figure 2.8 below. However, the problem with these architectures is usually that they not only depend heavily on the leader, but it appears that little discussion of the case where the leader adjusts its state based on feedback of the totality of the states of the network has taken place, and most of the systems dealt with in that context are linear.




Figure 2.8: Illustration of three steps of a typical leader following based control algorithm. With the system in a given position (step 1), the leader moves somewhere (step 2) in response to which the other agents move to follow him (step 3).

Virtual structures

In the virtual structure approach, the entire network of agents is treated as a single entity, the virtual structure. The desired behaviour is then assigned to the virtual structure relative to which each member controls its own behaviour. This approach is illustrated in Figure 2.9 on the facing page.




Figure 2.9: Illustration of three steps of a typical virtual structure based control algorithm. With the system in a given position (step 1), the virtual structure is moved (step 2) in response to which all agents move to follow their assigned positions relative to the virtual structure (step 3).

Behaviour based

In the behavioural approach, each agent's behaviour is based on a combination (e.g. weighted sum) of a number of desired behaviours, such as goal seeking, formation keeping, obstacle and collision avoidance, *etc.* An example for this is shown in Figure 2.10 below. A typical application of these techniques are *rendez-vous* problems with obstacle and collision avoidance, where the agents are to meet in a certain place, but avoid running into obstacles or crashing into each other during the approach.

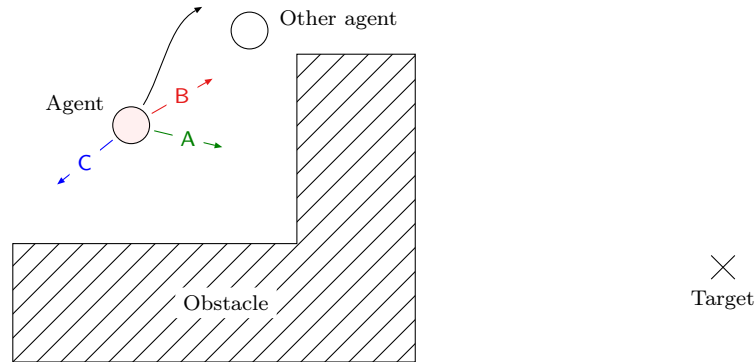


Figure 2.10: Illustration of the behaviour based approach, where the agent's final action is a combination of three desirable behaviours: Goal seeking (A), obstacle avoidance (B) and collision avoidance with other agents (C).

It is in this third class that our work later in Chapters 4 and 5 could be placed, as the desired behaviour of the agents in our networks is both a combination of localised and global constraints.

With these remarks we close this literature review section and move on to present our first sets of results for switched positive systems.

Switching

This chapter develops necessary and sufficient conditions for the existence of common linear co-positive Lyapunov functions first for the state-dependent and then the arbitrary switching case for sets of positive LTI systems, both in continuous-time and discrete-time. Additionally, numeric methods for checking these conditions are provided, we discuss what can be done if the conditions are violated, and also provide a few preliminary examples for our results.

Chapter contents

- 3.1 Introduction
 - 3.2 Preliminaries
 - 3.3 State-dependent switching case
 - 3.4 Arbitrary switching case
 - 3.5 Discrete-time switched positive systems
 - 3.6 Examples of usage
 - 3.7 Conclusion
-

3.1 Introduction

The focus of this chapter will be on switched positive linear time-invariant (LTI) systems, and in particular on the existence of common linear co-positive Lyapunov functions (CLCLF). It presents joint work with Dr. O. Mason and Prof. R. Shorten and has been published in Knorn *et al.* (2009a,b).¹

In some sense, such systems may be interpreted as a (possibly dense) interconnection of scalar systems, where the graph describing the system interactions changes abruptly over time. Now, recall the well known result that the existence of a linear co-positive Lyapunov function is both necessary and sufficient for the exponential stability of a positive linear time-invariant (LTI) system, Farina and Rinaldi (2000). In light of our earlier remarks concerning *common* Lyapunov functions in general it may appear overly conservative to

¹ It should also be noted that Theorem 3.2 may be deduced from the independent, more general results on **P**-matrix sets given in Song *et al.* (1999), of which the author was unaware of when the result was developed.

study the existence of such Lyapunov functions for switched systems. However, establishing conditions under which such functions exist is nonetheless a natural place to begin the study of stability of switched positive linear systems.

For one, common Lyapunov functions are very useful since existence of such functions implies exponential stability of the overall switched system, Fornasini and Valcher (2011). Additionally, many of the interesting properties of positive systems can be attributed to the existence of linear co-positive Lyapunov functions. Of particular interest is the work by Haddad and Chellaboina (2004), in which the existence of such a function was related to delay independent stability properties that are possessed by many positive systems. Exploiting these properties further, we will later demonstrate the use of one of the main results in the applications chapter (Section 6.1 on page 117).

Contributions

Inspired by this and related work, the main contributions of this present chapter will be the derivation of tractable conditions for the existence of a *common* linear co-positive Lyapunov function for a finite number of LTI systems that are associated either with the entire positive orthant (arbitrary switching) or with polyhedral regions partitioning the positive orthant (state-dependent switching). In both cases, compact and easily verifiable conditions are obtained. We also show that our results directly carry over to the discrete-time case.

Structure

The rest of this chapter is structured as follows: The next section sets up the notation and defines linear co-positive Lyapunov functions. We then present our main results both for the case of state-dependent switching (Section 3.3), and for arbitrary switching (Section 3.4). Next, we shall discuss how these results can easily be applied to discrete-time systems. Finally, before making some concluding statements, Section 3.6 highlights the significance of our results and gives a number of examples that motivate their use.

3.2 Preliminaries

3.2.1 Notation

For general notational conventions, please take note of the Notation section on page 153. We say that matrices or vectors are positive (non-negative) if all their entries are positive (non-negative); this is written as $\mathbf{A} \succ \mathbf{0}$ resp. $\mathbf{A} \succeq \mathbf{0}$, where $\mathbf{0}$ is the zero-matrix of appropriate dimension. A matrix \mathbf{A} is said to be *Hurwitz stable* (or just “Hurwitz”) if all its eigenvalues lie in the open left half of the complex plane. A matrix is said to be *Metzler*

(in the literature also referred to as *essentially non-negative*) if all its off-diagonal entries are non-negative (Metzler, 1945).

Also, let $\mathcal{C} \subseteq \mathbb{R}^n$ be a *closed, pointed, solid convex cone* (or *proper convex cone*) if and only if its interior is not empty and $\alpha\mathbf{x} + \beta\mathbf{y} \in \mathcal{C}$ for any $\mathbf{x}, \mathbf{y} \in \mathcal{C}$ and non-negative scalars α, β . Such cone is called *polyhedral* if and only if it can be written as the intersection of finitely many closed half spaces, each containing the origin on its boundary, Berman and Plemmons (1979). In other words, it has finitely many *extremal rays* (or *generators*). Figure 3.1 below gives an illustration of a polyhedral proper convex cone in $\mathbb{R}_{\geq 0}^3$ with three extremal rays.




Figure 3.1: Illustration of a polyhedral proper convex cone in $\mathbb{R}_{\geq 0}^3$ with three generators.

3.2.2 Definitions

A dynamic system is called *positive*² if and only if, for any non-negative initial condition, all its states remain in the closed positive orthant throughout time (irrespective of the system being stable or not). A classic result for LTI systems shows that a necessary and sufficient condition for this to hold true is that the system matrix \mathbf{A} is a Metzler matrix: In that case (and only that case) $e^{\mathbf{A}t}$, which characterises the solution of the differential equation, is non-negative for all $t \geq 0$ (Luenberger, 1979), implying that all solutions starting from non-negative initial conditions remain non-negative.

We now define the class of switched positive linear systems considered in the following.

Definition 3.1 (Switched positive linear system, continuous time) _____

A switched positive linear time-invariant system with N modes and of dimension n is a dynamical system of the form

$$\dot{\mathbf{x}}(t) = \mathbf{A}_{s(\mathbf{x}(t), t)} \mathbf{x}(t) \quad \text{with} \quad \mathbf{x}(t=0) = \mathbf{x}_0 \succeq \mathbf{0} \quad (3.1)$$

² Technically, one may also use the word “non-negative”, which would be more accurate, but the term “positive” is typically used.

where $s : \mathbb{R}^n \times \mathbb{R} \rightarrow \{1, \dots, N\}$ is some piecewise constant switching signal (or switching function or switching sequence) which may or may not depend on the state vector $\mathbf{x}(t)$, and where $\mathbf{A}_1, \dots, \mathbf{A}_N \in \mathbb{R}^{n \times n}$ are the system matrices of the constituent systems (or subsystems or modes).

Furthermore, we will always assume that all the \mathbf{A}_i matrices are Metzler matrices (in order to ensure positivity of the system) and Hurwitz matrices (in order to ensure stability of each individual mode).

Finally, we define the following type of Lyapunov function:

Definition 3.2 (Linear co-positive Lyapunov function)

The function $v(\mathbf{x}) = \mathbf{v}^\top \mathbf{x}$ is said to be a linear co-positive Lyapunov function (LCLF) for the positive LTI system $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}$ if and only if $v(\mathbf{x}) > 0$ and $\dot{v}(\mathbf{x}) < 0$ for all $\mathbf{x} \succ \mathbf{0}$, or, equivalently, $\mathbf{v} \succ \mathbf{0}$ and $\mathbf{v}^\top \mathbf{A} \prec \mathbf{0}$.

For more background on Lyapunov theory and related concepts, especially in the context of switched systems, please refer to the references presented in the literature review on page 8.

3.3 State-dependent switching case

We first consider necessary and sufficient conditions for the existence of common linear co-positive Lyapunov functions (CLCLF) for sets of positive LTI systems where each constituent system is associated with a closed convex region of the closed positive orthant.

3.3.1 Main result

Consider the following partition of the state-space: Assume that there exist N — possibly overlapping — proper convex cones $\mathcal{C}_i \subseteq \mathbb{R}_{\geq 0}^n$ such that the closed positive orthant $\mathbb{R}_{\geq 0}^n$ can be written as $\mathbb{R}_{\geq 0}^n = \cup_{i=1}^N \mathcal{C}_i$. Moreover, assume that there are N stable positive LTI subsystems $\dot{\mathbf{x}} = \mathbf{A}_i \mathbf{x}$ such that the i th mode can only be active when the state vector is in the cone \mathcal{C}_i .

Our first main result gives a necessary and sufficient condition for the existence of a CLCLF for this type of switched positive linear system with state-dependent switching. Formally, we provide a condition for the existence of a vector $\mathbf{v} \succ \mathbf{0}$ such that $\mathbf{v}^\top \mathbf{A}_i \mathbf{x}_i < 0$ for all non-zero $\mathbf{x}_i \in \mathcal{C}_i$ for $i = 1, \dots, N$.

Theorem 3.1 (Existence CLCLF, state-dependent switching)

Given N Metzler and Hurwitz matrices $\mathbf{A}_1, \dots, \mathbf{A}_N \in \mathbb{R}^{n \times n}$ and N proper convex cones $\mathcal{C}_1, \dots, \mathcal{C}_N \subseteq \mathbb{R}_{\geq 0}^n$ such that $\mathbb{R}_{\geq 0}^n = \cup_{i=1}^N \mathcal{C}_i$, precisely one of the following two statements is true:

- (i) *There is a positive vector $\mathbf{v} \in \mathbb{R}^n$ such that $\mathbf{v}^\top \mathbf{A}_i \mathbf{x}_i < 0$ for all non-zero $\mathbf{x}_i \in \mathcal{C}_i$ and $i = 1, \dots, N$.*
- (ii) *There are vectors $\mathbf{x}_i \in \mathcal{C}_i$, with $i = 1, \dots, N$, not all zero such that $\sum_{i=1}^N \mathbf{A}_i \mathbf{x}_i \succeq \mathbf{0}$.*

Proof (ii) \Rightarrow \neg (i):³ Assume that (ii) holds. Then, for any $\mathbf{v} \succ \mathbf{0}$ we have

$$\mathbf{v}^\top \mathbf{A}_1 \mathbf{x}_1 + \dots + \mathbf{v}^\top \mathbf{A}_N \mathbf{x}_N \geq 0 \quad (3.2)$$

which immediately implies that (i) cannot hold.

\neg (ii) \Rightarrow (i): Assume that (ii) does not hold, i.e. there are no vectors $\mathbf{x}_i \in \mathcal{C}_i$ not all zero such that $\sum_{i=1}^N \mathbf{A}_i \mathbf{x}_i \succeq \mathbf{0}$. This means that the following intersection of convex cones is empty:

$$\underbrace{\left\{ \sum_{i=1}^N \mathbf{A}_i \mathbf{x}_i : \mathbf{x}_i \in \mathcal{C}_i, \text{ not all zero} \right\}}_{\mathcal{O}_1} \cap \underbrace{\left\{ \mathbf{x} \succeq \mathbf{0} \right\}}_{\mathcal{O}_2} = \emptyset. \quad (3.3)$$

By scaling appropriately it is easy to see that this is equivalent to:

$$\underbrace{\left\{ \sum_{i=1}^N \mathbf{A}_i \mathbf{x}_i : \mathbf{x}_i \in \mathcal{C}_i, \sum_{i=1}^N \|\mathbf{x}_i\|_1 = 1 \right\}}_{\bar{\mathcal{O}}_1} \cap \underbrace{\left\{ \mathbf{x} \succeq \mathbf{0} \right\}}_{\mathcal{O}_2} = \emptyset \quad (3.4)$$

where $\|\cdot\|_1$ denotes the usual spatial 1-norm. Now, $\bar{\mathcal{O}}_1$ and \mathcal{O}_2 are disjoint non-empty closed convex sets and additionally $\bar{\mathcal{O}}_1$ is bounded. Thus, we can apply Corollary 4.1.3 from Hiriart-Urruty and Lemaréchal (2001) which guarantees the existence of a vector $\mathbf{v} \in \mathbb{R}^n$ such that

$$\max_{\mathbf{y} \in \bar{\mathcal{O}}_1} \mathbf{v}^\top \mathbf{y} < \inf_{\mathbf{y} \in \mathcal{O}_2} \mathbf{v}^\top \mathbf{y} \quad (3.5)$$

As the zero vector is in \mathcal{O}_2 , it follows $\inf_{\mathbf{y} \in \mathcal{O}_2} \mathbf{v}^\top \mathbf{y} \leq 0$. However, as \mathcal{O}_2 is the cone $\{\mathbf{x} \succeq \mathbf{0}\}$ it also follows that $\inf_{\mathbf{y} \in \mathcal{O}_2} \mathbf{v}^\top \mathbf{y} \geq 0$. Thus, $\inf_{\mathbf{y} \in \mathcal{O}_2} \mathbf{v}^\top \mathbf{y} = 0$. Hence, $\mathbf{v}^\top \mathbf{y} \geq 0$ for all $\mathbf{y} \in \mathcal{O}_2$ and it follows that $\mathbf{v} \succeq \mathbf{0}$. Moreover, from (3.5), we can conclude that for any $i = 1, \dots, N$ and any $\mathbf{x}_i \in \mathcal{C}_i$ with $\|\mathbf{x}_i\|_1 = 1$, $\mathbf{v}^\top \mathbf{A}_i \mathbf{x}_i < 0$. As $\mathcal{C}_i \cap \{\mathbf{x} \succeq \mathbf{0} : \|\mathbf{x}\|_1 = 1\}$ is compact, it follows from continuity that by choosing $\epsilon > 0$ sufficiently small, we can guarantee that $\mathbf{v}_\epsilon := \mathbf{v} + \epsilon \mathbf{1} \succ \mathbf{0}$ satisfies $\mathbf{v}_\epsilon^\top \mathbf{A}_i \mathbf{x}_i < 0$ for all $\mathbf{x}_i \in \mathcal{C}_i \cap \{\mathbf{x} \succeq \mathbf{0} : \|\mathbf{x}\|_1 = 1\}$ and all $i = 1, \dots, N$.

Finally, it is easy to see that $\mathbf{v}_\epsilon^\top \mathbf{A}_i \mathbf{x}_i < 0$ is true even without the norm requirement on \mathbf{x}_i .

This completes the proof of Theorem 3.1. □

³ That is, we show that if (ii) is true, then (i) cannot hold.

Comment The theorem thus provides a necessary and sufficient condition for the existence of a CLCLF. Condition (ii) basically means that *if* (and only if) there is a non-trivial linear combination of the all the columns of the different constituent system matrices (using vectors taken from the corresponding cones) that yields a non-negative value *then* no CLCLF exists for the switched system. Unfortunately, to the best of the author's knowledge, this condition in its present form is difficult to check numerically. However, a slight reformulation changes this. /

3.3.2 Numerical test based on a linear program

To establish a simple numerical test, we note that polyhedral proper convex cones \mathcal{C} with k extremal rays in the non-negative orthant of the $\mathbb{R}_{\geq 0}^n$ can be expressed as

$$\mathcal{C} := \left\{ \mathbf{x} \mid \mathbf{x} = \sum_{i=1}^k \alpha_i \mathbf{Q}^{(i)}, \alpha_i \geq 0, i = 1, \dots, k \right\} \quad (3.6)$$

where $\mathbf{Q} \in \mathbb{R}_{\geq 0}^{n \times k}$ is the (non-singular) *generating matrix* of the cone, and $\mathbf{Q}^{(i)}$ denotes the i th column of \mathbf{Q} . This generating matrix can then be included in the second condition of the previous theorem to yield the following corollary

Corollary 3.1 (Existence CLCLF, state-dependent switching, polyhedral cones) _____

Given N Metzler and Hurwitz matrices $\mathbf{A}_1, \dots, \mathbf{A}_N \in \mathbb{R}^{n \times n}$ and N polyhedral proper convex cones \mathcal{C}_i of the type (3.6) such that $\mathbb{R}_{\geq 0}^n = \cup_{i=1}^N \mathcal{C}_i$, precisely one of the following two statements is true:

- (i) *There is a positive vector $\mathbf{v} \in \mathbb{R}^n$ such that $\mathbf{v}^\top \mathbf{A}_i \mathbf{x}_i < 0$ for all non-zero $\mathbf{x}_i \in \mathcal{C}_i$ and $i = 1, \dots, N$.*
- (ii) *There are vectors $\mathbf{w}_i \succeq \mathbf{0}$ not all zero such that $\sum_{i=1}^N \mathbf{B}_i \mathbf{w}_i \succeq \mathbf{0}$, where $\mathbf{B}_i := \mathbf{A}_i \mathbf{Q}_i$.*

Proof Virtually identical to that of Theorem 3.1.

The advantage of this reformulation now is that condition (ii) can be checked efficiently by running a simple feasibility check on a suitably defined linear program, Bertsekas *et al.* (2003). For example, it is straightforward to see that (ii) is fulfilled if and only if the following linear program is feasible:

$$\begin{aligned} \text{argmax} \quad & \mathbf{1}^\top \tilde{\mathbf{w}} \\ \text{subject to} \quad & \tilde{\mathbf{B}} \tilde{\mathbf{w}} \succeq \mathbf{0}, \quad \tilde{\mathbf{w}} \succeq \mathbf{0}, \quad \tilde{\mathbf{w}} \preceq \mathbf{1} \end{aligned} \quad (3.7)$$

where $\tilde{\mathbf{B}}$ corresponds to the horizontally concatenated \mathbf{B}_i , and $\tilde{\mathbf{w}}$ to the vertically stacked \mathbf{w}_i . It is then straightforward to run a feasibility check on this linear program, to provide

an answer in polynomial time. For similar results, the reader may refer to Rami and Tadeo Rico (2007).

Comment As we noted before in the literature review, such numerical tests will certainly be useful in practical applications. However, their major drawback is that they typically give little insight as to *why* a system may be stable or not. They only answer the stability question with “yes” or “no”, but in case the answer is “no”, do not help establishing why this may be the case. /

In the following section, we will present an analytical test for the arbitrary switching case, completing initial work reported in Mason and Shorten (2007). Furthermore, we shall also comment on how it can give more extensive insights in the stability question.

3.4 Arbitrary switching case

An important special case of the previous results is when each of the cone generating matrices \mathbf{Q}_i are the identity matrix. In that case, each switching restricting cone is the positive orthant itself, meaning that there are no more switching restraints and arbitrary switching between the modes is allowed. Then, condition (ii) of the corollary above offers another interpretation: The convex hull of the (polyhedral convex) cone generated by all the columns of the \mathbf{A}_i must not intersect the closed positive orthant except in the origin in order for a CLCLF to exist.

However, additional necessary and sufficient conditions for the existence of a CLCLF for each of the constituent systems can be derived — guaranteeing stability of the overall system under arbitrary switching. This will be given by Theorem 3.2 below.

3.4.1 Main result

Before stating Theorem 3.2, we need a technical result which will simplify its proof significantly. The following lemma is in fact very similar to Theorem 3.1, when each of the generating matrices \mathbf{Q}_i is the identity matrix.

Lemma 3.1

Given N Metzler and Hurwitz matrices $\mathbf{A}_1, \dots, \mathbf{A}_N \in \mathbb{R}^{n \times n}$ the following two statements are equivalent:

- (i) *There is a non-zero $\mathbf{v} \succeq \mathbf{0}$ such that $\mathbf{v}^\top \mathbf{A}_i \preceq \mathbf{0}$ for all $i = 1, \dots, N$.⁴*
- (ii) *There are no $\mathbf{w}_i \succ \mathbf{0}$ such that $\sum_{i=1}^N \mathbf{A}_i \mathbf{w}_i = \mathbf{0}$.*

⁴ Note that with the assumptions of the lemma, $\mathbf{v}^\top \mathbf{A}_i$ will always be non-zero for a non-zero $\mathbf{v} \succeq \mathbf{0}$.

Proof (i) \Rightarrow (ii): Assume there is a non-zero vector $\mathbf{v} \succeq \mathbf{0}$ such that $\mathbf{v}^\top \mathbf{A}_i \preceq \mathbf{0}$ for all $i = 1, \dots, N$. Thus,

$$\mathbf{v}^\top \mathbf{A}_1 + \dots + \mathbf{v}^\top \mathbf{A}_N \preceq \mathbf{0} \quad (3.8)$$

and for any set of strictly positive vectors $\mathbf{w}_i \succ \mathbf{0}$,

$$\mathbf{v}^\top \mathbf{A}_1 \mathbf{w}_1 + \dots + \mathbf{v}^\top \mathbf{A}_N \mathbf{w}_N < 0 \quad (3.9)$$

$$\mathbf{v}^\top (\mathbf{A}_1 \mathbf{w}_1 + \dots + \mathbf{A}_N \mathbf{w}_N) < 0 \quad (3.10)$$

so that

$$\mathbf{A}_1 \mathbf{w}_1 + \dots + \mathbf{A}_N \mathbf{w}_N \neq \mathbf{0} \quad (3.11)$$

In other words, there are no vectors $\mathbf{w}_i \succ \mathbf{0}$ such that $\sum_{i=1}^N \mathbf{A}_i \mathbf{w}_i = \mathbf{0}$.

(ii) \Rightarrow (i): Assuming that there are no vectors $\mathbf{w}_i \succ \mathbf{0}$ such that $\sum_{i=1}^N \mathbf{A}_i \mathbf{w}_i = \mathbf{0}$, we can write

$$\{\mathbf{A}_1 \mathbf{w}_1 + \dots + \mathbf{A}_N \mathbf{w}_N : \mathbf{w}_i \succ \mathbf{0}\} \cap \{\mathbf{0}\} = \emptyset \quad (3.12)$$

Since the \mathbf{A}_i are all Metzler and Hurwitz matrices, it is easy to show that this implies

$$\underbrace{\{\mathbf{A}_1 \mathbf{w}_1 + \dots + \mathbf{A}_N \mathbf{w}_N : \mathbf{w}_i \succ \mathbf{0}\}}_{\mathcal{O}_1} \cap \underbrace{\{\mathbf{x} \succ \mathbf{0}\}}_{\mathcal{O}_2} = \emptyset \quad (3.13)$$

This corresponds to the intersection of two open convex cones, \mathcal{O}_1 and \mathcal{O}_2 . As this intersection is empty, the two cones are disjoint and there must exist a separating hyperplane between them, see for instance Rockafellar (1970). In other words, there is a vector $\mathbf{v} \in \mathbb{R}^n$ such that

$$\mathbf{v}^\top \mathbf{y} < 0 \text{ for all } \mathbf{y} \in \mathcal{O}_1 \quad \text{and} \quad \mathbf{v}^\top \mathbf{y} > 0 \text{ for all } \mathbf{y} \in \mathcal{O}_2 \quad (3.14)$$

From the second inequality we get that \mathbf{v} has to be non-negative (and non-zero). The first inequality, in turn, can be written as

$$\mathbf{v}^\top \mathbf{A}_1 \mathbf{w}_1 + \dots + \mathbf{v}^\top \mathbf{A}_N \mathbf{w}_N < 0 \quad \text{for all } \mathbf{w}_i \succ \mathbf{0} \quad (3.15)$$

Furthermore, since $\mathbf{v} \succeq \mathbf{0}$, and since the inequality has to hold for any choice of (strictly positive) vectors \mathbf{w}_i , each individual summand must be less than or equal to zero. However, this can only be the case if $\mathbf{v}^\top \mathbf{A}_i \preceq \mathbf{0}$ for $i = 1, \dots, N$, which completes the proof of Lemma 3.1 on the preceding page. \square

Some additional notation is also required for the presentation of our second main result. Let the set containing all possible mappings $\sigma : \{1, \dots, n\} \rightarrow \{1, \dots, N\}$ be called $\mathcal{S}_{n,N}$,

for positive integers n and N . Given N matrices \mathbf{A}_i , these mappings will then be used to construct matrices $\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$ in the following way:

$$\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N) := \begin{bmatrix} \mathbf{A}_{\sigma(1)}^{(1)} & \mathbf{A}_{\sigma(2)}^{(2)} & \dots & \mathbf{A}_{\sigma(n)}^{(n)} \end{bmatrix} \quad (3.16)$$

that is, the i th column $\mathbf{A}_\sigma^{(i)}$ of \mathbf{A}_σ is the i th column of one of the $\mathbf{A}_1, \dots, \mathbf{A}_N$ matrices, depending on the mapping $\sigma \in \mathcal{S}_{n,N}$.

We can now state the following theorem giving a necessary and sufficient condition for the existence of a linear co-positive Lyapunov function for arbitrary switching between finitely many positive LTI systems of finite dimension:

Theorem 3.2 (CLCLF existence, arbitrary switching) _____

Given a finite number of Hurwitz and Metzler matrices $\mathbf{A}_1, \dots, \mathbf{A}_N \in \mathbb{R}^{n \times n}$, the following statements are equivalent:

- (i) *There is a strictly positive vector $\mathbf{v} \in \mathbb{R}^n$ such that $\mathbf{v}^\top \mathbf{A}_i \prec \mathbf{0}$ for all $i = 1, \dots, N$.*
- (ii) *$\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$ is Hurwitz for all $\sigma \in \mathcal{S}_{n,N}$.*

Proof (i) \Rightarrow (ii): Assuming that there exists a positive vector $\mathbf{v} \in \mathbb{R}^n$ such that $\mathbf{v}^\top \mathbf{A}_i \prec \mathbf{0}$ for all $i = 1, \dots, N$, this of course implies, when looking at the columns of the matrices \mathbf{A}_i , that $\mathbf{v}^\top \mathbf{A}_i^{(j)} < 0$ for any $i = 1, \dots, N$ and $j = 1, \dots, n$. Thus, it follows that $\mathbf{v}^\top \mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N) \prec \mathbf{0}$ for all $\sigma \in \mathcal{S}_{n,N}$. Next, we note that since the $\mathbf{A}_1, \dots, \mathbf{A}_N$ are all Metzler matrices, by construction so must be all the $\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$, $\sigma \in \mathcal{S}_{n,N}$. Finally, applying Theorem 2.5.3 from Horn and Johnson (1991), we have that all matrices $\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$, $\sigma \in \mathcal{S}_{n,N}$, must be Hurwitz.

$\neg(i) \Rightarrow \neg(ii)$: We show that if there does not exist a vector \mathbf{v} as described in (i), then at least one of the matrices $\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$ is not a Hurwitz matrix for some $\sigma \in \mathcal{S}_{n,N}$.

To begin, assume that there is no non-zero $\mathbf{v} \succeq \mathbf{0}$ such that $\mathbf{v}^\top \mathbf{A}_i \preceq \mathbf{0}$ for all $i = 1, \dots, N$ (note that this is a stronger assumption than the non-existence of a strictly positive vector \mathbf{v} , as stated in (i); we will relax this assumption below). From Lemma 3.1 on page 39 we then know that there is at least one set of vectors $\mathbf{w}_i \succ \mathbf{0}$ such that

$$\mathbf{A}_1 \mathbf{w}_1 + \dots + \mathbf{A}_N \mathbf{w}_N = \mathbf{0} \quad (3.17)$$

Next, we express $\mathbf{w}_2, \dots, \mathbf{w}_N$ in terms of \mathbf{w}_1 using diagonal matrices: $\mathbf{w}_i = \mathbf{D}_i \mathbf{w}_1$ where $\mathbf{D}_i = \text{diag} \{d_i^{(jj)}\}$ and $d_i^{(jj)} > 0$ for all $i = 1, \dots, N$ and $j = 1, \dots, n$. We can then rewrite Equation (3.17) as

$$\mathbf{A}_1 \mathbf{D}_1 \mathbf{w}_1 + \mathbf{A}_2 \mathbf{D}_2 \mathbf{w}_1 + \dots + \mathbf{A}_N \mathbf{D}_N \mathbf{w}_1 = \mathbf{0} \quad (3.18)$$

$$(\mathbf{A}_1 \mathbf{D}_1 + \dots + \mathbf{A}_N \mathbf{D}_N) \mathbf{w}_1 = \mathbf{0} \quad (3.19)$$

and thus, since $\mathbf{w}_1 \succ \mathbf{0}$, we must have for the determinant

$$\det(\mathbf{A}_1 \mathbf{D}_1 + \dots + \mathbf{A}_N \mathbf{D}_N) = 0 \quad (3.20)$$

To simplify notation, define for each mapping $\sigma \in \mathcal{S}_{n,N}$ the following product

$$p_\sigma := \prod_{j=1}^n d_{\sigma(j)}^{(jj)} \quad (3.21)$$

for which we note that $p_\sigma > 0$ for all $\sigma \in \mathcal{S}_{n,N}$ since $d_i^{(jj)} > 0$ for all i and j . Using the fact that the determinant of a matrix is multilinear in the columns of that matrix, we can now express the left-hand side of (3.20) as

$$\det(\mathbf{A}_1 \mathbf{D}_1 + \dots + \mathbf{A}_N \mathbf{D}_N) = \sum_{\sigma \in \mathcal{S}_{n,N}} p_\sigma \det(\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)) \quad (3.22)$$

Recall that the determinant of any square matrix is equal to the product of its eigenvalues. Since the eigenvalues of a Hurwitz matrix in $\mathbb{R}^{n \times n}$ have strictly negative real parts, its determinant will either be strictly positive (when n is even) or strictly negative (when n is odd), but never zero. Thus, using (3.22) in (3.20), we conclude that there must be at least one $\sigma \in \mathcal{S}_{n,N}$ for which $\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$ is not a Hurwitz matrix.

To recapitulate, we have shown so far that if there is no non-zero $\mathbf{v} \succeq \mathbf{0}$ such that $\mathbf{v}^\top \mathbf{A}_i \preceq \mathbf{0}$ for all i , then at least one of the $\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$ matrices has to be non-Hurwitz. However, in order to finish the proof, we need to extend this result to strictly positive \mathbf{v} , as stated in the theorem. So let us assume that there is no common $\mathbf{v} \succ \mathbf{0}$ such that $\mathbf{v}^\top \mathbf{A}_i \prec \mathbf{0}$ for all $i = 1, \dots, N$. If, additionally, there was no $\mathbf{v} \succeq \mathbf{0}$ either such that $\mathbf{v}^\top \mathbf{A}_i \preceq \mathbf{0}$ for all i , then the desired result follows from the above discussion. However, if there was such a $\mathbf{v} \succeq \mathbf{0}$, an additional argument is needed.

Assume that no $\mathbf{v} \succ \mathbf{0}$ satisfies $\mathbf{v}^\top \mathbf{A}_i \prec \mathbf{0}$ for all i . Letting $\mathbf{A}_i(\varepsilon) := \mathbf{A}_i + \varepsilon \mathbf{1}_{n \times n}$ where $\varepsilon > 0$ and $\mathbf{1}_{n \times n}$ is the $n \times n$ matrix of all ones, it then follows that there cannot be a non-zero $\mathbf{v} \succeq \mathbf{0}$ achieving $\mathbf{v}^\top \mathbf{A}_i(\varepsilon) \preceq \mathbf{0}$ for all i . This can be proved by contradiction: Assume there was such a vector $\mathbf{v} \succeq \mathbf{0}$ for which $\mathbf{v}^\top \mathbf{A}_i(\varepsilon) \preceq \mathbf{0}$ for all i and $\varepsilon > 0$. Then

$$\mathbf{v}^\top (\mathbf{A}_i + \varepsilon \mathbf{1}_{n \times n}) \preceq \mathbf{0} \quad (3.23)$$

$$\mathbf{v}^\top \mathbf{A}_i \preceq \mathbf{0} - \varepsilon \mathbf{v}^\top \mathbf{1}_{n \times n} \quad (3.24)$$

$$\mathbf{v}^\top \mathbf{A}_i \prec \mathbf{0} \quad (3.25)$$

for $\varepsilon > 0$ and $i = 1, \dots, N$, which contradicts the first assumption; thus, there is no non-zero $\mathbf{v} \succeq \mathbf{0}$ so that $\mathbf{v}^\top \mathbf{A}_i(\varepsilon) \preceq \mathbf{0}$ for all $i = 1, \dots, N$.

Now, choosing $\varepsilon > 0$ small enough to ensure all $\mathbf{A}_j(\varepsilon)$ are still Hurwitz and Metzler matrices, it follows from our earlier argument that there is at least one $\sigma \in \mathcal{S}_{n,N}$ so that $\mathbf{A}_\sigma(\mathbf{A}_1(\varepsilon), \dots, \mathbf{A}_N(\varepsilon))$ is non-Hurwitz.

Finally consider a sequence of ε_k such that $\varepsilon_k \rightarrow 0$ as $k \rightarrow \infty$ and where the ε_k are small enough so that all $\mathbf{A}_j(\varepsilon_k)$ are still Hurwitz and Metzler matrices. Since these matrices and

thus all $\mathbf{A}_\sigma(\mathbf{A}_1(\varepsilon_k), \dots, \mathbf{A}_N(\varepsilon_k))$ depend continuously on ε_k , it follows for all $\sigma \in \mathcal{S}_{n,N}$ that

$$\mathbf{A}_\sigma(\mathbf{A}_1(\varepsilon_k), \dots, \mathbf{A}_N(\varepsilon_k)) \rightarrow \mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N) \quad \text{as} \quad \varepsilon_k \rightarrow 0 \quad (3.26)$$

And since there is at least one $\sigma \in \mathcal{S}_{n,N}$ for which $\mathbf{A}_\sigma(\mathbf{A}_1(\varepsilon_k), \dots, \mathbf{A}_N(\varepsilon_k))$ is non-Hurwitz this will also be the case for $\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$.

This completes the proof of Theorem 3.2. □

3.4.2 Remarks

Theorem 3.2 states that N positive LTI systems have a common linear co-positive Lyapunov function $v(\mathbf{x}) = \mathbf{v}^\top \mathbf{x}$ if and only if all the $\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$ matrices are Hurwitz matrices, for all $\sigma \in \mathcal{S}_{n,N}$. We recall that in this case any switched system formed with any number of these subsystems would be uniformly asymptotically stable under arbitrary switching.

We note also that if the $\mathbf{A}_i \mathbf{Q}_i$ matrices from Section 3.3 are all square Metzler and Hurwitz matrices, then this Hurwitz condition can also be used to give a solution to the state-restricted switching problem.

A piece of MATLAB[®] code to conveniently check the Hurwitz condition (ii) of Theorem 3.2 is given at the very end of this chapter. Note that this requires the computation of the spectra of N^n matrices of dimension $n \times n$. This may, on a computational level, be significantly more expensive (and possibly even infeasible) compared to the linear program based test described earlier. However, a very recent paper by Narendra and Shorten (2010) provides an efficient, necessary and sufficient test for Hurwitz stability of Metzler matrices. The test involves recursively checking the sign of main diagonal entries of a sequence of lower dimensional matrices that are created by adding two matrices and is thus very inexpensive to perform.

Finally, as stated earlier, the above result may also be derived from the independent, more general results on \mathbf{P} -matrix set by Song *et al.* (1999).

3.4.3 Insights from Hurwitz condition

We stated earlier that analytical results as shown above can lead to more insights into the stability problem as compared to numerical tests. Hence, before extending our results to discrete-time systems, we would like to give an example in support of this claim.

Assume a set of matrices does not pass the stability test given by statement (ii) of Theorem 3.2. In particular, assume that it is the matrix \mathbf{A}_{σ_0} that is not Hurwitz stable, $\sigma_0 \in \mathcal{S}_{n,N}$. If one has some form of control over the entries in the original matrices $\mathbf{A}_1, \dots, \mathbf{A}_N$, what can be done so that \mathbf{A}_{σ_0} may eventually become Hurwitz? Clearly, sufficiently decreasing the entries on the main diagonal and/or the off-diagonal entries will eventually make the matrix become Hurwitz stable. While this is straightforward to show

(see for instance Horn and Johnson, 1991, Chapter 2.5), it is also somewhat intuitive given the fact that \mathbf{A} is Hurwitz if and only if there exists a vector $\mathbf{v} \succ \mathbf{0}$ such that $\mathbf{A}\mathbf{v} \prec \mathbf{0}$, and hence decreasing the non-negative off-diagonals as well as decreasing the negative diagonal elements works toward satisfaction of that inequality. An additional argument is given by the following observations.

Assuming we have some form of control over the matrix entries, another question one may now ask is which matrix element in particular to manipulate first?⁵ In this context, it is useful to note that, by construction, any Metzler matrix \mathbf{A} can be written as

$$\mathbf{A} = \mathbf{P} - \alpha \mathbf{I} \quad \text{with } \mathbf{P} \succeq \mathbf{0} \text{ and for some } \alpha \geq 0 \quad (3.27)$$

and Hurwitz stability of \mathbf{A} is equivalent to $\alpha > \rho(\mathbf{P})$. Thus, if \mathbf{A} is not Hurwitz, $\rho(\mathbf{A})$ is “too large” for the given α . Now, to work toward satisfaction of the inequality, the question is which element $p^{(ij)}$ would have (locally) the biggest impact on $\rho(\mathbf{P})$ in order to decrease it? Assuming \mathbf{A} is irreducible, we can give the following argument. Given the irreducibility assumption, the non-negative matrix \mathbf{P} will also be irreducible. Application of the Perron-Frobenius theorem then guarantees that its Perron root will be algebraically simple (Horn and Johnson, 1985, Theorem 8.4.4) and the corresponding left- and right Perron eigenvectors will be strictly positive. This allows us to apply a standard result (see for instance Stewart, 1973) concerning the partial derivatives of simple eigenvalues of a matrix with respect to the matrix entries: Given some matrix $(p^{(ij)}) = \mathbf{P} \in \mathbb{R}^{n \times n}$ with a simple eigenvalue λ and corresponding normalised left- and right eigenvectors $\boldsymbol{\eta}$ and $\boldsymbol{\xi}$ such that $\boldsymbol{\eta}^\top \boldsymbol{\xi} = 1$, then

$$\frac{\partial \lambda}{\partial p^{(ij)}} = \eta^{(i)} \xi^{(j)} \quad \text{locally, for each } i, j = 1, \dots, n \quad (3.28)$$

This means, in the case where \mathbf{P} is non-negative and irreducible, that the Perron root will always decrease if any element in the matrix is decreased (which is consistent with our earlier remarks). But furthermore, if both Perron eigenvectors can be computed, one immediately knows which entry (i, j) to target first — namely that where $\eta^{(i)} \xi^{(j)}$ is largest.

Application of this result to the original problem and \mathbf{A}_{σ_0} gives an indication which entry $a_k^{(ij)}$ where $k = \sigma_0(j)$ in the original system matrix $\mathbf{A}_k \in \mathbf{A}_1, \dots, \mathbf{A}_N$ to modify first. However, this is only a local result, i. e. having reduced $a_k^{(ij)}$ somewhat may suddenly cause a different entry (potentially in a different system matrix) to have the largest impact on driving \mathbf{A}_{σ_0} toward Hurwitz stability. In fact, the off-diagonal elements can only be reduced to zero but not beyond (in order for the matrix to stay Metzler) — and even if one particular off-diagonal element is reduced to zero the matrix may still not be Hurwitz. Lastly, one may wonder what the impact of reducing $a_k^{(ij)}$ might have on other matrices in \mathbf{A}_σ that include it? Clearly, our earlier observations guarantee that reducing entries in the matrices always makes them “more stable”, in other words decreasing the elements in one matrix will never destroy the Hurwitz stability of other matrices in \mathbf{A}_σ .

⁵ The author is very grateful to Prof. S. Kirkland for pointing him in this direction.

3.5 Discrete-time switched positive systems

As we mentioned earlier, most of the results for continuous-time switched positive linear systems can easily be applied to discrete-time systems as well, Fornasini and Valcher (2010). The discrete time version of the system given in Definition 3.1 on page 35 would be

$$\mathbf{x}(k+1) = \mathbf{A}_{s(\mathbf{x}(k),k)} \mathbf{x}(k) \quad \text{with} \quad \mathbf{x}(k=0) = \mathbf{x}_0 \succeq \mathbf{0} \quad (3.29)$$

where $s : \mathbb{R}^n \times \mathbb{R} \rightarrow \{1, \dots, N\}$ is again some piecewise constant switching function that may or may not depend on the state vector $\mathbf{x}(k)$, and where the system matrices $\mathbf{A}_i \in \mathbb{R}^{n \times n}$ for each $i = 1, \dots, N$ must now be non-negative in order to ensure positivity, and Schur-stable (i.e. all their eigenvalues must lie inside the unit circle).

A linear co-positive Lyapunov function $v(\mathbf{x}) = \mathbf{v}^\top \mathbf{x}$ for such systems would then have to fulfil

$$v(\mathbf{x}) > 0 \quad \text{for all non-zero} \quad \mathbf{x} \succeq \mathbf{0} \quad (3.30)$$

$$v(\mathbf{x}(k+1)) - v(\mathbf{x}(k)) < 0 \quad \text{for all} \quad k \geq 0 \quad \text{and} \quad \mathbf{x}(k) \succ \mathbf{0} \quad (3.31)$$

Clearly, it will be a CLCLF for the switched system (3.29) if (and only if) it is a LCLF for each constituent system, that is if and only if

$$v(\mathbf{A}_i \mathbf{x}) - v(\mathbf{x}) = \mathbf{v}^\top (\mathbf{A}_i - \mathbf{I}) \mathbf{x} \prec \mathbf{0} \quad \text{for all} \quad i = 1, \dots, N \quad \text{and non-zero} \quad \mathbf{x} \succ \mathbf{0}$$

Thus, by letting $\tilde{\mathbf{A}}_i := (\mathbf{A}_i - \mathbf{I})$ for $i = 1, \dots, N$, all our earlier results directly apply to the discrete-time case as well, noting that all $\tilde{\mathbf{A}}_i$ will of course be Metzler (the off-diagonal elements remain non-negative after subtraction of the identity matrix) and Hurwitz (since the spectral radius of the \mathbf{A}_i is strictly less than one, subtracting the identity matrix will shift all eigenvalues into the open left half of the complex plane).

3.6 Examples of usage

While we will give in Chapter 6 an in-depth discussion of an application where our results are used to prove stability of a power control algorithm for wireless networks, we still would like to give a few examples here at this point to illustrate our above results.

3.6.1 Numerical example

As a short example for Theorem 3.2, consider three Metzler and Hurwitz matrices

$$\mathbf{A}_1 = \begin{bmatrix} -12 & 6 & 6 \\ 1 & -10 & 2 \\ 5 & 3 & -10 \end{bmatrix}, \quad \mathbf{A}_2 = \begin{bmatrix} -12 & 4 & 0 \\ 6 & -10 & 9 \\ 4 & 3 & -13 \end{bmatrix}, \quad \mathbf{A}_3 = \begin{bmatrix} -9 & 2 & 8 \\ 6 & -10 & 4 \\ 3 & 0 & -11 \end{bmatrix}$$

It turns out that the $\mathbf{A}_\sigma(\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3)$ are all Hurwitz matrices, for any $\sigma \in \mathcal{S}_{3,3}$; hence a switched positive linear system with these matrices will be uniformly asymptotically

stable under arbitrary switching. If, however, the (3,1)-element of \mathbf{A}_3 is changed from 3 to 5 — note that after change \mathbf{A}_3 is still a Metzler and Hurwitz matrix — then the matrix $\mathbf{A}_{(3,1,3)} = [\mathbf{A}_3^{(1)} \mathbf{A}_1^{(2)} \mathbf{A}_3^{(3)}]$ will have an eigenvalue $\lambda \approx 0.06$ which violates the Hurwitz condition.

3.6.2 Switched positive systems with multiplicative noise

Consider the class of switched positive systems

$$\dot{\mathbf{x}} = \mathbf{A}(t)\mathbf{x}, \quad \mathbf{A}(t) \in \{\mathbf{A}_1, \dots, \mathbf{A}_N\}$$

If all N constituent systems share a co-positive linear Lyapunov function, then it follows that the system

$$\dot{\mathbf{x}} = \mathbf{A}(t)\mathbf{D}(t)\mathbf{x}, \quad \mathbf{A}(t) \in \{\mathbf{A}_1, \dots, \mathbf{A}_N\}$$

where $\mathbf{D}(t) = \text{diag}\{d^{(ii)}(t)\}$ for $i = 1, \dots, n$ is a diagonal matrix, is also exponentially stable, provided that the $d^{(ii)}(t)$ are strictly positive and bounded for all t and i . Systems of this type arise in situations where the state is reset (for example, by quantisation).

3.6.3 Robustness of switched positive systems with channel dependent multiplicative noise

An important class of positive systems is the class that arises in certain networked control problems. Here, the system of interest has the form:

$$\dot{\mathbf{x}} = \mathbf{A}(t, \mathbf{x})\mathbf{x} + [\mathbf{C}_1(t, \mathbf{x}) + \dots + \mathbf{C}_n(t, \mathbf{x})]\mathbf{x}$$

where we assume $(\mathbf{A}(t, \mathbf{x}) + \mathbf{C}_1(t, \mathbf{x}) + \dots + \mathbf{C}_n(t, \mathbf{x}))$ to be always Metzler and Hurwitz (for all t and $\mathbf{x} \in \mathbb{R}_{\geq 0}^n$), where $\mathbf{A}(t, \mathbf{x}) \in \mathbb{R}^{n \times n}$ is Metzler, and where $\mathbf{C}_i(t, \mathbf{x}) \succeq \mathbf{0}$ is an $n \times n$ matrix that describes the communication path from the network states to the i th state; namely it is a matrix of unit rank with only one non-zero row. Further, we allow the network interconnection structure to vary with time between N different configurations, so that $\mathbf{A}(t, \mathbf{x}) \in \{\mathbf{A}_1, \dots, \mathbf{A}_N\}$ and $\mathbf{C}_i(t, \mathbf{x}) \in \{\mathbf{C}_{i1}, \dots, \mathbf{C}_{iN}\}$ for $i = 1, \dots, n$. Our principal result can then be used to give conditions such that this system is exponentially stable. Further, by exploiting simple properties of Metzler matrices (all off-diagonal entries are non-negative), we get the robust stability of the related system:

$$\dot{\mathbf{x}} = \mathbf{A}(t, \mathbf{x})\mathbf{x} + [\mathbf{C}_1(t, \mathbf{x})\mathbf{D}_1(t) + \dots + \mathbf{C}_n(t, \mathbf{x})\mathbf{D}_n(t)]\mathbf{x}$$

where $\mathbf{D}_i(t)$ is a non-negative diagonal matrix whose diagonal entries are strictly positive, but with entries bounded less than one, $i = 1, \dots, n$.

3.7 Conclusion

In this chapter our main results were two necessary and sufficient conditions for the existence of a certain type of Lyapunov function for switched positive linear systems, namely common linear co-positive Lyapunov functions (CLCLF). As we noted earlier, results of this type are very useful as, loosely speaking, existence of such functions implies exponential stability of the overall switched system.

First, we considered the case where the switching rule of the system depends on or is restricted by the system state. More concretely, the state space was assumed to be partitioned by (possibly overlapping) proper convex cones that were each associated with one of the constituent subsystems (but multiple cones could be associated with the same mode). Then, with the system's state being in a given location of the state space, the system could only be in the mode(s) associated with the cone(s) that included that location. For this setting, two necessary and sufficient conditions were given for the existence of CLCLFs: The first one applied to any type of proper convex cone (provided they are convex), while the second one required the cones to be polyhedral. The latter result had the advantage that it directly allowed a simple linear program to be defined whose feasibility was then equivalent to the Lyapunov function existence. However, both cases gave little insight into the overall existence problem and in particular what could be done if the condition was violated.

This led to a second result which applied to the general, arbitrary switching case (in which, of course, the constrained switching cases is included). We showed that existence of CLCLFs is equivalent to a Hurwitz condition on a set of matrices that contains all matrices that can be created by recombining the columns of the original system matrices. Apart from being very general, this algebraic condition had the additional benefit of giving insights into what could be done (and to which subsystem) if the condition was violated.

Finally, after commenting on how our results directly carry over to the discrete time case, three examples were given to illustrate some of the implications of our work.

At this point, we shall leave the domain of switched positive systems for now and consider cooperative control problems in the next two chapters. Although our subsequent results apply to general (not necessarily positive) systems, they may be interpreted as adding an additional feedback loop to a system that switches between different topologies.

* * *

Chapter appendix

The following MATLAB[®] code for easy checking of the Hurwitz condition Theorem 3.2(ii) can be obtained from <http://goo.gl/JM31u>.

```
function [result,perm] = check_hurwitz(Ac)

%Checks Hurwitz condition for all column permutations
%
%   [result,perm] = check_hurwitz(Ac)  where
%       Ac - cell array with the A_j matrices in it
%       result - TRUE (all matrices are Hurwitz), FALSE if not
%       perm - indices of all the permutations of columns for
%             which the condition is violated
%
% Florian Knorn, florian@knorn.org, 14 April 2011

%% Some error catching
if nargin ~= 1
    error('Please provide cell array with matrices');
end
if ~iscell(Ac)
    error('Please provide *cell* array with matrices');
end

%% Initialisations
result = true;
N = length(Ac);
n = length(Ac{1});
perm = [];
maxrho = -1e10;
rhoperm = [];
sigmas = char(zeros(N^n,n));

%% Generate permutations
for i = 1:length(sigmas) % count from 1 to N^n in base N
    sigmas(i,:) = dec2base(i-1,N,n);
end

% Convert strings generated by dec2base back to numbers
sigmas = abs(sigmas) - 47; % numbers
sigmas(sigmas>10) = sigmas(sigmas>10) - 7; % letters
```

```
% Iterate through permutations
for i = 1:length(sigmas)

    % create A_sigma for Hurwitz test
    A_sigma = zeros(n,n);
    for j = 1:n % columns
        temp = Ac{sigmas(i,j)};
        A_sigma(:,j) = temp(:,j);
    end

    % perform Hurwitz test
    rho = max(real(eig(A_sigma)));
    if rho > maxrho
        maxrho = rho;
        rhoperm = sigmas(i,:);
    end
    if max(real(eig(A_sigma)))>0
        result = false;
        perm = [perm;sigmas(i,:)];
    end
end

end
```


Switching and Feedback

This chapter presents a new paradigm for cooperative control and consensus in multi-agent networks with switching topologies. We present and prove stability of three algorithms in this framework that make different assumptions on the overall setting and available information in the network, and provide several simulation results to demonstrate their use.

Chapter contents

- 4.1 Introduction
 - 4.2 Preliminaries
 - 4.3 Algorithm 1: Complete knowledge of system
 - 4.4 Algorithm 2: System only partially known
 - 4.5 Algorithm 3: Dynamics and controllers
 - 4.6 Extension to asynchronous state updates
 - 4.7 Conclusion
 - 4.A Chapter appendix
-

4.1 Introduction

The objective of this chapter is to develop a novel cooperative control scheme that applies to a very general class of problems. It presents joint work with Prof. M. Corless and Prof. R. Shorten and has been published in Knorn *et al.* (2011a,b). On a very abstract level, our overall approach may well be interpreted as a switched system with an added feedback loop.

While the overall setting will be introduced properly in Section 4.2, let us briefly state it here. Consider a system that consists of a large number of interconnected agents (say, a fleet of cars with inter-car communication capabilities) that all have some form of local behaviour (driving speed). This local behaviour has both a local and global effect (CO₂ emissions locally, which result in the total emissions produced by the fleet globally). The objective now is twofold: (i) regulate the global effect or behaviour of the network (such as limit the overall emissions) subject to (ii) some additional local constraint in the form of an inter-agent agreement on some quantity of interest that depends on each node's own

behaviour (equalise emissions between cars for instance which depend on the car’s driving speed). This very general setting is encountered in many more situations, such as:

- cooperative charging of electric vehicles in smart grids
(global constraint: total power available, local constraint: charging time);
- regulation of inflation in economic networks
(global constraint: inflation, local constraint: inter-bank interest rates);
- distributed Quality-of-Service control in cloud computing applications
(global constraint: total bandwidth, local constraint: server load, see Stanojević and Shorten 2009a);
- thermal aware load balancing in large data centres
(global constraint: total work load, local constraint: server temperatures)

Clearly, while cooperative control and the control of networked systems are active topics of research across various disciplines, many fundamental questions remain unanswered. Our objective in this chapter is to provide a new cooperative control paradigm that addresses problems of this type. To do this we exploit the fact that there is usually a non-unique solution to the global regulation problem. In the CO₂ emissions example for instance, the aggregate emissions are just the sum of the individual emissions and hence there is no unique distribution of individual contributions that results in one particular amount of global emissions. Indeed, the key idea will be to use this degree of freedom to solve the global regulation problem while at the same time fulfilling some additional local constraints. For example, in each of the above applications, not only do we seek a certain global behaviour, but we also require some level of inter-agent fairness (in the CO₂ example for instance we wish to regulate CO₂ emissions such that each car is equally polluting).

The idea of inter-agent fairness or “agreement” immediately brings about the notion of *consensus* and *coordination* in multi-agent networks. However, as discussed in the literature review, most of the work in this area assumes bi-directional communications (undirected communication graphs) between agents, often does not cater for time changing topologies in the communication network, and, in many cases, does not consider dynamics involved in state changes (or only very specific types of linear dynamics for specific applications). Most importantly, however, while many consensus schemes will correctly produce an agreement, it appears little work has been done to control and use this consensus value in order to influence the overall network behaviour and achieve some form of “common goal”.

Contributions

In the present work we thus not only attempt to be free of these commonly made assumptions — in particular the graph symmetry assumption upon which much of the underlying

mathematical machinery of the previous work is based — but aim at additionally influencing the consensus value reached in order to meet a global objective. To achieve this, we start with a classic consensus scheme, but add an external input to regulate the consensus value according to a global performance measure that depends on the entirety of the network’s states. Our results will be applicable to a wide range of situations, in particular when only limited knowledge about the network is available.

Structure

The remainder of this chapter is structured as follows: The next section will introduce the problem setting more concretely and define some necessary notation and assumptions. This is followed by three algorithms and convergence proofs thereof (together with a number of comments and simulations) that give a solution to the problem making different assumptions on the problem setting. These form the main contributions of this chapter. Finally, after extending our results to the case of asynchronous communications, we will draw some conclusions, discuss open questions and suggest some future directions.

4.2 Preliminaries

4.2.1 Overall setting and problem statement

We consider the following situation. In a network with $n > 1$ agents or “nodes” and a number of directed communication links¹ that may change over time, each node i has a *physical state* (or just “state”) that it can change, either directly or indirectly through certain dynamics. Furthermore, associated to each node is also what we call a *utility value*: This value directly depends on the node’s physical state and represents some particular quantity of interest that is somehow related to, but usually different from, the physical state. This dependence is given by each node’s *utility function*, which is generally assumed to differ between nodes.

Additionally, we define a certain *global value* that depends directly on *all* the nodes’ physical states; this dependency is given by the *global function*. By suitable means of communication (or decentralised estimation) either all or just some nodes in the network have access to this global value.² Finally, we assume that the agents (locally) share their current utility value through (directed) communication links. This set-up is illustrated in Figure 4.1 on the next page.

¹ This could be due to each node broadcasting information about its state at regular intervals, and other nodes in proximity picking up this broadcast — but these nodes do not necessarily have to communicate back.

² That is, either the global value can be measured or estimated locally by the nodes, or it will be communicated to them by some form of “external” broadcast (for instance sent from a base station that itself can estimate or measure that value).

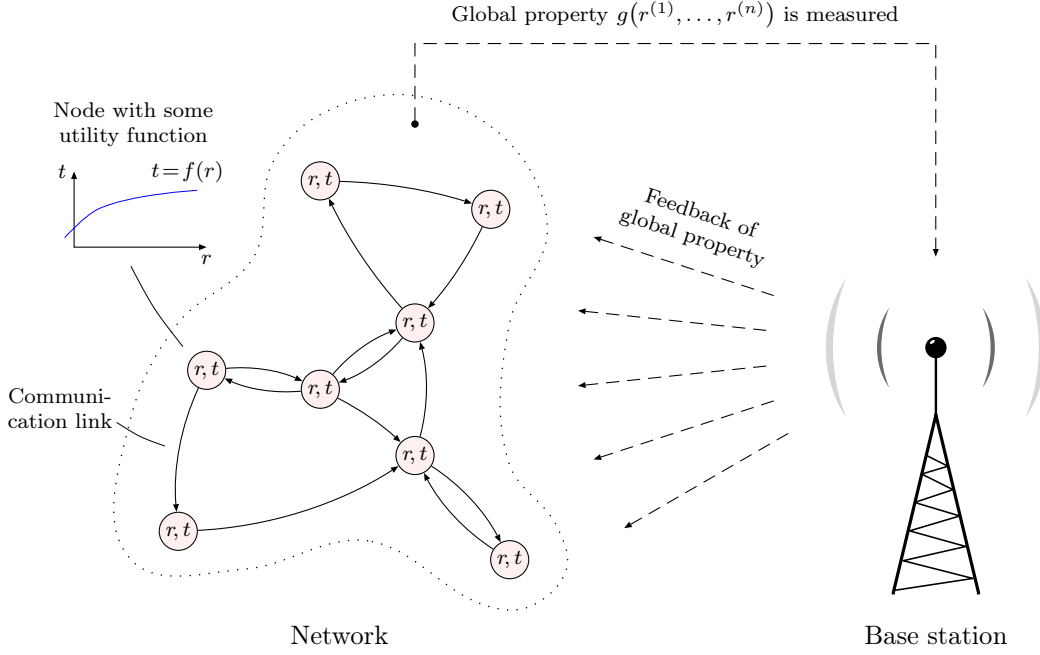


Figure 4.1: Illustration of the basic setting. Each node has a state r and a utility function $f(r)$ associated with it which describes the utility value's dependency on the state. The global property $g(r^{(1)}, \dots, r^{(n)})$ depends on all the network states.

Problem statement

The objective is now for all nodes in the network to reach consensus on their utility values, while also, jointly, driving the global value to a prescribed, “desired” value. This should be achieved in a fully decentralised way, using simple algorithms that will operate in a variety of settings, including non-linear utility functions that are only known approximately, when not all nodes have access to the global value and when the state updates are not necessarily performed synchronously.

Solutions to the problem

To address this problem setting, we will develop and prove convergence of three different decentralised algorithms that are designed to achieve the objectives in three different situations:

- (i) *Physical state:* No dynamics involved, can be changed instantly.
Utility functions: Must be perfectly known.
Global value: All nodes must have knowledge of.

- (ii) *Physical state*: No dynamics involved, can be changed instantly.
Utility functions: Only lower and upper growth bounds must be known.
Global value: Not all nodes must have knowledge of.
- (iii) *Physical state*: Dynamics may be involved in state change.
Utility functions: Only approximate knowledge required, can be filtered values.
Global value: Not all nodes must have knowledge of (but at least one).

Additionally, in each case the underlying communication network can be directed and time varying, both the utility functions as well as the global quantity's dependence on the network states can be non-linear, and the state updates in the network must not necessarily be performed synchronously (in other words, asynchronous communications are covered by our approach as well).

4.2.2 Notation

Our problem setting is best described using typical notions from graph theory, Harary (1969). Let $\mathcal{V} = \{1, \dots, n\}$ be the vertex set of the network and let $\mathcal{A}_k \in \mathcal{V} \times \mathcal{V}$ be the edge set representing the (directed) communication links at time $k = 0, 1, \dots$ between the nodes. We shall always assume that each node can also communicate with itself, i.e. there is always a self-loop on each node. The overall directed graph describing the communication structure of the network at time k is the pair $\mathcal{G}_k = (\mathcal{V}, \mathcal{A}_k)$, where we explicitly assume that the communication links may change over time, but not the node set. The set of (in-)neighbours of node i is called $\mathcal{N}_k^{(i)}$; it contains all the nodes j that can send information to node i (which also includes node i itself), i.e. $\mathcal{N}_k^{(i)} = \{j \mid (j, i) \in \mathcal{A}_k\}$. In a slight abuse of notation we then define the graph's *adjacency matrix* \mathbf{A} as follows: $a_k^{(ij)} = 1$ if $j \in \mathcal{N}_k^{(i)}$, and $a_k^{(ij)} = 0$ otherwise. Strictly speaking, this would be the transpose of the adjacency matrix as defined in the standard literature. Similarly, we say that \mathcal{G}_k is the graph of a non-negative square matrix \mathbf{S}_k if for each $i, j = 1, \dots, n$, $s_k^{(ij)} \neq 0$ if and only if $j \in \mathcal{N}_k^{(i)}$.

The network is called *connected* (in the literature also referred to as *strongly connected*) if there exists a path from every node to every other node in the network, respecting the orientation of the edges. This is the case if and only if the adjacency matrix is *irreducible* (Horn and Johnson, 1985, Theorem 6.2.24). We shall either assume in the following that all networks dealt with are strongly connected, or, if this is not the case, we use the concept of joint connectivity: A set of graphs is called *jointly (strongly) connected* if the union of those graphs is (strongly) connected.³

A matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$ is called *row-stochastic* if all its entries are non-negative and all its row-sums equal one, in other words $p^{(ij)} \geq 0$ and $\mathbf{P}\mathbf{1} = \mathbf{1}$. Similarly, *row sub-stochastic* matrices are defined to be real valued, non-negative matrices whose row-sums are less than

³ The union of a set of graphs on a common vertex set is defined as the graph consisting of that vertex set and whose edge set is the union of the edge sets of the constituent graphs.

or equal to one (but with at least one row-sum strictly less than one). A *strictly row sub-stochastic matrix* is a row sub-stochastic matrix where *all* row-sums are strictly less than one.

Let $r_k^{(i)} \in \mathbb{R}$ be the *physical state* of node i at time k where $k = 0, 1, \dots$, so that \mathbf{r}_k forms the state vector of the network. Node i 's *utility value* $t_k^{(i)} \in \mathbb{R}$ depends on the physical state *via* a continuous and strictly increasing *utility function* $f^{(i)} : \mathbb{R} \rightarrow \mathbb{R}$, that is $t_k^{(i)} = f^{(i)}(r_k^{(i)})$. Further properties of the utility functions (such as invertibility) will be assumed where necessary. Note that for convenience we will often write the utility functions in vector form, i.e. we use $\mathbf{r}_k = \mathbf{f}(\mathbf{t}_k)$ to represent $t_k^{(i)} = f^{(i)}(r_k^{(i)})$ for each i . Furthermore, let $g : \mathbb{R}^n \rightarrow \mathbb{R}$ be a *global function* that depends on all the states, which we assume to be element-wise strictly increasing. Desired values are usually denoted with subscript asterisks, so that, for example, the desired value for the global function is always denoted by g_* . Based on this desired value, the solution to the problem thus consists of a vector \mathbf{r}_* for which $f^{(i)}(r_*^{(i)}) = t_*$ for all i and $g(\mathbf{r}_*) = g_*$.

4.2.3 Growth conditions

Throughout we shall assume that the utility functions and the global function are continuous and satisfy the following growth conditions

Assumption 4.1 (Bounded growth rates)

There are positive constants $\underline{d}^{(i)}, \bar{d}^{(i)}, \underline{h}^{(i)}, \bar{h}^{(i)}$ such that

$$\underline{d}^{(i)} \leq \frac{f^{(i)}(r_a) - f^{(i)}(r_b)}{r_a - r_b} \leq \bar{d}^{(i)} \quad \text{for all } r_a, r_b \in \mathbb{R} \text{ with } r_a \neq r_b \quad (4.1a)$$

$$\underline{h}^{(i)} \leq \frac{g(\mathbf{r} + \Delta \mathbf{r} \mathbf{e}_i) - g(\mathbf{r})}{\Delta r} \leq \bar{h}^{(i)} \quad \text{for all } \mathbf{r} \in \mathbb{R}^n \text{ and all } \Delta r \in \mathbb{R} \text{ with } \Delta r \neq 0 \quad (4.1b)$$

for all $i = 1, \dots, n$.

Loosely speaking, the growth conditions require the utility functions to be strictly increasing with a rate that is bounded away from zero and upper bounded; the global function must also be strictly increasing with a non-zero but also upper bounded rate coordinate-wise.

4.2.4 Feasibility and existence of unique solution

Before presenting our main results we need to first show that indeed a unique solution to the overall regulation problem exists. As we show next, the existence of such a solution is guaranteed by the above growth conditions.

First, we note that the conditions on the continuous utility functions guarantee that they are strictly monotone increasing and unbounded; hence each utility function has a

continuous inverse $\varphi^{(i)}$ so that

$$\varphi^{(i)}(f^{(i)}(r)) = r \quad \text{and} \quad f^{(i)}(\varphi^{(i)}(t)) = t \quad (4.2)$$

for all $t, r \in \mathbb{R}$. Let $\boldsymbol{\varphi}(\mathbf{t}) := [\varphi^{(1)}(t^{(1)}), \dots, \varphi^{(n)}(t^{(n)})]^\top$ denote the inverse of $\mathbf{f}(t)$. Furthermore, with $\mathbf{t}_* = t_* \mathbf{1}$ define

$$\theta(t_*) := g(\boldsymbol{\varphi}(\mathbf{t}_*)) = g(\boldsymbol{\varphi}(t_* \mathbf{1})) \quad (4.3)$$

In order for our problem to have a solution, it is thus necessary and sufficient that the equation $\theta(t_*) = g_*$ has a solution for t_* for all g_* , that is, the function θ is invertible.

When a solution for t_* exists, the solution for the state vector is given by $\mathbf{r}_* = \boldsymbol{\varphi}(t_* \mathbf{1})$. We now show that, as a consequence of the growth conditions, the function θ is indeed invertible. Using the result in Section 4.A.1 on page 80, we obtain that, for any $t_a, t_b \in \mathbb{R}$,

$$\theta(t_a) - \theta(t_b) = g(\boldsymbol{\varphi}(t_a \mathbf{1})) - g(\boldsymbol{\varphi}(t_b \mathbf{1})) \quad (4.4)$$

$$= \sum_{i=1}^n c^{(i)}(t_a - t_b) \quad \text{where} \quad \frac{h^{(i)}}{d^{(i)}} \leq c^{(i)} \leq \frac{\bar{h}^{(i)}}{\underline{d}^{(i)}} \quad (4.5)$$

From this it follows that θ satisfies the growth condition

$$0 < \underline{c} \leq \frac{\theta(t_a) - \theta(t_b)}{t_a - t_b} \leq \bar{c} \quad \text{for all } t_a, t_b \in \mathbb{R} \text{ with } t_a \neq t_b \quad (4.6)$$

where

$$\underline{c} = \sum_{i=1}^n \frac{h^{(i)}}{d^{(i)}} \quad \text{and} \quad \bar{c} = \sum_{i=1}^n \frac{\bar{h}^{(i)}}{\underline{d}^{(i)}} \quad (4.7)$$

Satisfaction of the above growth condition implies that θ is invertible, hence, our problem always has a unique feasible solution.

With all these definitions given we are now ready to derive the main contributions of this chapter. At the heart of each of the algorithms presented in the following sections will be a recursive update law according to which the nodes are to adjust their physical state. We would like to emphasise the fact that these update laws indeed represent a decentralised approach — they only require locally available information from neighbouring nodes, and the global term (which is assumed, ideally, to be estimated in a decentralised fashion as well).

4.3 Algorithm 1: Complete knowledge of system

The first algorithm provides a control law that will be suitable for situations where the utility functions are invertible functions and are perfectly known to the designer. Situations like these are encountered, for instance, in the computer communication networks space,

Stanojević and Shorten (2008). Also, the value of the global function must also be accessible to all nodes at all times.

Before stating it, let us present the following lemma which will simplify the proof of our main result.

Lemma 4.1 (Consensus with common input)

Let $\mathbf{P}_k \in \mathbb{R}^{n \times n}$ be a sequence of matrices taken from a finite set of primitive, row-stochastic matrices with strictly positive main diagonal entries, and $\vartheta(\mathbf{x}_k, k)$ a sequence of real numbers.

If $\mathbf{x}_k = (x_k^{(1)}, \dots, x_k^{(n)})^\top$ evolves for some $\mathbf{x}_{k=0} = \mathbf{x}_0 \in \mathbb{R}^n$ according to

$$\mathbf{x}_{k+1} = \mathbf{P}_k \mathbf{x}_k + \vartheta(\mathbf{x}_k, k) \mathbf{1} \quad (4.8)$$

then the elements of \mathbf{x}_k will approach each other over time, that is

$$\lim_{k \rightarrow \infty} x_k^{(i)} - x_k^{(j)} = 0 \quad (4.9)$$

for all $i, j \in \{1, \dots, n\}$.

Proof For $k \geq 1$, define

$$\tilde{\mathbf{x}}_k := \sigma_k \mathbf{1} \quad \text{where} \quad \sigma_k := \sum_{i=0}^{k-1} \vartheta(\mathbf{x}_i, i) \quad (4.10)$$

Since \mathbf{P}_k is row-stochastic,

$$\mathbf{P}_k \tilde{\mathbf{x}}_k = \mathbf{P}_k [\sigma_k \mathbf{1}] = \sigma_k \mathbf{P}_k \mathbf{1} = \sigma_k \mathbf{1} = \tilde{\mathbf{x}}_k \quad (4.11)$$

Hence

$$\begin{aligned} \tilde{\mathbf{x}}_{k+1} &= \tilde{\mathbf{x}}_k + \vartheta(\mathbf{x}_k, k) \mathbf{1} \\ &= \mathbf{P}_k \tilde{\mathbf{x}}_k + \vartheta(\mathbf{x}_k, k) \mathbf{1} \end{aligned} \quad (4.12)$$

Letting $\mathbf{y}_k = \mathbf{x}_k - \tilde{\mathbf{x}}_k$, it follows from (4.8) and (4.12) that $\mathbf{y}_{k+1} = \mathbf{P}_k \mathbf{y}_k$. Since all the \mathbf{P}_k are taken from a finite set of primitive and row-stochastic matrices, there exists a constant scalar $\bar{\vartheta}$ such that

$$\lim_{k \rightarrow \infty} \mathbf{y}_k = \bar{\vartheta} \mathbf{1} \quad (4.13)$$

see for instance Hartfiel (1998). This means that as $k \rightarrow \infty$, the elements in \mathbf{y}_k approach a common value, $\bar{\vartheta}$. Since $\mathbf{x}_k = \mathbf{y}_k + \sigma_k \mathbf{1}$ the desired result follows.

 \square

4.3.1 Main result

The basic idea of the following algorithm consists of running a classical consensus scheme directly on the utility values with an additional global term added in each node. The actual required state update is then calculated (“reverse engineered”) from these new utility values using the inverse utility function.

Theorem 4.1 (Algorithm 1: Complete knowledge of system)

Consider the standard situation as described in the Notation section and assume that the utility functions $f^{(i)}$ and the global function g are continuous and satisfy the growth condition. Furthermore, assume that each node, using the inverse of its utility function, can calculate its physical state corresponding to a particular utility value.

For any initial condition $\mathbf{r}_{k=0} = \mathbf{r}_0 \in \mathbb{R}^n$, and any sequence of strongly connected communication graphs, suppose that the nodes iteratively update their physical states based on

$$t_{k+1}^{(i)} = t_k^{(i)} + \eta \sum_{j \in \mathcal{N}_k^{(i)}} \left(t_k^{(j)} - t_k^{(i)} \right) + \mu (g_* - g(\mathbf{r}_k)) \quad (4.14a)$$

$$r_{k+1}^{(i)} = \varphi^{(i)}(t_{k+1}^{(i)}) \quad (4.14b)$$

for some

$$0 < \eta < \frac{1}{n-1} \quad \text{and} \quad 0 < \mu < \frac{2}{c} \quad (4.15)$$

Then, the physical state vector \mathbf{r}_k converges asymptotically to \mathbf{r}_ for which $f^{(i)}(r_*^{(i)}) = t_*$ for all i and $g(\mathbf{r}_*) = g_*$.*

Proof The control equation (4.14a) can be expressed as

$$\mathbf{t}_{k+1} = \mathbf{S}_k \mathbf{t}_k + \mu [g_* - g(\varphi(\mathbf{t}_k))] \mathbf{1} \quad (4.16)$$

where

$$s_k^{(ij)} = \begin{cases} 1 - \sum_{j \in \mathcal{N}_k^{(i)}} \eta & \text{if } j = i \\ \eta & \text{if } j \in \mathcal{N}_k^{(i)} \\ 0 & \text{otherwise} \end{cases} \quad (4.17)$$

Clearly \mathbf{S}_k is a row-stochastic matrix. The bounds in (4.15) on η guarantee that, for all i , the elements $s_k^{(ii)}$ and $s_k^{(ij)}$ are positive for $j \in \mathcal{N}_k^{(i)}$. Thus the graph corresponding to \mathbf{S}_k is the (strongly) connected communication graph at time step k ; this implies that \mathbf{S}_k is irreducible. Furthermore, since the main diagonal entries of \mathbf{S}_k are all strictly positive, this matrix is *primitive* (Horn and Johnson, 1985, Lemma 8.5.5). Noting that the number

of strongly connected graphs on n nodes is finite, it follows that all the \mathbf{S}_k matrices are contained in a finite set.

Having shown these properties of the \mathbf{S}_k matrices we can now readily apply Lemma 4.1 on page 58 which guarantees that $t_k^{(i)} - t_k^{(j)} \rightarrow 0$ as k grows. Considering the fact that in any practical implementation of this algorithm quantisation effects will inevitably occur, this implies that the evolution of each utility $t^{(i)}$ will eventually be described by

$$\bar{t}_{k+1} = \bar{t}_k + \underbrace{\mu[g_* - g(\varphi(\bar{t}_k \mathbf{1}))]}_{=: \psi(\bar{t}_k)} \quad (4.18)$$

It is well known that such one-dimensional iterated maps have a unique and globally asymptotically stable fixed point $t_* = \psi(t_*)$ if

$$\left| \frac{\psi(t_a) - \psi(t_b)}{t_a - t_b} \right| \leq \beta < 1 \quad (4.19)$$

for any $t_a, t_b \in \mathbb{R}$ and $t_a \neq t_b$, Hilborn (1994). So let us determine suitable bounds for μ so that the above inequality is satisfied and the system will indeed converge to a fixed point. Considering any $t_a, t_b \in \mathbb{R}$ with $t_a \neq t_b$, we have

$$\psi(t_a) - \psi(t_b) = t_a - t_b - \mu[\theta(t_a) - \theta(t_b)] \quad (4.20)$$

where $\theta(t) = g(\varphi(t \mathbf{1}))$. We have already shown that

$$0 < \underline{c} \leq \frac{\theta(t_a) - \theta(t_b)}{t_a - t_b} \leq \bar{c} \quad (4.21)$$

from which the following bounds can be established

$$1 - \mu\bar{c} \leq \frac{\psi(t_a) - \psi(t_b)}{t_a - t_b} \leq 1 - \mu\underline{c} < 1 \quad (4.22)$$

Thus, condition (4.19) holds if $1 - \mu\bar{c} > -1$, that is, $\mu < 2/\bar{c}$ which is one of the hypotheses of the theorem. Convergence of the one-dimensional system (4.19) to t_* corresponds to all the utility values of all nodes converging to the same value t_* ; since $\psi(t_*) = t_*$ can only be the case if $g(\varphi(t_* \mathbf{1})) = g_*$, we obtain the result that $g(\mathbf{r}_*) = g_*$ where $\mathbf{r}_*^{(i)} = f^{(i)}(t_*)$, i. e. the original system converges to the desired solution.

This concludes our proof of Theorem 4.1. □

Comment The control law (4.14a) has two components: One aimed at achieving consensus on the utility values, the other at regulating the global value. In order to make this control law easier to understand and implement we suggested a rather specific form for the consensus part — it only involves one parameter (the gain η) together with the summation over the differences of utility values. As we stated earlier, the bounds on the gain η are used ensure that this formulation results in primitive, row-stochastic averaging matrices \mathbf{S}_k

so that Lemma 4.1 on page 58 can be applied. Clearly, this specific formulation does not necessarily have to be used, and Theorem 5.1 on page 97 in the next chapter, whose claims are similar to Lemma 4.1 on page 58, would allow for a much broader class of averaging schemes to be employed. \swarrow

4.3.2 Simulations

To produce time varying graphs for our simulations, we made use of random geometric graphs with time varying *connection radii* (or *distance parameters*), see Penrose (2003); Santi (2005). A geometric graph is created by distributing nodes over a defined area (typically, the unit square is used), associating a connection radius to each node i and then connecting it to all the nodes j that lay within node i 's connection radius (which could be thought of as a “broadcast radius”, that is an area within which other nodes j can receive information from node i). In all the examples here, each node's physical state is interpreted as its connection radius,⁴ and thus, as the states change so will the network's topology. All examples use graphs with $n = 25$ nodes.




Figure 4.2: *Simulation of Algorithm 1.*

The global and utility functions used for the simulation of Algorithm 1 were of the quadratic type, see Section 4.A.2 and Figure 4.8 on page 81. For these functions it is

⁴ However, if a state is less than 0 or larger than 1.5, it is interpreted as 0 or 1.5 respectively to determine the graph topology.

straightforward to determine the growth-bounds as required by the theorem and calculate the bounds on the gains μ and η used in the update equation.

Figure 4.2 on the preceding page shows the results for a desired global value of $g_* = 54$, when the network was initialised with a common physical state of $r_0^{(i)} = 0.35$ for $i = 1, \dots, 25$. The subplots show the evolution over time of the value of the global term (with the desired value marked by the dashed line), the physical states and the utility values, respectively. As can be seen, the network quickly reaches consensus on the utility values. The general increase in the physical state values is driven by the, initially, lower than desired global value, which then pushes the global value towards its target value. The physical states (interpreted as the connection radii for the underlying communication graph) remained large enough for the network to be strongly connected in each time step.

In closing, note that the theorem requires a very precise setting where perfect knowledge of the utility functions (and their inverses in particular) is required. Additionally, every node needs to have access to the value of the global term which may not be possible in all applications. In that regard, the algorithm and its generalisation developed in the next section requires weaker assumptions on the setting and thus is relevant to a much larger class of applications.

4.4 Algorithm 2: System only partially known

In this section we present a second, more general algorithm for consensus and cooperative control of a global goal, together with an extension (presented after some simulation results) that allows it to work even in the case where not all nodes have access to the global value. Also, as shown in Section 4.6 on page 75, it can be easily extended further to situations where the communication network is not necessarily strongly connected (which allows the algorithm to handle asynchronous communications, or to tolerate a certain amount of communication failures).

4.4.1 Main result

The implementation of this method only requires limited knowledge of the utility functions as well as the global function. By limited, we mean that only the growth bounds need to be known, not the actual functions itself.

Theorem 4.2 (Algorithm 2: System only partially known)

Consider the standard situation as described in the Notation section and assume that the utility functions $f^{(i)}$ and the global function g are continuous and satisfy the growth condition. For any initial condition $\mathbf{r}_{k=0} = \mathbf{r}_0 \in \mathbb{R}^n$ and any sequence of strongly connected communication graphs, suppose that the nodes iteratively update their physical states based

on

$$r_{k+1}^{(i)} = r_k^{(i)} + \sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} (t_k^{(j)} - t_k^{(i)}) + \mu_k^{(i)} \sigma_k \quad (4.23)$$

where

$$\sigma_k = \begin{cases} g_* - g(\mathbf{r}_{k+1-M}) & \text{if } k+1 \text{ is a multiple of } M := n-1 \\ 0 & \text{otherwise} \end{cases} \quad (4.24)$$

and there exist constants $\varepsilon_1, \varepsilon_2, \underline{\mu}, \bar{\mu} > 0$ such that

$$\eta_k^{(ij)} \geq \varepsilon_1 \quad \text{for } j \in \mathcal{N}_k^{(i)}, \quad \text{and} \quad \sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} \leq \frac{1}{\bar{d}^{(i)}} - \varepsilon_2 \quad (4.25)$$

and

$$0 < \underline{\mu} \leq \mu_k^{(i)} \leq \bar{\mu} \quad (4.26)$$

Then, if $\bar{\mu} > 0$ is sufficiently small, the state vector \mathbf{r}_k converges asymptotically to the vector \mathbf{r}_* for which $f^{(i)}(r_*^{(i)}) = t_*$ for all i and $g(\mathbf{r}_*) = g_*$.

Proof Using the growth properties of the utility functions, we have

$$t_{k+1}^{(i)} - t_k^{(i)} = d_k^{(i)} (r_{k+1}^{(i)} - r_k^{(i)}) \quad \text{where} \quad 0 < \underline{d}^{(i)} \leq d_k^{(i)} \leq \bar{d}^{(i)}. \quad (4.27)$$

Hence, multiplication of update law (4.23) by $d_k^{(i)}$ results in

$$t_{k+1}^{(i)} = t_k^{(i)} + d_k^{(i)} \sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} (t_k^{(j)} - t_k^{(i)}) + d_k^{(i)} \mu_k^{(i)} \sigma_k \quad (4.28)$$

that is,

$$t_{k+1}^{(i)} = s_k^{(ii)} t_k^{(i)} + \sum_{j \in \mathcal{N}_k^{(i)}} s_k^{(ij)} t_k^{(j)} + d_k^{(i)} \mu_k^{(i)} \sigma_k \quad (4.29)$$

where

$$s_k^{(ij)} = \begin{cases} 1 - d_k^{(i)} \sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} & \text{if } j = i \\ d_k^{(i)} \eta_k^{(ij)} & \text{if } j \in \mathcal{N}_k^{(i)} \\ 0 & \text{otherwise} \end{cases} \quad (4.30)$$

Using the result in Section 4.A.1 on page 80 again, we obtain that

$$g_* - g(\mathbf{r}_{k+1-M}) = g(\varphi(\mathbf{t}^*)) - g(\varphi(\mathbf{t}_{k+1-M})) = \sum_{i=1}^n c_k^{(i)} (t_* - t_{k+1-M}^{(i)}) \quad (4.31)$$

where

$$0 < \frac{\underline{h}^{(i)}}{\underline{d}^{(i)}} \leq c_k^{(i)} \leq \frac{\bar{h}^{(i)}}{\bar{d}^{(i)}} \quad (4.32)$$

This allows us to rewrite Equation (4.28) as

$$\mathbf{t}_{k+1} = \begin{cases} \mathbf{S}_k \mathbf{t}_k - \bar{\mu} \mathbf{Q}_k (\mathbf{t}_{k+1-M} - \mathbf{t}_*) & \text{if } k+1 \text{ is a multiple of } M \\ \mathbf{S}_k \mathbf{t}_k & \text{otherwise} \end{cases} \quad (4.33)$$

where, for each $i = 1, \dots, n$,

$$q_k^{(ij)} = \mu_k^{(i)} d_k^{(i)} c_k^{(j)} / \bar{\mu} \quad (4.34)$$

Since $\mathbf{S}_k \mathbf{t}_* = \mathbf{S}_k \mathbf{1} t_* = \mathbf{1} t_* = \mathbf{t}_*$, we subtract \mathbf{t}_* from both sides of (4.33) and define $\Delta \mathbf{t}_k := \mathbf{t}_k - \mathbf{t}_*$ to get the following reformulation of (4.23)

$$\Delta \mathbf{t}_{k+1} = \begin{cases} \mathbf{S}_k \Delta \mathbf{t}_k - \bar{\mu} \mathbf{Q}_k \Delta \mathbf{t}_{k+1-M} & \text{if } k+1 \text{ is a multiple of } M \\ \mathbf{S}_k \Delta \mathbf{t}_k & \text{otherwise} \end{cases} \quad (4.35)$$

We can now use this expression to show that $\Delta \mathbf{t}_k$ converges to the zero vector — which, of course, implies that the states converge to the desired solution of the problem.

If the system starts at $k = 0$ then, after M iterations, Equation (4.35) results in

$$\Delta \mathbf{t}_M = \overbrace{\mathbf{S}_{n-2} \mathbf{S}_{n-3} \dots \mathbf{S}_0}^{=: \bar{\mathbf{S}}_0} \Delta \mathbf{t}_0 - \bar{\mu} \mathbf{Q}_{n-2} \Delta \mathbf{t}_0 \quad (4.36)$$

$$= \underbrace{(\bar{\mathbf{S}}_0 - \bar{\mu} \mathbf{Q}_{n-2})}_{=: \mathbf{Z}_0} \Delta \mathbf{t}_0 \quad (4.37)$$

and after another $n-1$ steps

$$\Delta \mathbf{t}_{2n-2} = \bar{\mathbf{S}}_1 \Delta \mathbf{t}_{n-1} - \bar{\mu} \mathbf{Q}_{2n-3} \Delta \mathbf{t}_{n-1} \quad (4.38)$$

$$= \mathbf{Z}_1 \Delta \mathbf{t}_{n-1} \quad (4.39)$$

In general, for $l = 0, 1, \dots$, we have

$$\Delta \mathbf{t}_{(l+1)M} = \mathbf{Z}_l \Delta \mathbf{t}_{lM} \quad (4.40)$$

where

$$\mathbf{Z}_l = \bar{\mathbf{S}}_l - \bar{\mu} \mathbf{Q}_{(l+1)M-1} \quad \text{and} \quad \bar{\mathbf{S}}_l = \mathbf{S}_{(l+1)(n-1)-1} \dots \mathbf{S}_{l(n-1)} \quad (4.41)$$

The evolution of the $\Delta \mathbf{t}_k$ vectors is thus governed by the product of \mathbf{Z}_l matrices, at which we must hence take a closer look.

To this end, we first show that the $\bar{\mathbf{S}}_l$ matrices are row-stochastic and positive. To do this we first show that the \mathbf{S}_k matrices are primitive and thus *fully indecomposable*

row-stochastic matrices whose non-zero elements are uniformly bounded away from zero. It is clear from (4.30) that \mathbf{S}_k is a row-stochastic matrix. Now note that

$$s_k^{(ii)} = 1 - d_k^{(i)} \sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} \geq 1 - \bar{d}^{(i)} \sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} \geq \bar{d}^{(i)} \varepsilon_2 > 0 \quad (4.42)$$

Also, when $j \in \mathcal{N}_k^{(i)}$ we have

$$s_k^{(ij)} = d_k^{(i)} \eta_k^{(ij)} \geq \underline{d}^{(i)} \varepsilon_1 \geq \underline{d} \varepsilon_1 > 0 \quad (4.43)$$

where

$$\underline{d} := \min_i \{ \underline{d}^{(i)} \} \quad \text{and} \quad \bar{d} := \max_i \{ \bar{d}^{(i)} \} \quad (4.44)$$

The above positive lower bounds on the elements $s_k^{(ij)}$ for $j \in \mathcal{N}_k^{(i)}$ imply that the graph corresponding to \mathbf{S}_k is the (strongly) connected communication graph at time step k . Since the diagonal elements of \mathbf{S}_k are positive this implies that \mathbf{S}_k is primitive (Horn and Johnson, 1985, Lemma 8.5.5). Applying Theorem 2.2 of Brualdi and Liu (1991) we can thus note that \mathbf{S}_k is fully indecomposable for all k . However, a product of the $n-1$ fully indecomposable $n \times n$ matrices yields a strictly positive matrix (Hartfiel, 2002, Corollary 2.5) and hence the $\bar{\mathbf{S}}_l$ are all strictly positive (row-stochastic) matrices.

We now obtain a lower bound on the elements of every $\bar{\mathbf{S}}_l$. It follows from (4.42) and (4.43) that the non-zero elements $s_k^{(ij)}$ of \mathbf{S}_k must satisfy

$$s_k^{(ij)} \geq s_{\min} \quad \text{where} \quad s_{\min} := \min \{ \underline{d} \varepsilon_1, \bar{d}^{(1)} \varepsilon_2, \dots, \bar{d}^{(n)} \varepsilon_2 \} \quad (4.45)$$

Since each element of $\bar{\mathbf{S}}_l$ is the sum of a number of positive terms, where each term is the product of at most M elements of \mathbf{S}_k matrices, and $s_{\min} \leq 1$, we must have

$$\bar{s}_k^{(ij)} \geq (s_{\min})^M =: \bar{s}_{\min} \quad (4.46)$$

for all i, j and k .

Regarding the \mathbf{Q}_k matrices, it follows from (4.34) that, for all k ,

$$0 < \underline{d} \left(\frac{\underline{\mu}}{\bar{\mu}} \right) \frac{\underline{h}^{(j)}}{\bar{d}^{(j)}} \leq q_k^{(ij)} \leq \bar{d} \frac{\bar{h}^{(j)}}{\underline{d}^{(j)}} \quad (4.47)$$

or

$$0 < q_{\min} \leq q_k^{(ij)} \leq q_{\max} \quad (4.48)$$

where

$$q_{\min} := \underline{d} \left(\frac{\underline{\mu}}{\bar{\mu}} \right) \min_i \left\{ \frac{\underline{h}^{(i)}}{\bar{d}^{(i)}} \right\} \quad \text{and} \quad q_{\max} := \bar{d} \max_i \left\{ \frac{\bar{h}^{(i)}}{\underline{d}^{(i)}} \right\} \quad (4.49)$$

Thus, provided

$$0 < \bar{\mu} \leq \frac{\bar{s}_{\min}}{q_{\max}} = \frac{(s_{\min})^M}{q_{\max}} \quad (4.50)$$

every \mathbf{Z}_l matrix will be non-negative; furthermore, since $\bar{\mathbf{S}}_l$ is row-stochastic the row sum of every row of \mathbf{Z}_l will be bounded above by

$$\kappa := 1 - \bar{\mu}q_{\min} < 1 \quad (4.51)$$

This implies that the \mathbf{Z}_l matrices are strictly row *sub*-stochastic and thus satisfy

$$\|\mathbf{Z}_l \Delta \mathbf{t}\|_\infty \leq \kappa \|\Delta \mathbf{t}\|_\infty \quad (4.52)$$

where $\|\cdot\|_\infty$ denotes the usual maximum-norm. It now follows from (4.40) that

$$\|\Delta \mathbf{t}_{(l+1)M}\|_\infty \leq \kappa \|\Delta \mathbf{t}_{lM}\|_\infty \quad (4.53)$$

for all l ; hence

$$\|\Delta \mathbf{t}_{lM}\|_\infty \leq \kappa^l \|\Delta \mathbf{t}_0\|_\infty \quad (4.54)$$

Since each \mathbf{S}_k matrix is non-negative and row-stochastic, it satisfies $\|\mathbf{S}_k \Delta \mathbf{t}\|_\infty \leq \|\Delta \mathbf{t}\|_\infty$; hence

$$\|\Delta \mathbf{t}_k\|_\infty \leq \kappa^l \|\Delta \mathbf{t}_0\|_\infty \quad \text{when} \quad lM \leq k \leq (l+1)M - 1. \quad (4.55)$$

Thus $\Delta \mathbf{t}_k$ converges to zero as k goes to infinity.

This concludes the proof of Theorem 4.2. □

Speaking loosely, an implementation of Algorithm 2 would look as follows. In each time step, a node listens to the utility values broadcast by other nodes in the vicinity, and also broadcasts its own. It then takes the weighted average of these values and updates its own physical state. If additionally the time step is a multiple of $n - 1$, it would also listen for the global broadcast of the global value (or the node estimates it, if this is possible in the application), and stores this value. At the same time, it fetches the global value from $n - 1$ iterations ago and incorporates it in the state update.

Let us now provide some simulation results of this procedure.

4.4.2 Simulations of Algorithm 2

For the simulations of the algorithm based on Theorem 4.2, which are shown in Figure 4.3 on the next page, we used piecewise linear utility functions; the global function was selected to be of affine form, see again Section 4.A.2 and Figure 4.8(b) on page 81. The parameter bounds were chosen within certain bounds based on which the growth-bounds as required by the theorem were derived.

As described in Theorem 4.2, the states only incorporate the value of the global term every $n - 1 = 24$ time steps. These updates are marked by the dashed, vertical lines in the second subplot.




Figure 4.3: *Simulation of Algorithm 2.*

While in each time step the averaging scheme pulls together the utility values, each update with the global term pulls them apart again (but brings the global value closer to its desired value). As the targeted value is approached, however, the influence of the global term gets smaller and smaller and eventually the averaging scheme brings a “lasting” consensus to the utility values, at a point where the global term has reached the desired value.

4.4.3 Extension when access to the global term is limited

The previous result assumes that all nodes always have access to the global value when it is needed. In order to make our results also relevant to applications where this assumption may not always be practical or possible to guarantee (for instance in the presence of communication failures), we provide the following corollary to Theorem 4.2. It relaxes the assumptions to the more general setting where not all nodes have access to the global term. In fact, it is sufficient for only one single node to have access to the global value. This “special” node could for instance be placed in a strategic position where it can either measure or determine the global value itself, or receive from an external source (“bridge node”).

To model this more general scenario, consider any time step k where the global term $g(\mathbf{r}_{k+1-M})$ is needed and let $\mathcal{I}_k \subseteq \{1, \dots, n\}$ be the non-empty set of nodes which incorpo-

rate the global term in their state update at time k . Then, recalling the original algorithm in (4.23), the more general algorithm under consideration is modelled by

$$r_{k+1}^{(i)} = r_k^{(i)} + \sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} (t_k^{(j)} - t_k^{(i)}) + \mu_k^{(i)} \sigma_k^{(i)} \quad (4.56)$$

where

$$\sigma_k^{(i)} = \begin{cases} g_* - g(\mathbf{r}_{k+1-M}) & \text{if } k+1 \text{ is a multiple of } M = n-1 \text{ and } i \in \mathcal{I}_k \\ 0 & \text{otherwise} \end{cases} \quad (4.57)$$

for all $k = 0, 1, \dots$. We have now the following result.

Corollary 4.1 (Restricted access to global term) _____

The results of Theorem 4.2 on page 62 still hold when not all (but at least one) node includes the global term in the state update whenever it is required. _____

Proof The proof of the corollary is almost identical to that of Theorem 4.2; only some small modifications are needed. Proceeding as before, the algorithm can still be described by (4.35) where \mathbf{S}_k is the same as before and the rows of \mathbf{Q}_k corresponding to the nodes which update with the global term at k are the same as before; however the rows of \mathbf{Q}_k corresponding to those nodes which cannot incorporate the global term at k are zero. Thus \mathbf{Q}_k is not necessarily strictly positive. However, since the assumptions of the corollary guarantee at least one positive row, the \mathbf{Z}_l matrices defined in (4.41) will still be row sub-stochastic but not necessarily *strictly* row sub-stochastic (as they were under the hypotheses of Theorem 4.2).

However, as we show now, products of the form $\mathbf{Z}_{l+1}\mathbf{Z}_l$ are strictly row sub-stochastic. To this end suppose that $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times n}$ are positive, row-stochastic or row sub-stochastic matrices, and at least one row-sum in \mathbf{B} is strictly less than one. We show that then the product \mathbf{AB} must be strictly row sub-stochastic. Let $\mathbf{b} = \mathbf{B}\mathbf{1}$ and $\mathbf{w} = \mathbf{AB}\mathbf{1}$ be the vectors containing the row-sums of \mathbf{B} and the product \mathbf{AB} respectively. Since \mathbf{B} is row-stochastic or sub-stochastic, we have $b^{(j)} \leq 1$ for all j and, by assumption, there is at least one j_0 for which $b^{(j_0)} < 1$. Since $\mathbf{w} = \mathbf{AB}\mathbf{1} = \mathbf{Ab}$, it follows from the definition of the matrix product that for each $i = 1, \dots, n$, $w^{(i)} = \sum_{j=1}^n a^{(ij)}b^{(j)}$. As all elements in \mathbf{A} are positive, $\sum_{i=1}^n a^{(ij)} \leq 1$, $b^{(j)} \leq 1$ for all j and $b^{(j_0)} < 1$, we must have

$$\begin{aligned} w^{(i)} &= \sum_{j=1}^n a^{(ij)} - \sum_{j=1}^n a^{(ij)}(1 - b^{(j)}) \\ &\leq 1 - a^{(ij_0)}(1 - b^{(j_0)}) \\ &< 1 \end{aligned} \quad (4.58)$$

In other words, the product \mathbf{AB} is strictly row sub-stochastic.

Using (4.40) we obtain that for $l = 0, 2, 4, \dots$,

$$\Delta \mathbf{t}_{(l+2)M} = (\mathbf{Z}_{l+1} \mathbf{Z}_l) \Delta \mathbf{t}_{lM} \quad (4.59)$$

Since the elements of each matrix \mathbf{Z}_l are uniformly bounded away from zero and each matrix has at least one row whose sum is uniformly bounded above by a number less than one, it follows that the matrix product $\mathbf{Z}_{l+1} \mathbf{Z}_l$ is positive, strictly row sub-stochastic with row sums uniformly bounded above by some $\kappa < 1$. As demonstrated in proof of Theorem 4.2, one can now prove again convergence of $\Delta \mathbf{t}_k$ to zero.

This concludes the proof of Corollary 4.1. □

Simulations of this extension to Algorithm 2 where only a small number of nodes have access to the global value are given in Section 4.6.3 on page 77.

We shall now move on to our third main result that makes even less assumptions on the utility functions.

4.5 Algorithm 3: Dynamics and controllers

While the third proposed algorithm shares some similarities with the previous two, it differs conceptually from them in that it is more abstract, modular and allows different nodes to use different controllers to adjust their physical state. In fact, the combination of controller and utility function (the “control loop”) may even have a dynamic behaviour, and can be heterogeneous (that is, different nodes may use completely different controller types or utility functions).

The following approach can be interpreted as “decoupling” the adjusting of the physical states (control action) from the iterative calculation of “target utility values” that are designed to converge to the actual solution of the problem. As is the case with the previous two algorithms, this algorithm is also intended to be implemented in a fully decentralised way.

Concretely, we envisage the following structure: First, in a distributed averaging step the current utility values are averaged using some distributed averaging scheme. To this, if k is a multiple of $M = n - 1$, a term $\mu \sigma_k^{(i)}$ which is proportional to the error between desired global value and actual global value is added. This yields the *target utility values* $\tilde{t}^{(i)}$:

$$\tilde{t}_{k+1}^{(i)} = a_k^{(ii)} t_k^{(i)} + \sum_{j \in \mathcal{N}_k^{(i)}} a_k^{(ij)} t_k^{(j)} + \mu \sigma_k^{(i)} \quad (4.60)$$

where $\mathbf{A}_k = (a^{(ij)})$ represents the distributed averaging scheme, $\sigma_k^{(i)}$ is as defined in (4.57) for Algorithm 2 and $\mu > 0$ is a sufficiently small gain which is to be determined.

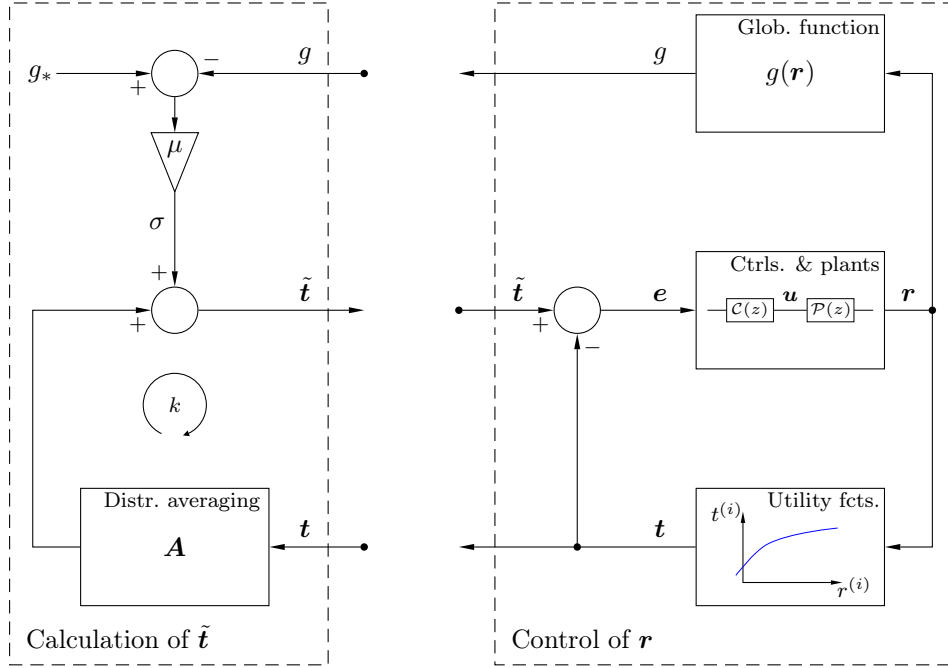


Figure 4.4: Illustration of the interplay of calculation of the target utility values \tilde{t} and the control action to adjust the physical states r accordingly.

Each node then passes its target utility value to its controller, which (over a certain finite time span) manipulates the physical state $r^{(i)}$ in order to drive the node's utility value toward its target value. After that control action, new target values will be calculated based on the resulting new utility values as well as the value of the global function, and so on. This interplay of calculating the target values and then adjusting the states accordingly is shown in Figure 4.4 above.

To leave this third approach as modular as possible, we will not specify any specific averaging scheme or controller type. All that will be required for convergence is that it must be possible to express the averaging scheme as multiplication by row-stochastic matrices with non-zero entries uniformly bounded away from zero, and that the controllers reduce the control error to within some specified range.

4.5.1 Main result

As for the previous two algorithms the question is again: Does there exist a gain μ such that the resulting system is stable and converges to the desired solution?

Theorem 4.3 (Algorithm 3: Dynamics and controllers)

Consider the standard situation as described in the Notation section and assume that the utility functions $f^{(i)}$ and the global function g are continuous and satisfy the growth

condition. Suppose that the communication structure of the network allows it to run a distributed averaging scheme on the utility values. Furthermore, each node is assumed to use a controller that is designed to adjust the node's physical state in such a way as to drive its utility value towards the target utility value.

If the averaging scheme can be represented in each time step as a non-negative row-stochastic matrix \mathbf{A}_k whose graph is strongly connected and with all non-zero elements uniformly bounded away from zero by some $\gamma > 0$, and if the controllers guarantee

$$\underline{\alpha}_i(t_k^{(i)} - \tilde{t}_{k+1}^{(i)}) \leq t_{k+1}^{(i)} - \tilde{t}_{k+1}^{(i)} \leq \bar{\alpha}_i(t_k^{(i)} - \tilde{t}_{k+1}^{(i)}) \quad (4.61)$$

in each control phase for some constant $\underline{\alpha}^{(i)}, \bar{\alpha}^{(i)}$ which satisfy

$$-\gamma/(1-\gamma) < \underline{\alpha}^{(i)} \leq \bar{\alpha}^{(i)} < 1 \quad (4.62)$$

then a positive gain μ can be found for any initial condition $\mathbf{r}_{k=0} = \mathbf{r}_0$ so that the system converges to $\mathbf{t}_k \rightarrow t_* \mathbf{1}$ and $g(\mathbf{r}_k) \rightarrow g_*$.

Proof We will show that any algorithm under consideration here can be reduced to one considered in Corollary 4.1. Satisfaction of the inequalities in (4.61) is equivalent to writing

$$t_{k+1}^{(i)} - \tilde{t}_{k+1}^{(i)} = \beta_k^{(i)}(t_k^{(i)} - \tilde{t}_{k+1}^{(i)}) \quad \text{with} \quad \underline{\alpha}^{(i)} \leq \beta_k^{(i)} \leq \bar{\alpha}^{(i)} \quad (4.63)$$

that is,

$$t_{k+1}^{(i)} = t_k^{(i)} + (1 - \beta_k^{(i)})(\tilde{t}_{k+1}^{(i)} - t_k^{(i)}) \quad (4.64)$$

Recall that

$$\tilde{t}_{k+1}^{(i)} = a_k^{(ii)} t_k^{(i)} + \sum_{j \in \mathcal{N}_k^{(i)}} a_k^{(ij)} t_k^{(j)} + \mu \sigma_k^{(i)} \quad (4.65)$$

Since \mathbf{A}_k is row-stochastic, we must have $a_k^{(ii)} = 1 - \sum_{j \in \mathcal{N}_k^{(i)}} a_k^{(ij)}$ and hence

$$\tilde{t}_{k+1}^{(i)} - t_k^{(i)} = \sum_{j \in \mathcal{N}_k^{(i)}} a_k^{(ij)} (t_k^{(j)} - t_k^{(i)}) + \mu \sigma_k^{(i)} \quad (4.66)$$

Recalling (4.64) now results in

$$t_{k+1}^{(i)} = t_k^{(i)} + \sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} (t_k^{(j)} - t_k^{(i)}) + \mu_k^{(i)} \sigma_k^{(i)} \quad (4.67)$$

where

$$\eta_k^{(ij)} = (1 - \beta_k^{(i)}) a_k^{(ij)} \quad \text{and} \quad \mu_k^{(i)} = (1 - \beta_k^{(i)}) \mu \quad (4.68)$$

Thus the algorithm is an example of those considered in Corollary 4.1. We now show that the hypotheses of Corollary 4.1 hold. First note that $\eta_k^{(ij)} \geq (1 - \bar{\alpha}^{(i)})\gamma > 0$; hence

$$\eta_k^{(ij)} \geq \varepsilon_1 \quad \text{where} \quad \varepsilon_1 = \gamma \min_i \left\{ (1 - \bar{\alpha}^{(i)}) \right\} > 0 \quad (4.69)$$

We also note that

$$\sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} = (1 - \beta_k^{(i)}) \sum_{j \in \mathcal{N}_k^{(i)}} a_k^{(ij)} \quad (4.70)$$

$$= \underbrace{(1 - \beta_k^{(i)})}_{\leq 1 - \underline{\alpha}^{(i)}} \underbrace{(1 - a_k^{(ii)})}_{\leq 1 - \gamma} \quad (4.71)$$

$$\leq 1 - [\gamma + \underline{\alpha}^{(i)}(1 - \gamma)] \quad (4.72)$$

Since $\gamma + \underline{\alpha}^{(i)}(1 - \gamma) > 0$ for $j \in \mathcal{N}_k^{(i)}$, we obtain the desired result that

$$\sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} \leq 1 - \varepsilon_2 \quad \text{where} \quad \varepsilon_2 = \gamma + (1 - \gamma) \min_i \left\{ \underline{\alpha}^{(i)} \right\} > 0 \quad (4.73)$$

We also obtain that

$$0 < \underline{\mu} \leq \mu_k^{(i)} \leq \bar{\mu} \quad (4.74)$$

where

$$\underline{\mu} = \mu \min_i \left\{ 1 - \bar{\alpha}^{(i)} \right\} \quad \text{and} \quad \bar{\mu} = \mu \max_i \left\{ 1 - \underline{\alpha}^{(i)} \right\} \quad (4.75)$$

So, clearly, $\bar{\mu}$ can be made sufficiently small by choosing μ sufficiently small. Application of Corollary 4.1 concludes the proof of 4.3. \square

Comment It is easy to see that the lower bound in (4.62) is automatically satisfied if the controllers are designed to produce no overshoot. By “no overshoot” we mean specifically that during each control phase the utility values never exceed the target values, in other words if for instance $t_k^{(i)} < \tilde{t}_{k+1}^{(i)}$ then the utility value during that control phase will always be less than or equal to $\tilde{t}_{k+1}^{(i)}$. \swarrow

4.5.2 Simulations of Algorithm 3

The set-up used for our simulations of the third algorithm was the following. In a network on $n = 10$ nodes the global functions were again of affine type (as for the simulations of Algorithm 1), the utility functions, in turn, were of quadratic type (as for Algorithm 2).

The averaging scheme in this example was based on random, strongly connected row-stochastic matrices with non-zero entries uniformly bounded below by $\gamma = 0.02$.

The controllers used were discrete-time implementations of PID and PI controllers,⁵ randomly assigned to nodes, Visioli (2006). For both controller types the parameters were tuned as to guarantee that the resulting closed loop system would not produce any overshoot. The gains in the PI controllers were intentionally reduced somewhat in order to produce a slightly slower step response and increase the heterogeneity between the controllers.




Figure 4.5: Step-response of the closed loop control part in the simulation of Algorithm 3.

As for the plants (that is, the physical state updates) we chose first order low-pass filters (see for instance Oppenheim *et al.*, 1996) with randomly chosen smoothing parameter $\zeta \in [0.55, 0.85]$ to simulate a system where the physical state cannot be changed instantly.⁶

To illustrate the behaviour of the resulting controller-plant combination (together with the non-linear utility functions), a step response of the closed loop system is shown in Figure 4.5 above: The system was initialised with a physical state distribution such that all the utility values would be equal. At $k_c = 10$ the target utility values were then set to $\tilde{t}^{(i)} = 4.5$. While the first two subplots showing the global value and physical states are not of particular interest here, the third subplot clearly reveals the two “groups” of nodes — those with the slower PI controllers and those with the faster PID controllers. At $k_c = 30$

⁵ For easier implementation, we used the “velocity formulation”, that is the output of each controller is calculated recursively with: $u_k^{(i)} = u_{k-1}^{(i)} + k_p(e_k^{(i)} - e_{k-1}^{(i)}) + k_i e_k^{(i)} + k_d(e_k^{(i)} - 2e_{k-1}^{(i)} + e_{k-2}^{(i)})$. For the PID-controllers, the parameters were set to $k_p = 0.10$, $k_i = 0.09$ and $k_d = 0.03$; for the PI-controllers in turn, $k_p = 0.02$, $k_i = 0.05$ and $k_d = 0$.

⁶ Specifically, the new states were calculated as $r_{k+1}^{(i)} = \zeta^{(i)} u_k^{(i)} + (1 - \zeta^{(i)}) r_k^{(i)}$.

(that is, after 20 control iterations), the error between actual utility value and target value relative to the initial value is less than 0.1% for each node.

While this observation does not guarantee that the control error is less than 0.1% at the end of *every* control phase (since the system is not necessarily in steady-state at the beginning of each control phase) it is still reasonable to assume that the error is reduced sufficiently in order to guarantee the bounds (4.62).

This closed loop based on 20 control iterations was then also used in the actual simulation of a system operating according to Algorithm 3, shown in Figure 4.6 below.

The dashed vertical lines in the third subplot indicate each time a new target utility value was calculated. The global term was incorporated every $(n - 1) \cdot 20 = 180$ time steps. Again, consensus is reached on the utility values and the global term reaches its target value of $g_* = 0.44$ as desired.




Figure 4.6: Simulation of Algorithm 3.

To round off this chapter, we note that until now we have only considered networks where the state updates are all performed in a synchronised fashion. That is, for a given time step k , the nodes first exchanged all the relevant state information with each other, and then, jointly, performed the update based on the state information at time k to reach the new state value at $k + 1$. However, this perfectly synchronised way of performing the updates may not always be easy to implement, or even guarantee at all. The next and final section of results in this chapter is to remedy that situation.

4.6 Extension to asynchronous state updates

We now extend our above results to asynchronous communications and state updates by no longer requiring the communication graphs representing the information flow in the network to be strongly connected in each time step (as above), but rather only *jointly* strongly connected over time, with a fixed and constant time horizon $m \geq 1$. In other words, it is only required that the union of any m consecutive graphs taken from that sequence must yield a strongly connected graph. That way, the communication between nodes can be “staggered out”, with nodes updating their state right after they have received information from a neighbour, rather than having to wait until they have received the states from all their neighbours and until all the other nodes are also “ready” to perform the (synchronised) update.

In each of our three results above, the update equations (or their transformed versions in the proofs) contain a consensus term based on row-stochastic and primitive matrices. In case of asynchronous updates, these matrices would also be row-stochastic, but not necessarily primitive. Rather, they would contain a number of rows that only have a 1 in the main diagonal entry and 0 everywhere else (corresponding to nodes that have not received any state information from any other nodes). The key idea of the following extension is that non-zero elements in these matrices do not “get lost” (thanks to the positive main diagonals); only new non-zero entries can appear. Hence, intuitively, all one needs to do is “wait long enough” until eventually these matrices become primitive. This is laid out in detail in the following sub-sections, with which we shall close this chapter.

4.6.1 Asynchronous version of Algorithm 1

Corollary 4.2 (Algorithm 1, asynchronous updates) _____

The results of Theorem 4.1 on page 59 still hold if the sequence of communication graphs is jointly strongly connected over some finite and constant time horizon $m \geq 1$. —

Proof Recall that Algorithm 1 given by Theorem 4.1 can be written as

$$\mathbf{t}_{k+1} = \mathbf{S}_k \mathbf{t}_k + \mu [g^* - g(\mathbf{r}_k)] \mathbf{1} \quad (4.76)$$

The proof of that theorem relied on the convergence result given by Lemma 4.1, the proof of which in turn required a sequence of *primitive* matrices so that Theorem 1.9 of Hartfiel (1998) could be used.

Now, in the case of only jointly strongly connected graphs, primitivity of individual \mathbf{S}_k matrix cannot be guaranteed. Rather, we need to interpret the product of the \mathbf{S}_k matrices as *blocks* of m matrices multiplied together, since only these “sub-products” yield primitive, row-stochastic matrices (thanks to the main diagonal entries in each matrix \mathbf{S}_k being strictly positive). Additionally, since there are only finitely many possible graph

topologies on n nodes, there can only be finitely many different m -blocks of \mathbf{S}_k matrices, which implies a uniform, non-zero lower bound on the non-zero matrix elements in all these m -blocks. Both properties make the use of Lemma 4.1 possible again to show that, ultimately, system (4.16) converges to a scalar system. The rest of the proof then follows again the lines of the proof of Theorem 4.1. \square

4.6.2 Asynchronous versions of Algorithms 2 and 3

Corollary 4.3 (Algorithms 2 and 3, asynchronous updates) \square

The results of Theorem 4.2 and Corollary 4.1 still hold if the sequence of communication graphs is jointly strongly connected over some finite and constant time horizon $m \geq 1$, provided $M = n - 1$ is replaced with $M = m(n - 1)$. \square

Proof Only a small modification to the proof of Theorem 4.2 and Corollary 4.1 is needed to show the above result. For $\tilde{k} = 0, 1, \dots$, let

$$\tilde{\mathbf{S}}_{\tilde{k}} = \mathbf{S}_{(m+1)\tilde{k}-1} \cdots \mathbf{S}_{\tilde{k}m} \quad (4.77)$$

Since all the \mathbf{S}_k matrices are non-negative row-stochastic matrices with strictly positive diagonal elements, each matrix $\tilde{\mathbf{S}}_{\tilde{k}}$ is row-stochastic, has positive diagonal elements and its graph corresponds to the collection of communication graphs from time step $\tilde{k}m$ to $(m+1)\tilde{k} - 1$. As any collection of m consecutive graphs is assumed to be jointly strongly connected, it follows that $\tilde{\mathbf{S}}_{\tilde{k}}$ is irreducible, and since it has positive diagonal elements, it is primitive and thus fully indecomposable. The algorithm under consideration still satisfies (4.40) where

$$\bar{\mathbf{S}}_l = \mathbf{S}_{(l+1)M-1} \cdots \mathbf{S}_{lM}; \quad (4.78)$$

However, here $M = m(n-1)$. Thus,

$$\bar{\mathbf{S}}_l = \tilde{\mathbf{S}}_{(l+1)(n-1)-1} \cdots \tilde{\mathbf{S}}_{l(n-1)} \quad (4.79)$$

Having established the above properties of the $\tilde{\mathbf{S}}_{\tilde{k}}$ matrices, the remainder of the proof follows Theorem 4.2 or Corollary 4.1. \square

Comment The generalised forms of Algorithm 2 given by Corollaries 4.1 and 4.3 are designed to tolerate certain communication problems. In the case of Corollary 4.1 this robustness is achieved at the cost of very small gains $\mu^{(i)}$ on the global term; see (4.50). Observing (4.51) and (4.55) it is not difficult to see that smaller gains produce slower convergence.

However, various simulations using sufficiently “general” graphs (rather than pathological cases like the directed n -cycle) have shown that those gains can, in fact, be set

significantly larger than required by the theoretical results above, which suggests that these bounds are loose and may be improved on. \swarrow

Our third algorithm can also be modified to accommodate for asynchronous communications in the same manner as described above; simply let $M = m(n - 1)$.

4.6.3 Simulations of the extensions of Algorithm 2

The communication graphs in the previous two simulations were, by design, all strongly connected. For the simulations of the modification of the second algorithm as presented in the corollary, we also used state-dependant disc graphs, but randomly removed, in each time step, a number of edges in order to deliberately disconnect the graphs. The amount of edges removed (in average 75% of the edges), however, was experimentally chosen in order to guarantee (almost always) that every set of $m = 3$ consecutive graphs would form a jointly connected graph, as required by the corollary. Thus, the updates using the global term were performed only every $3(n - 1) = 72$ time steps.




Figure 4.7: *Simulation of Algorithm 2 (corollary).*

Additionally, we also randomly prevented nodes from accessing the value of the global term (in average, only 25% of the nodes were allowed to use the global term during at each global term update step).

The results from the simulation under these harder conditions are shown in Figure 4.7 above and closely resemble that of the previous case. Due to the less frequent updates,

however, convergence to the desired global value takes much longer but is achieved nonetheless.

4.7 Conclusion

Consensus problems have attracted a large amount of attention in recent years. The present chapter's contributions in that area are three fully decentralised cooperative control algorithms that not only allow a network to reach consensus either directly or indirectly (that is, with or without utility functions involved), but also enable the nodes in the network to cooperate and achieve a global, common goal that depends on the aggregate behaviour of the network.

Our first result concerned the well-controlled case where the utility functions and their inverses are perfectly known *a priori*. The nodes then use the inverse utility functions to calculate the state updates.

The second contribution consisted of an algorithm that requires less precise knowledge of the problem setting and involved functions. All that needed to be known were upper and lower bounds on the growth rates of the global- and utility functions, but not the functions themselves. Also, through Corollary 4.1, we allowed for an even broader class of applications where not all nodes need to have access to the global value.

Our third piece of work took a somewhat different approach. The idea consisted of decoupling the adjusting of the physical state from the iterative calculation of desired values for the utility values. This enabled us to cater for networks where the state cannot change instantly, where only filtered versions of the state are available, but, probably most importantly, where different nodes may have completely different dynamics and controllers. The key property required for convergence in these networks was that the controllers must be designed so that they drive the physical states / utility values (in finite time) to within a certain range of the calculated target utility values.

Each of the three algorithms was accompanied by simulation results that demonstrated the effectiveness of our approach, and they were then extended to the case of asynchronous communications and state updates.

Applications for each of the three algorithms can be found for instance in the computer communication networks space (Algorithm 1, Stanojević and Shorten, 2008), emissions control of vehicles (Algorithm 2, see Chapter 6) or group coordination of mobile agents (Algorithm 3, Olfati-Saber, 2006).

Limitations

While the theoretical contributions of this chapter may well present a new paradigm for cooperative control, there are a number of limitations that should be resolved especially for practical applications. As we mentioned before, the gains required in our proofs are

much too small for any practical purpose. Since they are of very conservative nature, it should certainly be possible to improve on them. However, this may involve different mathematical approaches such as directly treating the problem as a switched system and subsequently searching for (common) Lyapunov functions.

Before moving on, we recall again that our general assumption was that all nodes (or at least one node) have access to the global term — typically provided to the network through some external entity that is able to determine, measure or estimate this value. However, there are situations where no such external entity may be available, feasible or even desirable (as it would constitute a single point of failure). To avoid such problems, the nodes would have to estimate the global property themselves. This would typically have to be done conjointly and in a distributed way in order to be more robust and, potentially, to also average out localised phenomena (it would, for instance, not be very accurate to measure the CO₂ levels in only a single location in the city — if a strongly polluting lorry had its engine running next to the CO₂ sensor then the measurement would clearly be biased and not representative for the city as a whole).

The next chapter will focus on one such application where it is actually possible for the nodes to estimate the global property on their own. In some sense, the approach we shall present next may be seen as a special instance of Algorithm 1 of this chapter.

* * *

4.A Chapter appendix

4.A.1 An expression for the global term

This proof has been moved here in order to improve the flow of the original chapter.

Given global and utility functions which satisfy the growth conditions, we show here that, for any $\mathbf{t}, \mathbf{t}_* \in \mathbb{R}^n$,

$$g(\varphi(\mathbf{t}_*)) - g(\varphi(\mathbf{t})) = \sum_{i=1}^n c^{(i)} (t_*^{(i)} - t^{(i)}) \quad \text{where} \quad \frac{\underline{h}^{(i)}}{\underline{d}^{(i)}} \leq c^{(i)} \leq \frac{\bar{h}^{(i)}}{\bar{d}^{(i)}} \quad (4.80)$$

Letting $\mathbf{r} = \varphi(\mathbf{t})$ and $\mathbf{r}_* = \varphi(\mathbf{t}_*)$, we start by showing that

$$\Delta g := g(\varphi(\mathbf{t}_*)) - g(\varphi(\mathbf{t})) = g(\mathbf{r}_*) - g(\mathbf{r}) \quad (4.81)$$

can be expressed as

$$\Delta g = \sum_{i=1}^n \tilde{c}^{(i)} \Delta r^{(i)} \quad \text{where} \quad \underline{h}^{(i)} \leq \tilde{c}^{(i)} \leq \bar{h}^{(i)} \quad (4.82)$$

and $\Delta r^{(i)} = r_*^{(i)} - r^{(i)}$. The change Δg corresponds to the change of the value of the global function when moving from \mathbf{r} to \mathbf{r}_* . Now, instead of going “directly” from \mathbf{r} to \mathbf{r}_* we can also reach \mathbf{r}_* by only changing one coordinate at a time, that is we basically break up the “cumulative change” Δg into the changes caused by moving along each coordinate. To express this mathematically, we recursively define the vectors $\mathbf{r}_0, \dots, \mathbf{r}_n$ by

$$\mathbf{r}_0 = \mathbf{r} \quad \text{and} \quad \mathbf{r}_i = \mathbf{r}_{i-1} + \Delta r^{(i)} \mathbf{e}_i \quad \text{for } i = 1, \dots, n. \quad (4.83)$$

Clearly, the \mathbf{r}_i vectors correspond to the “corner points” of the “path” if one starts at \mathbf{r} and then moves by $\Delta r^{(1)}$ along the first dimension, then by $\Delta r^{(2)}$ along the second and so on. By construction, in the end $\mathbf{r}_n = \mathbf{r}_*$.

As a consequence of the growth properties of g , we have

$$g(\mathbf{r}_i) - g(\mathbf{r}_{i-1}) = g(\mathbf{r}_{i-1} + \Delta r^{(i)} \mathbf{e}_i) - g(\mathbf{r}_{i-1}) = \tilde{c}^{(i)} \Delta r^{(i)} \quad (4.84)$$

where $\underline{h}^{(i)} \leq \tilde{c}^{(i)} \leq \bar{h}^{(i)}$, and since

$$\Delta g = g(\mathbf{r}_n) - g(\mathbf{r}_0) = \sum_{i=1}^n [g(\mathbf{r}_i) - g(\mathbf{r}_{i-1})] \quad (4.85)$$

the result (4.82) now follows.

Next, we replace the difference $\mathbf{r}_* - \mathbf{r}$ by the corresponding difference $\mathbf{t}_* - \mathbf{t}$. As a consequence of the growth properties of the utility functions $f^{(i)}$, we have

$$f^{(i)}(r_*^{(i)}) - f^{(i)}(r^{(i)}) = d^{(i)}(r_*^{(i)} - r^{(i)}) = d^{(i)} \Delta r^{(i)} \quad (4.86)$$

where $0 < \underline{d}^{(i)} \leq d^{(i)} \leq \bar{d}^{(i)}$. But since $t^{(i)} = f^{(i)}(r^{(i)})$ and $t_*^{(i)} = f^{(i)}(r_*^{(i)})$ we see that

$$t_*^{(i)} - t^{(i)} = d^{(i)} \Delta r^{(i)} \quad (4.87)$$

Hence,

$$\Delta r^{(i)} = \frac{t_*^{(i)} - t^{(i)}}{d^{(i)}} \quad \text{where} \quad 0 < \underline{d}^{(i)} \leq d^{(i)} \leq \bar{d}^{(i)} \quad (4.88)$$

Combining (4.82) and (4.88) now yields the desired result (4.80).

4.A.2 Global- and utility functions used in our simulations

For some simulations, the utility functions were chosen to be of quadratic form on the interval $[0, 1.5]$, and linear outside this range. Specifically, the functions were of the form

$$t^{(i)} = \begin{cases} \alpha_1^{(i)}(r^{(i)})^2 + \alpha_2^{(i)}r^{(i)} + \alpha_3^{(i)} & \text{if } 0 \leq r^{(i)} \leq 1.5 \\ \beta_1^{(i)}r^{(i)} + \beta_2^{(i)} & \text{if } r^{(i)} < 0 \\ \beta_3^{(i)}r^{(i)} + \beta_4^{(i)} & \text{otherwise} \end{cases} \quad (4.89)$$

where the coefficients $\alpha_1^{(i)}, \alpha_2^{(i)}, \alpha_3^{(i)}$ were chosen within appropriate bounds to guarantee invertibility on the interval $[0, 1.5]$. The coefficients $\beta_1^{(i)}, \dots, \beta_4^{(i)}$ were also chosen randomly, but in such a way as to guarantee that the overall function would be continuous (i.e. that the linear segments join up with the quadratic part). A set of 25 randomly generated functions of this type are shown in Figure 4.8(a).




Figure 4.8: Illustrations of the utility functions $t^{(i)} = f^{(i)}(r^{(i)})$ used in the simulations of Algorithms 1 and 2.

In other simulations we used piecewise linear utility functions, shown in Figure 4.8(b), also based on randomised coefficients.

The global functions used were also either of quadratic form

$$g(r) = d + \sum_{i=1}^n q^{(i)}(r^{(i)}) \quad (4.90)$$

where the $q^{(i)}$ were of a similar type as (4.89), or of affine form

$$g(\mathbf{r}) = d + \sum_{i=1}^n c^{(i)} r^{(i)} \quad (4.91)$$

where the parameters $c^{(i)} > 0$ and $d > 0$ were also chose at random.

All random parameters were chosen within certain bounds from which the required growth conditions for the theorems were then easily derived.

Switching, Feedback and Estimation

In this chapter, we add an estimation component to the general cooperative control problem, proposing a decentralised control scheme for regulating the topology of a wireless sensor network. First, an algorithm is developed that approximates the connectivity level as measured by the second largest eigenvalue of a stochastic normalisation of the system's adjacency matrix. These estimates are then used to inform a cooperative control algorithm that iteratively regulates the network's connectivity to some desired level.

Chapter contents

- 5.1 Introduction
 - 5.2 Preliminaries
 - 5.3 Decentralised estimation of the second eigenvalue
 - 5.4 Decentralised connectivity control
 - 5.5 Simulation results
 - 5.6 Conclusion
-

5.1 Introduction

The previous chapter presented a number of algorithms designed to solve a regulation problem involving both global and local constraints, operating in a variety of different settings with different assumptions. However, the common assumption throughout was that the global term is “provided” to one or more nodes in the networks so that they could integrate it into the control scheme. In contrast to this work, the present chapter now investigates an example where this assumption cannot be made. The global term thus needs to be estimated by the network itself. The following presents joint work with Dr. R. Stanojević, Prof. M. Corless and Prof. R. Shorten and has been published in Knorn *et al.* (2009c,d).

Recent years have witnessed a growing interest in the control community in problems that arise when dynamic systems evolve over graphs. But while the most high profile of these applications are clearly in consensus applications such as formation flying, synchronisation problems and sensor networks, there are also many other applications where the

manner in which the network topologies change affects the performance of algorithms that evolve over these graphs. In such applications, an essential requirement is that the topology of the graph be such that some basic properties required to support communication and control are satisfied, the most basic of these being that the network be connected. Considerations of this kind have given rise to the emerging field of network topology control.

The work in this chapter is inspired by the third motivating example we gave in Section 1.2.3 on page 4: A wireless sensor network that is based on stationary nodes (i.e. nodes that do not change their geographical location) that are able to adjust the transmit power in their radios and hence control the area over which they can broadcast information.¹ This means that by changing their *broadcast radius* (that is the distance from the transmitter up to which information can be reliably received) the nodes can directly influence the topology of the resulting communication network. Using the terminology of the previous chapter, each node's broadcast radius would be its physical state. No utility values will be considered in the present context, in other words the utility function is the identity function. The global quantity of interest will be the communication network's level of connectedness or an algebraic proxy thereof (this will be defined more precisely in Section 5.2).

Recall that, roughly speaking, a graph is (*strongly*) *connected* in the classic graph-theoretic sense if there exists at least one path from any one node in the network to any other. As we saw in the previous chapter, graph connectivity is an essential component in situations where a group of networked nodes must work together, in a decentralised manner, to achieve a common task. This issue of graph connectivity is therefore very important and has achieved much attention in various contexts. It appears that this work has followed three lines of enquiry. In the graph theory literature, attempts have been made to identify and grow graphs with pre-specified connectivity properties; see Fallat and Kirkland (1998); Ghosh and Boyd (2006); Boyd *et al.* (2004) and the references therein for an overview of this work. In the computer science and networking communities several attempts have also been made to identify local (node based) constraints that guarantee certain forms of graph connectivity. For example, the sector rule proposed in Wattenhofer *et al.* (2001) is one such rule that gives rise to certain types of connected graphs. Recently, work in this direction has been extended to reflect not only topological considerations, but also the effect of physical constraints such as power and interference, in achieving these objectives. Finally, a third strand of work has recently emerged in the control and robotics community. Roughly speaking, this work involved using feedback principles to achieve graphs with a desired topology. Examples of this work can be found in Ramanathan and Rosales-Hain (2000); Ji and Egerstedt (2005); Gennaro and Jadbabaie (2006); Cabrera *et al.* (2007); Dimarogonas and Kyriakopoulos (2008) and the references therein. In particular, Gennaro

¹ Such networks are widely used in many engineering problems, see for instance Akyildiz *et al.* (2002) for a very detailed survey of the area of wireless sensor networks.

and Jadbabaie have proposed an interesting approach to distributed control of the second smallest eigenvalue of the communication graph's Laplacian, Gennaro and Jadbabaie (2006). Those ideas were further developed in Yang *et al.* (2008). In this line of work, however, nodes have a fixed communication radius and change their positions relative to each other in order to achieve a desired connectivity level, with the consequence that the communication graphs are always undirected graphs.

Contributions

Clearly, regulation of the connectivity of a given graph is difficult because graph connectivity is a global property, whereas typically, nodes (or agents) can only act locally. Thus, any algorithm for maintaining graph connectivity must be decentralised if it is to be of any practical value. Our objective here is to propose one such algorithm; namely, a decentralised algorithm that is simple to implement yet efficiently regulates the connectivity level of a given graph to some pre-specified value. To that end we first develop and prove convergence of a decentralised estimation scheme whereby each node can estimate the level of graph connectivity (as a proxy for the level of connectivity we will use the second largest eigenvalue of a stochastic normalisation of the graph's adjacency matrix). We then present a control strategy to regulate the graph connectivity about a specified set-point. This approach may be seen as an adaptation of Algorithm 1 described in the previous chapter, but in contrast to our earlier work, the global function encountered here (which now describes the dependency of the eigenvalue on each node's broadcast radius) is neither continuous nor strictly monotone. Simulation results are also given to illustrate the theoretical contributions, and we present examples to show that our control framework is sufficiently general to allow other constraints such as local power, interference, or node density to be part of a connectivity/interference trade-off as well.

The work carried out in this chapter differs from that in the literature in a number of aspects. Firstly, some of the previous results are of a probabilistic nature, i.e. they draw statistical conclusions of the type “in average, roughly every third graph of this kind should be connected”. However, the application scenario that we have in mind consists of a concrete situation where a number of sensors are placed randomly in space (for example, a set of nodes dropped over a lake, each node communicating only with a subset of its neighbours). In this case, drawing probabilistic conclusions is of little help, as we would like to find results for particular instances of the problem. We are also interested in situations where information mixes quickly across the graph, which means that we must specifically account for the speed at which information passing takes place — and not just that the graph is connected (in the classic graph-theoretic sense). Finally, as before, we wish to develop algorithms that can be used irrespective of graph type where again we wish to break free of the assumption that the underlying graph structure is symmetric. This again delineates the work presented here from much of the recent results in the area.

Finally, we argue that our algorithms are very simple to implement and require minimal computational requirements, and give rise to graph growing techniques with truly scale-free properties.

Structure

In the next section we introduce the basic idea behind our approach and describe the general set-up and notation. We will then present our decentralised estimation scheme that iteratively approximates the second largest eigenvalue. We discuss in Section 5.4 how this value could be used to control the networks connectivity by proposing a simple controller based on these estimates, and determine the conditions for the stability of the decentralised closed loop system. Results from simulations are then presented in Section 5.5. Finally conclusions and future directions are given in the last section.

5.2 Preliminaries

5.2.1 Basic idea

Our basic idea for connectivity estimation is based on the observation that dynamic systems or algorithms evolving on graphs often reveal topological properties about the graph itself. One such algorithm is the *distributed averaging* or *consensus* algorithm, which is strongly related to the theory of Markov chains and to (non)homogeneous matrix products. While the primary focus of the work reported here is not on the dynamics of consensus algorithms, it is important to note here that the second eigenvalue of the averaging matrix (see notation section below) determines the rate at which the nodes in the network achieve consensus. Roughly speaking, as a graph becomes less connected this second eigenvalue becomes closer to unity, when rate of convergence is used as a measure of connectivity. Further, as we shall see, a simple algorithm can be used, together with elementary techniques from system identification, to locally estimate this eigenvalue in a decentralised manner.

Let us briefly illustrate these basic points in Figure 5.1 on the next page. Here, we show the average value of the second largest eigenvalue in magnitude of the averaging matrix of random (regular) graphs.² The averaging matrix was constructed directly from a stochastic normalisation of the adjacency matrix of the underlying graph. In the plot, the value of the second largest eigenvalue drops monotonically with increasing graph regularity (fixed number of neighbours per node). Although this is a very special type of graph, it shows that a single value can give an indication of the connectivity situation of a graph.

Comment Classically, the second smallest eigenvalue of the *Laplacian* (or *transition Laplacian*) matrix of a graph has been used as an algebraic measure for connectivity, Fiedler (1973); Chung (1997). However, usually Laplacians are only defined for undirected

² A d -regular graph is a graph where each node has exactly d neighbours (here chosen at random).




Figure 5.1: Average of the magnitude of the second largest eigenvalue of the averaging matrix of d -regular random graphs with 200 nodes.

graphs, and this is an unnatural restriction that we would like to eliminate. In contrast, the second largest eigenvalue (in magnitude) of an averaging matrix is also an excellent candidate to indicate the degree of connectivity of an entire graph (independent of the fact whether the underlying graph is directed or not) with the added benefit of being easily approximated locally in each node using computationally inexpensive estimation techniques as shown below. /

Knowledge of global information such as level of connectivity, based on purely *local* information, offers a wide range of local node actions with the objective of connectivity maintenance, one of which is will be presented in this chapter. For example, in the context of wireless networks, one possible action is for nodes to adjust the power of their radio transmissions, based on the local estimate of connectivity. Concretely, this could mean to reduce the communication radius if the connectivity is estimated to be larger than required (as decreasing the radius will lead to reducing the number of neighbours, hence reducing connectivity). A pseudo-protocol for such a strategy is given in Figure 5.2 on the following page.

That such a strategy is well posed is evident and follows from the basic observation that if all nodes increase their communication radii sufficiently, then the graph will eventually become more densely connected. The issues that make the realisation of such strategies challenging in a practical environment concern decentralised estimation of the second largest eigenvalue of the averaging matrix, and proving that the resulting closed loop strategy is robustly stable. Resolving these issues will be the main concern of this chapter.

-
- 1: Deploy pre-configured nodes and initialise network by choosing random initial communication radii such that network is connected.
 - 2: By running a consensus algorithm on the network, each node estimates the second largest eigenvalue of the averaging matrix based on the convergence of its own state.
 - 3: For each node, if the estimated eigenvalue is smaller than some desired value, decrease the broadcast radius; if the estimate is larger, increase the radius.
 - 4: Go to 2.
-

Figure 5.2: Pseudo-protocol for the overall scheme.

5.2.2 General setting

Building on Section 4.2.2, we assume that a consensus / averaging algorithm evolves on the graph \mathcal{G} . Formally, associate to each node $i = 1, \dots, n$ in the network a state $x^{(i)} \in \mathbb{R}$. The state of node i at time k is denoted $x_k^{(i)}$, and the network's state (i.e. the states of all the nodes combined) is the column vector $\mathbf{x}_k = (x_k^{(1)}, \dots, x_k^{(n)})^\top$. For each node $i = 1, \dots, n$, a distributed averaging scheme can then be written as

$$x_{k+1}^{(i)} = \sum_{j=1}^n p^{(ij)} x_k^{(j)} \quad \text{where} \quad \sum_{j=1}^n p^{(ij)} = 1 \quad \text{and} \quad \begin{cases} p^{(ij)} > 0 & \text{if } j \in \mathcal{N}^{(i)} \\ 0 & \text{otherwise} \end{cases} \quad (5.1)$$

for $k = 0, 1, 2, \dots$ with some initial condition $x_{k=0}^{(i)} = x_0^{(i)}$. It is easy to see that this relation can be written for the overall network as

$$\mathbf{x}_{k+1} = \mathbf{P} \mathbf{x}_k \quad \text{where} \quad \mathbf{x}_{k=0} = \mathbf{x}_0 \quad (5.2)$$

and where the stochastic, non-negative $\mathbf{P} = (p^{(ij)})$ is called the *averaging matrix*.

Let $\lambda^{(1)}, \dots, \lambda^{(n)}$ be the eigenvalues of \mathbf{P} and assume that they are ordered so that $|\lambda^{(i)}| \geq |\lambda^{(j)}|$ when $i \leq j$. To make matters more tractable we shall assume in the following that \mathbf{P} is always diagonalisable.³ Further, by making this assumption we have that \mathbf{P} has n linearly independent eigenvectors, $\boldsymbol{\nu}^{(1)}, \dots, \boldsymbol{\nu}^{(n)}$ corresponding to the eigenvalues $\lambda^{(1)}, \dots, \lambda^{(n)}$ (with a slight abuse of our usual notation, $\boldsymbol{\nu}^{(i)}$ and $\lambda^{(i)}$ denotes the i th eigenvector-eigenvalue pair). Thus these eigenvectors form a basis for \mathbb{R}^n and every initial state \mathbf{x}_0 can be uniquely expressed as

$$\mathbf{x}_0 = c^{(1)} \boldsymbol{\nu}^{(1)} + c^{(2)} \boldsymbol{\nu}^{(2)} + \dots + c^{(n)} \boldsymbol{\nu}^{(n)} \quad (5.3)$$

³ Since the set of diagonalisable matrices is dense in the set of stochastic matrices, this assumption is an entirely reasonable one to make.

for some scalars $c^{(1)}, \dots, c^{(n)}$. Since $\mathbf{P}\boldsymbol{\nu}^{(i)} = \lambda^{(i)}\boldsymbol{\nu}^{(i)}$,

$$\mathbf{x}_k = \mathbf{P}^k \mathbf{x}_0 = \mathbf{P}^k (c^{(1)}\boldsymbol{\nu}^{(1)} + c^{(2)}\boldsymbol{\nu}^{(2)} + \dots + c^{(n)}\boldsymbol{\nu}^{(n)}) \quad (5.4)$$

$$= c^{(1)}(\lambda^{(1)})^k \boldsymbol{\nu}^{(1)} + c^{(2)}(\lambda^{(2)})^k \boldsymbol{\nu}^{(2)} + \dots + c^{(n)}(\lambda^{(n)})^k \boldsymbol{\nu}^{(n)} \quad (5.5)$$

If the underlying graph is strongly connected, and since \mathbf{P} has positive entries along the main diagonal, it follows that \mathbf{P} is *primitive*, Horn and Johnson (1985). Thus, the Perron eigenvalue $\lambda^{(1)} = 1$ is simple and all other eigenvalues are smaller in magnitude. Also, $\boldsymbol{\nu}^{(1)} = \mathbf{1}$ as \mathbf{P} is row-stochastic, hence

$$\mathbf{x}_k = c^{(1)}\mathbf{1} + (\lambda^{(2)})^k \left[\sum_{j=2}^n c^{(j)} \left(\frac{\lambda^{(j)}}{\lambda^{(2)}} \right)^k \boldsymbol{\nu}^{(j)} \right] \quad (5.6)$$

and

$$\|\mathbf{x}_k - c^{(1)}\mathbf{1}\| \leq |\lambda^{(2)}|^k \beta(\mathbf{x}_0) \quad \text{with} \quad \beta(\mathbf{x}_0) = \sum_{j=2}^n |c^{(j)}| \cdot \|\boldsymbol{\nu}^{(j)}\| \quad (5.7)$$

where $\|\cdot\|$ denotes some norm.

In this case, \mathbf{x}_k converges exponentially to $c^{(1)}\mathbf{1}$ and the rate of convergence is bounded by $|\lambda^{(2)}|$. In other words, the rate of convergence of the distributed averaging can be measured by the magnitude of $\lambda^{(2)}$. Together with the intuition that the more the graph is connected the faster the averaging should converge, we can now see that $|\lambda^{(2)}|$ may very well be used as a proxy for the level of connectivity of the graph and the rate at which information can flow through it.

5.3 Decentralised estimation of the second eigenvalue

We now provide a simple method by which all nodes in the network may estimate $\lambda^{(2)}$ based only on local measurements.

Our basic idea is as follows. Once we know whether $\lambda^{(2)}$ is real or complex (non-real), different methods can be used to accurately estimate its magnitude based only on local measurements. For example, when $\lambda^{(2)}$ is real then the direct estimation method described by Proposition 5.1 will yield a correct estimate of $|\lambda^{(2)}|$. Also, the dynamic system that governs the evolution of $z_k^{(i)} := x_k^{(i)} - x_{k-1}^{(i)}$ can be modelled asymptotically as a first order linear system (with a noise term that decays to zero) if $\lambda^{(2)}$ is real valued. The parameters of that linear system can then be identified through an estimation method such as the classic recursive least squares algorithm (RLS, see for instance Haykin, 2002) providing another estimate of the absolute value of $\lambda^{(2)}$. When $\lambda^{(2)}$ is complex (non-real), a third estimation method, based on Proposition 5.2 below can be applied. Thus with appropriate numerical conditioning of the values of $z_k^{(i)}$, estimation of $\lambda^{(2)}$ can be carried out in a straightforward manner.

```

1:   $z_k^{(i)} = x_k^{(i)} - x_{k-1}^{(i)}$ 
2:   $A = \text{Estimate\_real}(z^{(i)}())$ 
3:   $B = \text{RLS\_real}(z^{(i)}())$ 
4:   $C = \text{Estimate\_complex}(z^{(i)}())$ 
5:  if  $|A - B| < \epsilon$ 
6:    return A
7:  else
8:    return C
9:  end if

```

Figure 5.3: Pseudocode for the overall estimation scheme of $|\lambda^{(2)}|$.

However, it is usually not clear *a priori* whether the averaging matrix \mathbf{P} has a real or complex (non-real) second eigenvalue (the exception being undirected graphs where $\lambda^{(2)}$ is always real valued). Thus we must develop a method for determining whether or not this eigenvalue is real or complex (non-real). To that end, we use the three estimators presented above and run them in parallel. Specifically, we first obtain estimates for $\lambda^{(2)}$ from the estimator based on Proposition 5.1 as well as the recursive least squares scheme, both of which are guaranteed to work only when $\lambda^{(2)}$ is real. If both estimates of $\lambda^{(2)}$ match up to a certain degree (that is, the absolute difference between the two values is less than some threshold ϵ), we assume that λ_2 is real and use these estimates. However, if the estimates do not match sufficiently, we consider $\lambda^{(2)}$ to be complex (non-real) and use the estimate obtained based on the Proposition 5.2 (which is guaranteed to converge to the correct value in that case). The pseudocode for this strategy is given in Figure 5.3.

In the rest of the section we provide the details explaining what each of the functions `Estimate_real()`, `RLS_real()` and `Estimate_complex()` do. All three functions require the distributed averaging algorithm to be run on the network, and each node is assumed to be able to store a small number of its own past states.

5.3.1 Estimate_real()

The following Proposition provides a method of estimating the value of the second largest eigenvalue of the averaging matrix provided the eigenvalue is real valued.

Proposition 5.1 (Decentralised estimation of real valued $\lambda^{(2)}$)

Let $\mathcal{G} = (\mathcal{V}, \mathcal{A})$ be a strongly connected network with averaging matrix \mathbf{P} such that its second largest eigenvalue in magnitude $\lambda^{(2)}$ is real and satisfies $|\lambda^{(2)}| > |\lambda^{(j)}|$ for all $j > 2$.

Consider any node i and let $z_k^{(i)} := x_k^{(i)} - x_{k-1}^{(i)}$ where x_k is determined by the distributed averaging algorithm (5.2) running on the network with a sufficiently general initial condi-

tion. Consider any positive integer m and for $k \geq m+1$, let

$$\tilde{\lambda}_k^{(i,2)} = \left| \frac{z_k^{(i)}}{z_{k-m}^{(i)}} \right|^{1/m} \quad (5.8)$$

be node i 's estimate of $|\lambda^{(2)}|$. Then $\lim_{k \rightarrow \infty} \tilde{\lambda}_k^{(i,2)} = |\lambda^{(2)}|$. _____

Proof Recall from (5.6) that for any node i :

$$x_k^{(i)} = c^{(1)} + (\lambda^{(2)})^k \underbrace{\left[\sum_{j=2}^n c^{(j)} \left(\frac{\lambda^{(j)}}{\lambda^{(2)}} \right)^k \nu^{(i,j)} \right]}_{=: \psi_k^{(i)}} \quad (5.9)$$

where $\nu^{(i,j)}$ denotes the i th element of the j th eigenvector of \mathbf{P} . We then have for $k > m+1$

$$\begin{aligned} \frac{z_k^{(i)}}{z_{k-m}^{(i)}} &= \frac{x_k^{(i)} - x_{k-1}^{(i)}}{x_{k-m}^{(i)} - x_{k-m-1}^{(i)}} = \frac{(\lambda^{(2)})^k \psi_k^{(i)} - (\lambda^{(2)})^{k-1} \psi_{k-1}^{(i)}}{(\lambda^{(2)})^{k-m} \psi_{k-m}^{(i)} - (\lambda^{(2)})^{k-m-1} \psi_{k-m-1}^{(i)}} \\ &= (\lambda^{(2)})^m \underbrace{\frac{\psi_k^{(i)} - (\lambda^{(2)})^{-1} \psi_{k-1}^{(i)}}{\psi_{k-m}^{(i)} - (\lambda^{(2)})^{-1} \psi_{k-m-1}^{(i)}}}_{=: w_{k,m}^{(i)}} \end{aligned} \quad (5.10)$$

and taking the m th root of the absolute values of both sides

$$\underbrace{\left| \frac{z_k^{(i)}}{z_{k-m}^{(i)}} \right|^{1/m}}_{\tilde{\lambda}_k^{(i,2)}} = |\lambda^{(2)}| \cdot |w_{k,m}^{(i)}|^{1/m} \quad (5.11)$$

From the last equation we can see that the estimate $\tilde{\lambda}_k^{(i,2)}$ approaches the true absolute value of the second largest eigenvalue if and only if the $|w_{k,m}^{(i)}| \rightarrow 1$, as k grows. Since

$$\psi_k^{(i)} = c^{(2)} \nu^{(i,2)} + \sum_{j=3}^n c^{(j)} \left(\frac{\lambda^{(j)}}{\lambda^{(2)}} \right)^k \nu^{(i,j)} \quad (5.12)$$

it will converge to $c^{(2)} \nu^{(i,2)}$ as k grows, as by assumption $\left| \frac{\lambda^{(j)}}{\lambda^{(2)}} \right|^k < 1$ for $j = 3, \dots, n$. For a general initial condition $c^{(2)} \nu^{(i,2)}$ is non-zero and, using (5.10), we now have that $|w_{k,m}^{(i)}| \rightarrow 1$ and thus $\tilde{\lambda}_k^{(i,2)} \rightarrow |\lambda^{(2)}|$ as $k \rightarrow \infty$. _____ \square

In summary, if the prerequisites are met, for $k \geq m+1$, each node can iteratively refine its estimate of $|\lambda^{(2)}|$ with (5.8) so that it converges to the true value of $|\lambda^{(2)}|$ as k grows.

Comment It also follows from the proof that larger the gap between $|\lambda^{(2)}|$ and $|\lambda^{(3)}|$ the faster the estimates $\tilde{\lambda}_k^{(i,2)}$ will converge to the true value of $|\lambda^{(2)}|$. /

5.3.2 RLS_real()

When $\lambda^{(2)}$ is real we can also use a recursive least squares algorithm for estimating $\lambda^{(2)}$. It can be seen from (5.10) that by letting $m = 1$ we have for $k = 1, 2, \dots$ the following relationship (asymptotically)

$$|z_{k+1}^{(i)}| \simeq |\lambda^{(2)}| \cdot |z_k^{(i)}| \quad (5.13)$$

Applying a suitably parametrised recursive least squares algorithm, see for instance Åström and Wittenmark (1997), should then also yield good estimates for $|\lambda^{(2)}|$.

5.3.3 Estimate_complex()

The next proposition provides a method for estimating the magnitude of a complex (non-real) valued $\lambda^{(2)}$. When $\lambda^{(2)}$ is complex (non-real), its complex conjugate $\bar{\lambda}^{(2)}$ is also an eigenvalue of \mathbf{P} with the same magnitude. If we assume that $|\lambda^{(2)}| > |\lambda^{(j)}|$ for all $j > 2$ then, recalling (5.6), it is straightforward to show that, for each node i , the variable $z_k^{(i)} = x_k^{(i)} - x_{k-1}^{(i)}$ can be written as

$$z_k^{(i)} = c^{(i)} (\lambda^{(2)})^k + \bar{c}^{(i)} (\bar{\lambda}^{(2)})^k + |\lambda^{(2)}|^k O_k^{(i)} \quad (5.14)$$

where $O_k^{(i)} \rightarrow 0$ as $k \rightarrow \infty$ and $c^{(i)}, \bar{c}^{(i)} \neq 0$ for a sufficiently general initial condition of the averaging algorithm.

Proposition 5.2 (Decentralised estimation of the magnitude of a complex (non-real) $\lambda^{(2)}$) —

Let $\mathcal{G} = (\mathcal{V}, \mathcal{A})$ be a strongly connected network with averaging matrix \mathbf{P} such that its second largest eigenvalue in magnitude $\lambda^{(2)}$ is complex with non-zero imaginary part and $|\lambda^{(2)}| = |\bar{\lambda}^{(2)}| > |\lambda^{(j)}|$ for $j > 2$.

Consider any node i and let

$$\zeta_k^{(i)} := z_k^{(i)} z_{k-2}^{(i)} - (z_{k-1}^{(i)})^2 \quad (5.15)$$

where $z_k^{(i)} = x_k^{(i)} - x_{k-1}^{(i)}$ and \mathbf{x}_k is determined by the distributed averaging algorithm (5.2) running on the network with a sufficiently general initial condition. Consider any positive integer m and for $k \geq m + 3$, let

$$\tilde{\lambda}_k^{(i,2)} = \left| \frac{\zeta_k^{(i)}}{\zeta_{k-m}^{(i)}} \right|^{\frac{1}{2m}} \quad (5.16)$$

be node i 's estimate of $|\lambda^{(2)}|$. Then $\lim_{k \rightarrow \infty} \tilde{\lambda}_k^{(i,2)} = |\lambda^{(2)}|$. _____

Proof For any node i , substituting expression (5.14) into (5.15), and dropping the superscripts “ (i) ” and “ (2) ” to increase legibility, yields

$$\begin{aligned}
\zeta_k &= \left[c\lambda^k + \bar{c}\bar{\lambda}^k + |\lambda|^k O_k \right] \left[c\lambda^{k-2} + \bar{c}\bar{\lambda}^{k-2} + |\lambda|^{k-2} O(k-2) \right] \dots \\
&\quad \dots - \left[c\lambda^{k-1} + \bar{c}\bar{\lambda}^{k-1} + |\lambda|^{k-1} O_{k-1} \right]^2 \\
&= c\bar{c} \left[\lambda^k \bar{\lambda}^{k-2} + \bar{\lambda}^k \lambda^{k-2} - 2\lambda^{k-1} \bar{\lambda}^{k-1} \right] + |\lambda|^{(2k-4)} \tilde{O}_k \\
&= |c|^2 |\lambda|^{(2k-4)} \left[\bar{\lambda}^2 + \lambda^2 - 2\lambda\bar{\lambda} \right] + |\lambda|^{(2k-4)} \tilde{O}_k \\
&= |\lambda|^{(2k-4)} \left[|c|^2 (\lambda - \bar{\lambda})^2 + \tilde{O}_k \right]
\end{aligned} \tag{5.17}$$

where

$$\begin{aligned}
\tilde{O}_k &= |\lambda|^2 \left\{ O_{k-2} \left[c \left(\frac{\lambda}{|\lambda|} \right)^k + \bar{c} \left(\frac{\bar{\lambda}}{|\lambda|} \right)^k \right] \dots \right. \\
&\quad \dots + O_k \left[c \left(\frac{\lambda}{|\lambda|} \right)^{k-2} + \bar{c} \left(\frac{\bar{\lambda}}{|\lambda|} \right)^{k-2} \right] + O_k O_{k-2} \dots \\
&\quad \left. \dots - 2 O_{k-1} \left[c \left(\frac{\lambda}{|\lambda|} \right)^{k-1} + \bar{c} \left(\frac{\bar{\lambda}}{|\lambda|} \right)^{k-1} \right] - O_{k-1}^2 \right\}
\end{aligned} \tag{5.18}$$

We note that since $O_k \rightarrow 0$ as $k \rightarrow \infty$, we also have

$$\lim_{k \rightarrow \infty} \tilde{O}_k = 0 \tag{5.19}$$

Furthermore, since $c, \bar{c} \neq 0$ and λ has nonzero imaginary part, $|c|^2 (\lambda - \bar{\lambda})^2$ is nonzero, and thus ζ_k in (5.17) is also non-zero for k sufficiently large. Finally,

$$\left| \frac{\zeta_k}{\zeta_{k-m}} \right| = |\lambda|^{2m} \left[\frac{|c|^2 (\lambda - \bar{\lambda})^2 + \tilde{O}_k}{|c|^2 (\lambda - \bar{\lambda})^2 + \tilde{O}_{k-m}} \right] \tag{5.20}$$

From this last expression and (5.19) we obtain that

$$\lim_{k \rightarrow \infty} \left| \frac{\zeta_k}{\zeta_{k-m}} \right|^{\frac{1}{2m}} = |\lambda| \tag{5.21}$$

which completes the proof. □

Based on this proposition, if $\lambda^{(2)}$ is complex (non-real) and each node calculates an estimate of $|\lambda^{(2)}|$ through (5.16) then the estimate will converge to the true value as k grows.

5.3.4 Remarks

The decision heuristic presented (Figure 5.3 on page 90) assumes that the first two estimators (which are designed for *real* valued $\lambda^{(2)}$ only) produce wrong and in particular *differently* wrong estimates, so that there is a sufficiently large disagreement between both

schemes so that it can be detected — clearly, if both schemes produced wrong but identical estimates, then our heuristic would consider these wrong estimates to be correct. However, our assumption of sufficiently different biases between the schemes is plausible given the fact that the first scheme only uses two samples for the estimation, whereas the recursive least squares scheme uses the entire history of samples attempting to minimise the square error between model and observed data.

An alternative approach can be used by employing several instances of the first estimation scheme, but using different m parameters. It can be shown (and this will indeed be observed in the simulations below) that in the presence of complex valued $\lambda^{(2)}$, the estimates produced by the scheme will exhibit some periodic, oscillatory behaviour. Roughly speaking, this oscillatory behaviour is due to the expression of the $|w_{k,m}^{(i)}|$ in (5.11) not converging to 1; rather it consists of a fraction of trigonometric functions that produces these oscillations (a similar behaviour can so also be shown for the RLS based estimator). In particular, the m parameter will affect the phase of these oscillations. Thus, using multiple instances of the first estimation scheme with different m parameters may be an alternative approach to detect whether $\lambda^{(2)}$ is real valued or not.

Next, by its very nature, when running the consensus algorithm over a connected network, the states of all nodes will converge to a common value. In that case, the difference in states $z_k^{(i)}$ will tend to zero. On the one hand, numerical calculation of the $z_k^{(i)}$ will be less and less precise as the $z_k^{(i)}$ approach zero, and on the other, when using the algorithms based on Propositions 5.1 and 5.2, the division of $z_k^{(i)}$ by $z_{k-m}^{(i)}$ resp. $\zeta_k^{(i)}$ by $\zeta_{k-m}^{(i)}$ will also become more and more numerically problematic. It is, however, not too difficult to solve these problems. Simply, whenever some node's state $x_k^{(i)}$ agrees with all of its neighbours on the top s digits, it shall stop broadcasting those top s digits and keep exchanging only the lower weight digits.

We must assume that in an actual implementation sufficiently exact numerical computations can be provided as the current approach does not take into account the inherently limited accuracy of numerical calculations in digital processors.

Finally, in this section we have assumed that there is a spectral gap between $\lambda^{(2)}$ (and its conjugates) and the remaining eigenvalues of the matrix \mathbf{P} . Since the set of matrices satisfying this property is dense in the set of stochastic matrices, this assumption is also entirely reasonable. However, the case where $|\lambda^{(2)}| = |\lambda^{(3)}| \geq |\lambda^{(4)}| \geq \dots$ can also be accommodated in our framework by including more estimators, similar to the ones presented above, and by modifying the logic described in Figure 5.3 on page 90 accordingly. This is omitted here for ease of exposition, and because the aforementioned case is a low probability event.

5.3.5 Demonstration of estimation

In the following two examples, we generated a two-dimensional random geometric graph with random connection radii for each node. These type of graphs are often used when modelling wireless networks, in particular wireless sensor networks, Penrose (2003); Santi (2005). A random *geometric graph* or *disc graph* is created as follows: Place n nodes uniformly distributed in the unit square, then interconnect the nodes based on the so-called *distance parameters* or *connection radii* of the nodes. That is, each node i has a parameter $r^{(i)}$ based on which it connects (or “sends information”) to other nodes that are closer than $r^{(i)}$ from it: if some node j is at (Euclidian) distance $d^{(ij)}$ from node i then there is an edge from node i to node j (i.e. node j is in reach) if and only if $d^{(ij)} \leq r^{(i)}$.




Figure 5.4: Comparison of the estimation schemes for real valued $\lambda^{(2)} \simeq 0.80$.

Figure 5.4 above and Figure 5.5 on the following page show the outputs of our three estimation schemes as well as their combination for two different situations: one where $\lambda^{(2)}$ is real, and one where $\lambda^{(2)}$ is complex (non-real). For each case we have plotted each nodes' estimates of $|\lambda^{(2)}|$ as a function of time (iterations of the estimation schemes), provided by the different algorithms, as well as the combination of using our proposed decision heuristic. From top to bottom, the subplots show the evolution of the estimates based on A) Proposition 1, B) recursive least squares and C) Proposition 2; as well as their combination in the last subplot. The true value of $|\lambda^{(2)}|$ is indicated by the dashed horizontal line.




Figure 5.5: Comparison of the estimation schemes for complex $\lambda^{(2)} \simeq 0.63 + 0.05i$.

Comment The following parameters were used. The random disc graphs on $n = 20$ nodes were created using connection radii $r^{(i)}$ uniformly distributed in the interval $[0.1, 0.6]$. We used $m = 5$ in the algorithm based on Proposition 5.1, and $m = 1$ in that based on Proposition 5.2. The initial estimates of the recursive least squares algorithm was set to 0.5. Finally, the combination of the estimates was done using the threshold $\epsilon = 0.005$. \swarrow

When the network has a real valued $\lambda^{(2)}$, it can be seen that the each node's estimates using the first two estimators converge quickly to the correct value. The estimates of the third estimator also converge, but to the wrong value.⁴ Since the two estimators targeted at a real valued $\lambda^{(2)}$ both converge to the same value, the error between them quickly drops below the preset threshold, and the combination scheme correctly switches to returning the value of the first estimator.

In the complex (non-real) case, Figure 5.5, the situation is different. Both the estimates of the estimators aimed at real valued $\lambda^{(2)}$ do not converge to the correct value of $|\lambda^{(2)}|$, but rather oscillate around it. The error between them is sufficiently large so that the combination scheme returns the value of the third estimator, which in turn now provides correct estimates.

⁴ In fact, it is not difficult to show that in this case the estimate which `Estimate_complex()` converges to will actually be $\lambda^{(2)} \sqrt{|\lambda^{(3)}|/|\lambda^{(2)}|}$.

5.4 Decentralised connectivity control

We now present our algorithm for decentralised connectivity control. Please note that, by an abuse of notation, we shall simply use λ in the remainder of this chapter to refer to $|\lambda^{(2)}|$. As mentioned already, we wish to adjust the communication radius of each sensor in the network, $\{r^{(1)}, \dots, r^{(n)}\}$ based on a local estimation of λ , with the ultimate objective of regulating λ to some neighbourhood of a target value; namely so that $|\lambda - \lambda_*| < \varepsilon$ for some $\lambda_* \in (0, 1)$ and $\varepsilon > 0$. Since we are trying to address situations in which individual sensors may fail resulting in a change in network connectivity, we are inherently dealing with situations where the network topology is slowly (but not constantly) changing. In what follows we therefore make the assumption of quasi-stationarity; specifically, we assume that the local estimators operate over very fast time scales when compared with the local control actions (local radius updates). This assumption greatly facilitates analytical tractability and makes our convergence proofs somewhat easier to develop. Finally, since there may be many sets of communication radii $\{r^{(1)}, \dots, r^{(n)}\}$ that guarantee $|\lambda - \lambda_*| < \varepsilon$, we shall make additional assumptions to guarantee that the closed loop algorithm converges to a *common* set of radii; namely, we seek a control action that guarantees convergence of all radii to the same value. We emphasise again that this assumption is made to facilitate analytical tractability, but it can also be motivated from a practical standpoint, where having all nodes use the same broadcast radius should contribute to similar battery lifetimes of the nodes. However, our framework is sufficiently general to allow other quantities of interest to be included (for instance, equal numbers of neighbours, maximum numbers of neighbours); although, the convergence proofs will change accordingly.

5.4.1 Consensus with input

Our control algorithm is again motivated by the intuitive idea that adding the same value to each member in a consensus scheme will not hinder the eventual agreement between the members. This was already noted in Lemma 4.1. However, this notion can be applied to a much broader class of consensus schemes as we show using the recent results of Moreau, 2005.

Theorem 5.1 (Generalised consensus with common input) _____

Let $\mathcal{G}_k = (\mathcal{V}, \mathcal{A}_k)$ be a sequence of strongly connected graphs, $\theta(\mathbf{x}_k, k)$ be a sequence of finite real numbers and f be a map on \mathcal{G}_k satisfying the following conditions. Associated to each directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{A})$ with node set $\mathcal{V} = \{1, \dots, n\}$, each node $i \in \mathcal{V}$ and each state $\mathbf{x} \in \mathcal{X}^n$, there is a compact set $\mathcal{E}^{(i)}(\mathcal{A})(\mathbf{x}) \subset \mathcal{X}$ satisfying:

- (i) $f^{(i)}(\mathbf{x}, k) \in \mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{x}) \quad \forall k \in \mathbb{N} \quad \forall \mathbf{x} \in \mathcal{X}^n,$
- (ii) $\mathcal{E}^{(i)}(\mathcal{A})(\mathbf{x}) = \{x^{(i)}\}$ whenever the states of node i and its neighbouring nodes j are all equal,

- (iii) $\mathcal{E}^{(i)}(\mathcal{A})(\mathbf{x})$ is contained in the relative interior of the convex hull of the states of node i and its neighbouring nodes j whenever the states of node i and its neighbouring nodes j are not all equal,
- (iv) $\mathcal{E}^{(i)}(\mathcal{A})(\mathbf{x})$ depends continuously on \mathbf{x} , that is, the set-valued function $\mathcal{E}^{(i)}(\mathcal{A}) : \mathcal{X}^n \rightrightarrows \mathcal{X}$ is continuous.⁵

Then, if $\mathbf{x}_k = (x_k^{(1)}, \dots, x_k^{(n)})^\top$ evolves for some $\mathbf{x}_{k=0} = \mathbf{x}_0$ according to

$$\mathbf{x}_{k+1} = f(\mathbf{x}_k, k) + \theta(\mathbf{x}_k, k)\mathbf{1} \quad (5.22)$$

the elements of \mathbf{x}_k will approach each other over time, that is

$$\lim_{k \rightarrow \infty} x_k^{(i)} - x_k^{(j)} = 0 \quad (5.23)$$

for all $i, j \in \{1, \dots, n\}$.

Proof Start by defining

$$\mathbf{y}_k := \mathbf{x}_k - \sigma_k \mathbf{1} \quad \text{where} \quad \sigma_k := \sum_{i=0}^{k-1} \theta(\mathbf{x}(i), i) \quad (5.24)$$

Then $\sigma_{k+1} = \sigma_k + \theta(\mathbf{x}_k, k)$ and

$$\begin{aligned} \mathbf{y}_{k+1} &= \mathbf{x}_{k+1} - \sigma_{k+1} \mathbf{1} \\ &\stackrel{(5.22)}{=} f(\mathbf{x}_k, k) + \theta(\mathbf{x}_k, k)\mathbf{1} - [\sigma_k + \theta(\mathbf{x}_k, k)]\mathbf{1} \\ &\stackrel{(5.24)}{=} \underbrace{f(\mathbf{y}_k + \sigma_k \mathbf{1}, k) - \sigma_k \mathbf{1}}_{:=g(\mathbf{y}_k, k)} \end{aligned} \quad (5.25)$$

Now, if g satisfies all of the assumptions (1)–(4) of the theorem, the results from Moreau (2005) guarantee that all entries in \mathbf{y}_k will converge to a common value, and hence, through (5.24), the values in \mathbf{x}_k have to approach each other. So let us test g for each of the four assumptions.

- (i) For all nodes $i \in \mathcal{V}$,

$$g^{(i)}(\mathbf{y}_k, k) = f(\mathbf{y}_k + \sigma_k \mathbf{1}, k) - \sigma_k \in \underbrace{\mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{y}_k + \sigma_k \mathbf{1}) - \sigma_k}_{=:\hat{\mathcal{E}}^{(i)}(\mathcal{A}_k)(\mathbf{y}_k)} \quad (5.26)$$

Clearly, if $f(\mathbf{x}_k, k) \in \mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{x}_k)$ for all $i \in \mathcal{V}$, $k \in \mathbb{N}$ and $\mathbf{x} \in \mathcal{X}^n$, and if $\mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{x}_k)$ is compact, then $\hat{\mathcal{E}}^{(i)}(\mathcal{A}_k)(\mathbf{y}_k)$ is also compact given σ is bounded.

⁵ Put simply, these four conditions require that the updated state of each node must be a strict convex combination of its own and its neighbours' states, and that the update function must be continuous.

- (ii) Whenever the states of node i and its neighbours are all equal, that is $y_k^{(i)} = y_k^{(j)}$ for all $j \in \mathcal{N}^{(i)}$,

$$\hat{\mathcal{E}}^{(i)}(\mathcal{A}_k)(\mathbf{y}_k) = \mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{y}_k + \sigma_k \mathbf{1}) - \sigma_k = \{y_k^{(i)} + \sigma_k\} - \sigma_k = \{y_k^{(i)}\} \quad (5.27)$$

- (iii) Assume the states of node i and its neighbours $j \in \mathcal{N}^{(i)}$ are not all equal. If $\mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{x}_k)$ is contained in the relative interior of the convex hull ($\text{conv}\{\cdot\}$) of the states of node i and its neighbours, we have

$$\begin{aligned} \mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{x}_k) &\subset \text{convh}_{j \in \mathcal{N}^{(i)}} \{x_k^{(j)}\} \\ \mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{x}_k - \sigma_k \mathbf{1} + \sigma_k \mathbf{1}) &\subset \text{convh}_{j \in \mathcal{N}^{(i)}} \{x_k^{(j)} + \sigma_k - \sigma_k\} \\ \mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{y}_k + \sigma_k \mathbf{1}) - \sigma_k &\subset \text{convh}_{j \in \mathcal{N}^{(i)}} \{y_k^{(j)} + \sigma_k\} - \sigma_k \end{aligned}$$

and with $\text{convh}\{\cdot\}$ being a linear operator

$$\begin{aligned} \mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{y}_k + \sigma_k \mathbf{1}) - \sigma_k &\subset \text{convh}_{j \in \mathcal{N}^{(i)}} \{y_k^{(j)} + \sigma_k - \sigma_k\} \\ \hat{\mathcal{E}}^{(i)}(\mathcal{A}_k)(\mathbf{y}_k) &\subset \text{convh}_{j \in \mathcal{N}^{(i)}} \{y_k^{(j)}\} \end{aligned} \quad (5.28)$$

- (iv) If $\mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{x}_k)$ depends continuously on \mathbf{x}_k , so will $\mathcal{E}^{(i)}(\mathcal{A}_k)(\mathbf{x}_k + \sigma_k \mathbf{1}) - \sigma_k \mathbf{1} = \hat{\mathcal{E}}^{(i)}(\mathcal{A}_k)(\mathbf{y}_k)$.

We have thus established, that the update map g satisfies Assumption 1. Assuming that the graphs never disconnect, we can now apply Theorem 1 from Moreau (2005). It guarantees that the entries in \mathbf{y}_k will converge to a common value, and thus, through (5.24) the states \mathbf{x}_k have to approach each other so that $x_k^{(i)} - x_k^{(j)} \rightarrow 0$ as $k \rightarrow \infty$. \square

5.4.2 Application to decentralised connectivity control

In the context of decentralised connectivity control, both Lemma 4.1 on page 58 and Theorem 5.1 are very useful. Roughly speaking, they indicate that consensus algorithms with an input term, that can depend on the consensus states, eventually become *scalar*.⁶ That is, their stability and convergence properties are eventually governed by the scalar equation

$$x_{k+1} = x_k + \theta(x_k, k) \quad (5.29)$$

Since the properties of such systems are well understood, the above theorems offer interesting possibilities for the design of control laws.

⁶ And while this convergence is asymptotic, in any practical implementation of this algorithm quantisation effects will be unavoidable, hence the system should become scalar in finite time.

With this in mind we propose updating individual radii using a convex combination of their neighbours' values, plus an input term that depends on the estimated second largest eigenvalue. Specifically, we propose the following decentralised control law

$$\mathbf{r}_{k+1} = \mathbf{P}_k \mathbf{r}_k + \mu [\lambda(\mathbf{r}_k) - \lambda_*] \mathbf{1} \quad (5.30)$$

for some $\mathbf{r}_{k=0} = \mathbf{r}_0$. Here \mathbf{P}_k is now a sequence of primitive, row-stochastic averaging matrices on the graphs induced by \mathbf{r}_k , $\lambda(\mathbf{r})$ is the magnitude of the second largest eigenvalue of the averaging matrix \mathbf{P} as in (5.2) for the graph induced by \mathbf{r} , and $\mu > 0$ is a suitable control gain. We are then guaranteed by Lemma 4.1 on page 58 that the radii will converge to a common value.

The next step is thus to determine conditions on the control gain μ so that $\lambda(\mathbf{r}_k)$ will indeed converge to (a desired neighbourhood of) λ_* .

Comments At this point, the similarities to the work from the previous chapter become evident. The proposed control law has a similar structure with its local and global component. However, the local component does not include utility functions (or rather, the utility functions are the identity function) and, most importantly, the global function is neither continuous nor monotonous.

We also note that any other consensus scheme (to which Theorem 5.1 can be applied) may be used here. The proposed controller is decentralised in that each node only requires the radius information of its neighbours, information that can easily be broadcast along the communication that is necessary to run the consensus algorithm needed to estimate λ_k in the first place.

Last, (5.30) has strong similarities with the *Lur'e problem*, see for instance Narendra and Taylor (1973); Vidyasagar (2002); Khalil (1992) and references therein for the precise problem statement and the wealth of results related to it. However, the classic results cannot be applied to the problem presented here since the non-linearity does not satisfy the continuity assumption that is usually made, nor does it guarantee a unique solution (as well shall see in the next section) which is also required to apply these results. \swarrow

5.4.3 Conditions for convergence of the decentralised control law

As we have shown, it follows from the closed loop dynamics that we can assume that eventually all radii have converged to a common value. In that case, (5.30) will be reduced to a scalar equation for the whole network:

$$r_{k+1} = r_k + \mu [\lambda(r_k) - \lambda_*] \quad (5.31)$$

for some $r_{k=0} = r_0$. Note that we write $\lambda(r_k)$ since the second largest eigenvalue of the averaging matrix of the network depends on the communication radius used by the nodes. Ideally we would like to ensure that $\lambda(r_k)$ asymptotically approaches λ_* under




Figure 5.6: Plot of $\lambda(r)$, the magnitude of the second largest eigenvalue of the averaging matrix of a random (undirected) disc graph on 20 nodes as a function of the (common) communication radius r .

the assumption that the estimation part of the algorithm can be completely decoupled from the closed loop control. As we shall see, even under this considerable simplification, proving stability is nontrivial. In particular, two practical issues arise.

Quantisation The first complication arises from the following observation. Normally, with problems of this type, one makes use of the fact that the eigenvalues of the consensus matrix vary as a continuous function of the matrix entries. In what we are proposing, the entries of \mathbf{P} are either zero, or jump to some non-zero value as we adjust the communication radius of each node. In other words, the matrix entries vary abruptly as a result of the control action; consequently, the result of this is that $\lambda(r)$ also changes abruptly. Thus, it is clear that not every arbitrary second largest eigenvalue value in the $(0, 1)$ interval can be achieved through feedback of the proposed type. Rather, the network can only produce a finite set of values, corresponding to the (limited number of) different possible topologies of the network with a fixed number of nodes in fixed locations. This fact is depicted in Figure 5.6 above. The plot shows how the magnitude of the second largest eigenvalue changes with the (common) communication radius for a given random disc graph on 20 nodes. Note that the curve is not continuous, but broken up into segments. A given magnitude of the second largest eigenvalue never corresponds to just a single radius, but a range of radii. Thus the best we can hope for is to converge to some neighbourhood of λ_* . Of course, for our application, this is entirely satisfactory as both connectivity and bounds on rates of information transmission in the network are controlled using this strategy.

Monotonicity A second complication arises due to the fact that we do not precisely know the relationship between $\lambda(r)$ and r . In fact, the previous example shows that this relationship need not even be monotonic. However, it is reasonable to assume that the aforementioned relationship is approximately monotonic. This follows from the following argument. Our strategy is motivated by the intuition that as the radii of the individual nodes increase (decrease), roughly speaking, the second largest eigenvalue of \mathbf{P} also will decrease (increase). Referring to Hartfiel (1998), we know that the coefficient of ergodicity of a stochastic matrix is an upper bound on the magnitude of the second largest eigenvalue, so $|\lambda| \leq \tau(\mathbf{P})$. Recall that for a stochastic matrix \mathbf{P} , using the 1-norm, $\tau(\mathbf{P})$ is defined as

$$\tau(\mathbf{P}) = \frac{1}{2} \max_{i \neq j} \left\| \mathbf{P}^{(i)} - \mathbf{P}^{(j)} \right\|_1 \quad (5.32)$$

where $\mathbf{P}^{(i)}$ denotes the i th row of \mathbf{P} . Thus, when the rows of \mathbf{P} become ever closer to each other as measured by the 1-norm, $\tau(\mathbf{P})$ decreases, and thus the magnitude of the second eigenvalue will also eventually decrease. So even though we are not assured of a locally monotonic relationship, in principle it should still be possible to regulate the magnitude of this second eigenvalue to a neighbourhood around some target value, if we have some knowledge of the approximate manner in which $\lambda(r)$ varies with r .




Figure 5.7: Illustration of a monotonic $\lambda(r)$ curve with some relevant points highlighted relative to λ_* highlighted.

Before we present our convergence results, some further notation is helpful. Once again, to ease exposition please refer to Figure 5.7 above as we give the following definitions. Let

$$\underline{\lambda} := \inf \{ \lambda(r) : \lambda(r) \geq \lambda_* \} \quad \text{and} \quad \overline{\lambda} := \sup \{ \lambda(r) : \lambda(r) \leq \lambda_* \} \quad (5.33)$$

Then $\overline{\lambda} \leq \lambda_* \leq \underline{\lambda}$. Put simply, for any λ_* there is a feasible λ “just above” and “just below”, called $\underline{\lambda}$ and $\overline{\lambda}$ respectively. Now define the following radii

$$\underline{r} := \inf \{ r : \lambda(r) \leq \underline{\lambda} \} \quad \text{and} \quad \bar{r} := \sup \{ r : \lambda(r) \geq \overline{\lambda} \} \quad (5.34)$$

With the above definitions, the following two theorems provide simple conditions on the controller gain μ so that the system (5.31) converges to within the interval \mathcal{R} (attractivity), and stays in that interval once it has entered it (invariance). Note that estimates of these bounds may be calculated *a priori* for graphs with typical geographic distributions (or they could be estimated in real time by each node in a decentralised fashion). The important point to note is that the convergence of the controlled system is guaranteed provided that the controller gain is small enough.

$$-\kappa^{(0)}(r-\underline{r}) \leq \lambda(r) - \lambda_* \leq -\kappa^{(0)}(r-\bar{r}) \quad \text{for} \quad \underline{r} \leq r \leq \bar{r} \quad (5.35)$$

Theorem 5.2 (Invariance of \mathcal{R}) _____
Consider a scalar system described by (5.31) and let $\kappa^{(0)}$ and the interval \mathcal{R} be as defined above. Suppose that the control gain $\mu \geq 0$ is chosen such that

Then whenever $r_{k=0} \in \mathcal{R}$, the resulting sequence r_k will stay in \mathcal{R} for all $k \geq 0$. _____

Proof Suppose that $r_k \in \mathcal{R}$. We need to show that $r_{k+1} \in \mathcal{R}$. Then we will have demonstrated invariance of \mathcal{R} . We first show that $r_{k+1} \leq \bar{r}$. Since $\mu \geq 0$ and $\mu\kappa^{(0)} \leq 1$, it follows from condition (5.35) and $r_k \in \mathcal{R}$ that

$$\begin{aligned} \mu[\lambda(r_k) - \lambda_*] &\leq -\mu\kappa^{(0)}[r_k - \bar{r}] \\ &\leq \mu\kappa^{(0)}[\bar{r} - r_k] \\ &\leq \bar{r} - r_k; \end{aligned} \tag{5.37}$$

hence

$$\begin{aligned} r_{k+1} &= r_k + \mu[\lambda(r_k) - \lambda_*] \\ &\leq r_k + \bar{r} - r_k \\ &\leq \bar{r} \end{aligned} \tag{5.38}$$

Next, we show that $r_{k+1} \geq \underline{r}$. Since $\mu \geq 0$ and $\mu\kappa^{(0)} \leq 1$, it follows from condition (5.35) and $r_k \in \mathcal{R}$ that

$$\begin{aligned} \mu[\lambda(r_k) - \lambda_*] &\geq -\mu\kappa^{(0)}[r_k - \underline{r}] \\ &\geq -[r_k - \underline{r}] \\ &\geq \underline{r} - r_k; \end{aligned} \tag{5.39}$$

hence

$$\begin{aligned} r_{k+1} &= r_k + \mu[\lambda(r_k) - \lambda_*] \\ &\geq r_k + \underline{r} - r_k \\ &\geq \underline{r} \end{aligned} \tag{5.40}$$

□

To discuss convergence of the solutions of system (5.31) to \mathcal{R} we let

$$d_k := \begin{cases} \underline{r} - r_k & \text{if } r_k < \underline{r} \\ 0 & \text{if } \underline{r} \leq r_k \leq \bar{r} \\ r_k - \bar{r} & \text{if } r_k > \bar{r} \end{cases} \tag{5.41}$$

be the distance of r_k from \mathcal{R} . Then we say that r_k converges to \mathcal{R} if $\lim_{k \rightarrow \infty} d_k = 0$.

The next theorem contains a condition on μ which guarantees that all solutions of the system converge to \mathcal{R} . Use of this theorem requires that λ satisfy the following sector conditions: There exist constants $\kappa^{(2)} \geq \kappa^{(1)} > 0$ such that

$$-\kappa^{(1)}(r - \underline{r}) \leq \lambda(r) - \lambda_* \leq -\kappa^{(2)}(r - r_m) \quad \text{for } 0 < r < \underline{r} \tag{5.42}$$

$$-\kappa^{(2)}(r - r_m) \leq \lambda(r) - \lambda_* \leq -\kappa^{(1)}(r - \bar{r}) \quad \text{for } \bar{r} < r \leq \sqrt{2} \tag{5.43}$$

where $r_m := (\underline{r} + \bar{r})/2$. An illustration of these sector bounds is given in Figure 5.8.

Theorem 5.3 (Attractivity of \mathcal{R})

Consider a scalar system described by (5.31) and let $\kappa^{(0)}, \kappa^{(1)}, \kappa^{(2)}$ and the interval \mathcal{R} be as defined above. Suppose that the control gain $\mu > 0$ is chosen such that $\mu\kappa^{(0)} \leq 1$ and

$$\mu\kappa^{(2)} < 2 \quad (5.44)$$

Then every solution of (5.31) converges to \mathcal{R} .

Proof Letting $\alpha := \max\{1 - \mu\kappa^{(1)}, \mu\kappa^{(2)} - 1\}$ we will show that

$$d_{k+1} \leq \alpha d_k \quad (5.45)$$

and hence $d_k \leq \alpha^k d_{k=0}$. Since by assumption $|\alpha| < 1$, we then obtain that $\lim_{k \rightarrow \infty} d_k = 0$. Since \mathcal{R} is invariant, we need only discuss the situations for which $r_k \notin \mathcal{R}$ as well as $r_{k+1} \notin \mathcal{R}$. There are four cases to consider.

(i) $r_k < \underline{r}$ and $r_{k+1} < \bar{r}$. In this case $d_k = \underline{r} - r_k$ and

$$\begin{aligned} d_{k+1} &= \underline{r} - r_{k+1} \\ &= \underline{r} - r_k - \mu[\lambda(r_k) - \lambda_*] \\ &\leq \underline{r} - r_k - \mu\kappa^{(1)}[\underline{r} - r_k] \\ &\leq (1 - \mu\kappa^{(1)})[\underline{r} - r_k] \\ &\leq (1 - \mu\kappa^{(1)})d_k \end{aligned} \quad (5.46)$$

that is $d_{k+1} \leq (1 - \mu\kappa^{(1)})d_k$, and thus (5.45) holds.

(ii) $r_k < \underline{r}$ and $r_{k+1} > \bar{r}$. In this case $d_k = \underline{r} - r_k$ and

$$\begin{aligned} d_{k+1} &= r_{k+1} - \bar{r} \\ &= r_k + \mu[\lambda(r_k) - \lambda_*] - \bar{r} \\ &\leq r_k + \mu\kappa^{(2)}[r_m - r_k] - \bar{r} \\ &\leq (1 - \mu\kappa^{(2)})r_k + \mu\kappa^{(2)}r_m - \bar{r} \end{aligned} \quad (5.47)$$

Recalling that $\mu\kappa^{(2)} < 2$ and $r_m = (\underline{r} + \bar{r})/2$, we can see that

$$\begin{aligned} \mu\kappa^{(2)}r_m - \bar{r} &= -\left(1 - \frac{\mu\kappa^{(2)}}{2}\right)\bar{r} + \frac{\mu\kappa^{(2)}}{2}\underline{r} \\ &\leq -\left(1 - \frac{\mu\kappa^{(2)}}{2}\right)\underline{r} + \frac{\mu\kappa^{(2)}}{2}\underline{r} \\ &\leq -(1 - \mu\kappa^{(2)})\underline{r} \end{aligned} \quad (5.48)$$

Hence

$$\begin{aligned} d_{k+1} &\leq (1 - \mu\kappa^{(2)})r_k - (1 - \mu\kappa^{(2)})\underline{r} \\ &\leq (\mu\kappa^{(2)} - 1)d_k \end{aligned} \quad (5.49)$$

and thus (5.45) holds.

(iii) $r_k > \bar{r}$ and $r_{k+1} < \bar{r}$. In this case $d_k = r_k - \bar{r}$ and

$$\begin{aligned}
 d_{k+1} &= \underline{r} - r_{k+1} \\
 &= \underline{r} - r_k - \mu \left[\lambda(r_k) - \lambda_* \right] \\
 &\leq \underline{r} - r_k - \mu \kappa^{(2)} [r_m - r_k] \\
 &\leq -(1 - \mu \kappa^{(2)}) r_k + \underline{r} - \mu \kappa^{(2)} r_m
 \end{aligned} \tag{5.50}$$

Again, we can see that since $\mu \kappa^{(2)} < 2$

$$\begin{aligned}
 \underline{r} - \mu \kappa^{(2)} r_m &= \left(1 - \frac{\mu \kappa^{(2)}}{2} \right) \underline{r} - \frac{\mu \kappa^{(2)}}{2} \bar{r} \\
 &\leq \left(1 - \frac{\mu \kappa^{(2)}}{2} \right) \bar{r} - \frac{\mu \kappa^{(2)}}{2} \bar{r} \\
 &\leq (1 - \mu \kappa^{(2)}) \bar{r}
 \end{aligned} \tag{5.51}$$

Hence

$$\begin{aligned}
 d_{k+1} &\leq -(1 - \mu \kappa^{(2)}) r_k + (1 - \mu \kappa^{(2)}) \bar{r} \\
 &\leq (\mu \kappa^{(2)} - 1) d_k
 \end{aligned} \tag{5.52}$$

and thus (5.45) holds.

(iv) $r_k > \bar{r}$ and $r_{k+1} > \bar{r}$. In this case $d_k = \underline{r}_k - \bar{r}$ and

$$\begin{aligned}
 d_{k+1} &= r_{k+1} - \bar{r} \\
 &= r_k + \mu \left[\lambda(r_k) - \lambda_* \right] - \bar{r} \\
 &\leq r_k - \bar{r} + \mu \kappa^{(1)} [\bar{r} - r_k] \\
 &\leq (1 - \mu \kappa^{(1)}) [r_k - \bar{r}] \\
 &\leq (1 - \mu \kappa^{(1)}) d_k
 \end{aligned} \tag{5.53}$$

that is $d_{k+1} \leq (1 - \mu \kappa^{(1)}) d_k$, and thus (5.45) holds.

□

In summary, the theorem gives a condition on the control gain so that the closed loop system (5.31) converges to the interval \mathcal{R} .

Comments If $\lambda(r)$ is not monotonic with r then it is possible that $\lim_{r \rightarrow \bar{r}_-} \lambda(r) > \lambda_*$ where the notation means that the limit is taken from the left; see Figure 5.9 on the facing page. If this occurs, one cannot satisfy (5.35) with any $\kappa^{(0)} > 0$. In this case (5.35) can be satisfied by replacing \bar{r} with \bar{r}_ε where $\bar{r}_\varepsilon = \bar{r} + \varepsilon$ and $\varepsilon > 0$; of course $\kappa^{(0)}$ will depend on ε ; see Figure 5.9. A similar remark holds if $\lim_{r \rightarrow \underline{r}_+} \lambda(r) < \lambda_*$.




Figure 5.9: Illustration of $\lambda(r)$ curve that is not monotonic.

Furthermore, it is possible that with the above control law the network may accidentally become disconnected. The closer λ_* is to one, the more likely this may happen: For instance, assume at time step k the estimated λ_k is smaller than λ_* . In that case, all the nodes will reduce their radius by a certain amount (that is, by $\mu[\lambda_k - \lambda_*]$). Now, if the updated radii are so small that a particularly “outlying” node becomes “out of reach”, the graph will disconnect.

However, in general, the disconnection of the graph can easily be prevented by setting a certain minimum radius that the nodes are allowed to use: this would be the smallest common radius (plus, maybe, a safety margin) that would still guarantee connectedness of the network, i.e. it would correspond to the largest inter-node distance. This information can either be pre-programmed into the nodes at the time of deployment (if a the corresponding maximum inter-node distance can be guaranteed), or after deployment. In any case, this only needs to be done once, as we assume that the nodes do not change their position after deployment. /

5.5 Simulation results

To conclude this section, we now present some simulation results. Most of the plots shown in this section are based on random disc graphs of 50 nodes, with initial radii uniformly distributed in $[0.1, 0.6]$, and $\lambda_* = 0.8$.

First we show a series of plots to illustrate the pertinent features of our stability proofs, then we show the general performance of our proposed controller, and finally examples of modified control objectives.

5.5.1 Example 1: Controller stability bounds

Figure 5.10 and Figure 5.11 show an experimentally obtained $\lambda(r)$ curve, the second figure being a close-up view of the first. Picking $\lambda_* = 0.8$ we indicate the values of $\underline{\lambda}$ and $\bar{\lambda}$, as well as \underline{r} , \bar{r} and r_m with dotted lines.




Figure 5.10: Actual $\lambda(r)$ of a random disc graph on 50 nodes, with an example of the bounds as required for by Theorems 5.2 and 5.3 drawn for $\lambda_* = 0.8$.

We then determined the bounds $\kappa^{(0)}$, $\kappa^{(1)}$ and $\kappa^{(2)}$ on the curve, which are indicated by the thicker lines, similar to Figure 5.8. The actual values of those bounds are $\kappa^{(0)} \simeq 14.3$, $\kappa^{(1)} \simeq 0.17$ and $\kappa^{(2)} \simeq 8.72$.⁷ When controlling the nodes' radii with (5.30), Theorem 5.2 requires that μ has to be less than $1/\kappa^{(0)} \simeq 0.067$ to guarantee invariance of the corresponding interval $\mathcal{R} \simeq [0.321, 0.322]$. Attractivity of \mathcal{R} according to Theorem 5.3 in turn requires μ to be less than $2/\kappa^{(2)} \simeq 0.23$.

Thus setting $\mu = 0.05$, we re-initialised the network with randomly distributed radii in the $[0.1, 0.6]$ and ran the controller on the network. As we can see in Figure 5.12 on the next page the convergence of both the radii and λ is smooth and fast.

⁷ Note that tighter bounds can be found.




Figure 5.11: Magnified view of the region around (λ_*, r_m) from the previous plot.




Figure 5.12: Evolution of λ_k and the individual nodes' radii $r_k^{(i)}$ in the 50 node network analysed in Figure 5.10, for $\lambda_* = 0.8$, with $\mu = 0.05$.

5.5.2 Example 2: Combining Control and Estimation

In the previous example we displayed the converged values of the estimation scheme. To show in more detail how estimation and control scheme work together, we present Figure 5.13. Plotted is again the evolution of the nodes' radii under control action (5.30) as well as the estimates of λ , shown in the upper subplot. These estimates were calculated as described in Section 5.3. We allowed 100 time steps for the estimation scheme to converge, before taking a control action based on the new estimates.

It can be seen that after every topology change all nodes' estimates converge to a common value and that the control scheme successfully regulates the second largest eigenvalue of the network to $\lambda_* = 0.8$.




Figure 5.13: Evolution of the estimates of λ_k and the individual nodes' radii $r_k^{(i)}$, as the controller updates the radii every 100 iterations of the estimation scheme, for $\lambda_* = 0.8$, with $\mu = 0.05$

5.5.3 Further Examples of control

Next, we present another example that depicts how the (true value of the) second largest eigenvalue in magnitude and the nodes' radii change over time, as the nodes control their radii using (5.30).

Figure 5.14 shows a situation where $\lambda_* = 0.5$ was required. As this represents a very densely connected network, all nodes had to increase their radius. In turn, in Figure 5.15 we start off with an extremely highly connected network (it was almost fully connected), and all nodes have to significantly decrease their radii to achieve the desired $\lambda_* = 0.8$.




Figure 5.14: Evolution of λ_k and the individual nodes' radii $r_k^{(i)}$ in a network of 50 nodes for $\lambda_* = 0.5$, with $\mu = 0.05$.




Figure 5.15: Evolution of λ_k and the individual nodes' radii $r_k^{(i)}$ in a network of 50 nodes with very large initial radii, for $\lambda_* = 0.8$, with $\mu = 0.05$.

The plots in Figure 5.16 on the following page show a scenario where the network had to react to a change in topology: At $k = 30$ we randomly removed half of the nodes from the network, thus reducing the graph size to 25 nodes. The resulting network's second largest eigenvalue in magnitude is larger than desired (i.e. it is less connected), and thus the controller compensates this by increasing the remaining nodes' radii until $\lambda_* = 0.8$ is achieved again.




Figure 5.16: Evolution of λ_k and the individual nodes' radii $r_k^{(i)}$ in a network of 50 nodes, where 25 nodes are removed at $k = 30$ (for $\lambda_* = 0.8$, with $\mu = 0.05$).

5.5.4 Validation of control results

In Figure 5.17 on the next page we compare the converged radii of our controller for several different λ_* (circles) with the second largest eigenvalue in magnitude of the averaging matrix of random disc graphs created with different initial radii (crosses). Until now we have only shown individual results from single instances of graphs. This plot is to demonstrate that our estimation and control scheme works over a whole range of set points, for any number of trials.

The data points marked by crosses were obtained as follows. Picking 17 different values of r we generated 1000 random geometric graphs (on 50 nodes) for each radius. Next, we calculated the second largest eigenvalue of the resulting averaging matrix of each graph $\lambda(r)$, and finally plotted the average value against the initial r value used. In turn, the data points marked by circles were generated by choosing 14 different values for λ_* , generating 1000 graphs and running the control algorithm on the network. The resulting converged (common) radii $r_{\text{conv}}(\lambda_*)$ were then averaged and the value plotted against the particular λ_* chosen.

As all points appear to lay on the same curve, the plot indicates that nodes radii set by the controller indeed converge to the corrected value over the entire range of sensible λ_* values.




Figure 5.17: Crosses \times : average $\lambda(r)$ of 1000 geometric graphs on 50 nodes created with common radius r . Circles \circ : Average converged radii r_{conv} after control targeted at different λ_* values, for 1000 trials each (where the initial radii were randomly distributed).

5.5.5 Examples of other control objectives

As we mentioned in Section 5.2, our control scheme is general enough to allow objectives other than a common radius while achieving a desired λ_* .

Imagine a situation in which some nodes are equipped with a longer-lasting power supply and we can allow those nodes to have a larger radius than most of the other nodes in the network. This would correspond to weighting the nodes' radii in the averaging scheme. It is possible to include such weighting in our framework, and all the proofs directly hold with but a small modification, Knorn *et al.* (2009c). An example of this is given in Figure 5.18, where by design we wish one node to have twice the radius as the others, and one node half the radius. As can be seen, the eigenvalue of the network converges quickly to its desired value of $\lambda_* = 0.8$, and the nodes radii converge to a common value with the exception of the two nodes of different weighting.

Comment Note that such a weighting will result — contrary to the other cases — in a *directed* network (that is, a non-symmetric averaging matrix), even in steady state. As we mentioned earlier it is an important feature of our algorithms that they work in both undirected and directed networks. /

Finally we now present an example where a completely different control objective is desired. Regulating the second largest eigenvalue in magnitude, here we do not care about the radii but rather about the number of neighbours of each node. In Figure 5.19 on the following page we required the nodes to achieve consensus on the number of neighbours,




Figure 5.18: Evolution of λ_k and the individual nodes' radii $r_k^{(i)}$ in a network of 50 where two nodes were to have twice resp. half the radius as their peers. Again, $\lambda_* = 0.8$ and $\mu = 0.05$.

rather than the radii. Although one needs to redo the proof of stability, we can see that the network converges to a stable solution.




Figure 5.19: Evolution of λ_k and the individual nodes' radii $r_k^{(i)}$, 50 nodes, consensus on number of neighbours.

5.6 Conclusion

In this chapter we have presented a general framework for controlling the topological properties of a network of distributed sensors. This work is closely related to the contents

of the previous chapter, with the important difference that the global term is not provided “externally” but estimated in a distributed fashion by the network agents themselves. As before, our framework breaks free of many of the assumptions of previous work such as graph symmetry, and utilises only simple ideas from control and estimation to regulate important graph properties. Conditions for the stability of our algorithms are presented. Roughly speaking, these results state that if the nodes are not too aggressive in the manner in which they expand or contract their neighbourhood set, stability is assured. This bears a strong resemblance with the growth bounds that were required in the previous chapter in order to calculate the controller gains.

Limitations

While the results reported in this chapter are certainly promising, there are a number of limitations to our theoretical contributions. The first concerns the estimation of the second largest eigenvalue in magnitude, where it would be beneficial if an estimation scheme could be found that can estimate it irrespective of whether it is real or complex (non-real) valued and hence does not require a decision-heuristic as presented here.

Next, the separation of time scales between estimation and control scheme may be an unnecessarily restrictive assumption. In fact, initial tests have shown that estimation and control scheme may well be “interleaved” in the sense that single iterations of each scheme can be performed in alternation without compromising convergence (provided the gains are small enough).

Last, the overall convergence proof here relies on the convergence to a scalar equation, which makes it difficult to derive precise convergence rates for the overall problem. It may be an interesting problem to attempt to prove convergence without this intermediate step and derive concrete convergence rates.

The last chapter of this thesis will discuss three applications where the main results from the previous chapters are applied to a number of real-world problems.

Applications

This chapter presents three practical applications for some of the results presented in the previous three chapters. In particular, they involve stability conditions for a power control algorithm (application of our CLCLF result), cooperative control of emissions in a fleet of plug-in hybrid electric vehicles as well as a real implementation of a small network of wireless motes (as applications for the cooperative control results).

Chapter contents

- 6.1 Stability of the Foschini-Miljanic algorithm
 - 6.2 Emissions control in a fleet of Hybrid Vehicles
 - 6.3 Real-world implementation of cooperative control
-

6.1 Stability of the Foschini-Miljanic algorithm

The first application we discuss uses one of our common linear co-positive Lyapunov function results from Chapter 3 to derive conditions for stability in the presence of time-varying time-delays and arbitrary switching in a popular distributed power control algorithm for wireless communication networks. This section is based on joint work with Dr. A. Zappavigna, Prof. P. Colaneri¹, Dr. T. Charalambous² and Prof. R. Shorten; it is accepted for publication in the *Automatica* journal, Zappavigna *et al.* (2011).

6.1.1 Introduction

Some *Code Division Multiple Access* (CDMA) based power control algorithms aim to assign power to wireless nodes in a distributed fashion, while guaranteeing a certain *Quality of Service* (QoS), Schulze and Lüders (2005). In real communication systems, especially ad-hoc networks, distributed algorithms require communication among the nodes. But processing time (coding and decoding), propagation delays and waiting for availability of

¹ Dr. Zappavigna and Prof. Colaneri are with the *Dipartimento di Elettronica e Informazione, Politecnico di Milano*, Italy.

² Dr. Charalambous was with the *Department of Computing, Imperial College London*, United Kingdom.

channels for transmission all introduce delays into the network. Additionally, the nodes may be mobile, entering or leaving the network, causing the network topology to change constantly. Hence, any stability analysis of distributed algorithms for such realistic situations should consider time-delays in the network and changing network topologies.

The authors in Foschini and Miljanic (1993) proposed a power control algorithm, the now well known Foschini-Miljanic (FM) algorithm, that provides distributed on-line power control of wireless networks with user-specific *Signal-to-Interference-and-Noise-Ratio* (SINR) requirements. Furthermore, this algorithm yields the minimum transmitter powers that satisfy these requirements.

Previous work

As we shall see, this study will involve switched positive systems where the states are delayed. Systems with time-delays naturally occur in many applications and have been studied extensively over the past few decades, see for instance the book by Lewis and Anderson (1980); Hale and Lunel (1993); Hennet and Tarbouriech (1997); Haddad and Chellaboina (2004); Hövel (2010) and the book by Mahmoud (2010).

In the context of switched systems, various types of delays are usually considered, in particular single, constant delays (Li *et al.*, 2009) or multiple but constant delays (Sun *et al.*, 2008; Liu *et al.*, 2008; Ding and Shu, 2010). The recent result by (Sun *et al.*, 2008) discusses switched systems with time-varying time-delays, but focuses on finding stabilising switching laws and hence does not cover the arbitrary switching case. Concerning the Foschini-Miljanic algorithm, it was recently shown in Charalambous *et al.* (2008) that it is globally asymptotically stable for arbitrarily large but *constant* time-delays, and the article did not consider time changing network topologies.

Contributions

In this section, making use of recent advancements in positive linear systems and in particular Theorem 3.2 from Chapter 3, we consider both the effects of time-varying delays and changing network topologies (in other words, arbitrary switching). For that we present a new theoretical result concerning the stability of such systems. This result is then used to show that the Foschini-Miljanic algorithm is globally asymptotically stable even under those harder, more realistic conditions, provided a condition similar to Theorem 3.2 is satisfied. Our results are of practical importance when designing wireless networks in changing environments, as is typically the case for CDMA networks.

Structure

The remainder of this application section is structured as follows: Section 6.1.2 provides some helpful mathematical preliminaries. Then, we introduce the channel model used

for modelling the wireless communications as well as the Foschini-Miljanic power control algorithm. In Section 6.1.4, a stability condition is derived for the Foschini-Miljanic algorithm, showing its stability under arbitrary switching and time-varying delays. Finally, an example as well as some concluding remarks are given.

6.1.2 Mathematical preliminaries

In what follows, we will establish the mathematical framework for our study and give a useful result on positive systems that is needed to prove our later results. We shall deviate slightly from our usual notation in that the variable t now denotes the (continuous) time variable, so that $\mathbf{x}(t)$ is the value of \mathbf{x} at time t . Subscripts are either used to index subsystems in a switched system, to indicate different delayed states, or to denote specific switching instants. In general, it should be clear from context and the explanations we give when defining new variables as to what the index is referring to.

We shall consider the following type of linear system with m different delayed states whose time-delays are time-varying:

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \sum_{k=1}^m \mathbf{B}_k \mathbf{x}(t - \tau_k(t)), \quad t \geq 0 \quad (6.1a)$$

$$\mathbf{x}(t) = \boldsymbol{\varphi}(t) \succeq \mathbf{0}, \quad t \in [-\bar{\tau}, 0] \quad (6.1b)$$

where $\mathbf{x}(t) \in \mathbb{R}_{\geq 0}^n$, $\mathbf{A} \in \mathbb{R}^{n \times n}$ is a Metzler matrix, $\mathbf{B}_k \in \mathbb{R}_{\geq 0}^{n \times n}$ are non-negative matrices for all $k = 1, \dots, m$, $\boldsymbol{\varphi}(\cdot)$ is a bounded, piecewise continuous vector function and the delays $\tau_k(t)$ are assumed to satisfy:

Assumption 6.1 (Bounded time-delays) _____

All the $k = 1, \dots, m$ time-varying time-delays $\tau_k(t)$ are bounded, piecewise continuous functions in t , satisfying

$$0 \leq \tau_k(t) \leq \bar{\tau}_k \leq \bar{\tau} \quad \text{for all } t \geq 0 \quad (6.2)$$

where $\bar{\tau} = \max_k \{\bar{\tau}_k\}$. _____

Comments Systems of the type (6.1) are referred to as *delay differential equations* or *functional differential equations*; an extensive overview over such systems can be found in Hale and Lunel (1993); Kuang (1993); Diekmann *et al.* (1995).

Furthermore, while for most practical applications piecewise continuity of both the initial condition function $\boldsymbol{\varphi}(\cdot)$ and the time-delays $\tau_k(t)$ will suffice, all results will in fact hold for locally Lebesgue integrable functions, Rudin (1976); Rami (2009). /

Recall that a dynamical system is said to be *positive* if its state trajectories remain in the positive orthant for all $t \geq 0$ (provided that the initial condition is positive). Thanks

to \mathbf{A} being Metzler and the \mathbf{B}_k being non-negative, it is easy to show that the system above is indeed positive, see for instance Rami (2009).

We can now present a useful result on switched positive systems with time-varying time-delays that are based on (6.1), where both the system matrix \mathbf{A} and delay matrices \mathbf{B}_k switch arbitrarily (but not infinitely fast). Given N constituent subsystems we make the common assumption that the switching instants are defined in all the real time axes and that $\inf_k(t_{k+1} - t_k) > 0$, where t_{k+1} and t_k are two consecutive switching instants, so that the switching rule has no accumulation points.

The following theorem states that the existence of a common linear co-positive Lyapunov function $v(\mathbf{x}) = \mathbf{c}^\top \mathbf{x}$ with $\mathbf{c} \succ \mathbf{0}$ for all un-delayed modes of the system is sufficient to guarantee the asymptotic stability of the system for bounded time-varying delays and arbitrary switching.

Theorem 6.1 (Stability of switched positive linear systems with time-varying delays) _____

Consider the switched positive system with time-varying time-delays for $t \geq 0$

$$\dot{\mathbf{x}}(t) = \mathbf{A}_{\sigma(t)} \mathbf{x}(t) + \sum_{k=1}^m \mathbf{B}_{k,\sigma(t)} \mathbf{x}(t - \tau_k(t)) \quad (6.3a)$$

$$\mathbf{x}(t) = \boldsymbol{\varphi}(t) \succeq \mathbf{0}, \quad t \in [-\bar{\tau}, 0] \quad (6.3b)$$

where $\mathbf{x}(t) \in \mathbb{R}_{\geq 0}^n$, $\sigma : \mathbb{R} \rightarrow \{1, \dots, N\}$ is some (piecewise constant and left-continuous) switching signal (defined in all the real time axes and with $\inf_k(t_{k+1} - t_k) > 0$, where t_{k+1} and t_k are two consecutive switching instants), $\mathbf{A}_i \in \mathbb{R}^{n \times n}$ are Metzler and $\mathbf{B}_{k,i} \in \mathbb{R}_{\geq 0}^{n \times n}$ are non-negative matrices, $i = 1, \dots, N$, and the delays $\tau_k(t)$ are assumed to satisfy Assumption 6.1. If there exists a strictly positive vector \mathbf{c} such that

$$\mathbf{c}^\top \left(\mathbf{A}_i + \sum_{k=1}^m \mathbf{B}_{k,i} \right) \prec \mathbf{0}, \quad \forall i = 1, \dots, N \quad (6.4)$$

then system (6.3) is asymptotically stable. _____

Proof The full proof of this theorem is given in Zappavigna *et al.* (2011). To give a rough outline, the main idea of the proof is to make use of certain monotonicity and order preserving properties exhibited by these systems and their counterparts with *constant* time-delays. The switched system is examined between each two consecutive switching instants and it is shown that it decreases exponentially in each of these time intervals, from which overall stability can then be deduced.

Comment Note that with the assumptions of the theorem, system (6.3) will also remain positive throughout time. /

Now, given this result, the question would be how to check for the existence of such Lyapunov function. From the third chapter, recall Theorem 3.2 which provided a (necessary

and sufficient) test for the existence of a common linear co-positive Lyapunov function. The following corollary is just a slight reformulation of that theorem in order to fit the current setting, reproduced here mainly for convenience:

Corollary 6.1 (CLCLF existence)

Given N Metzler matrices \mathbf{A}_i and $m \cdot N$ non-negative matrices $\mathbf{B}_{k,i}$, then there exists a strictly positive vector $\mathbf{c} \succ \mathbf{0}$ such that $\mathbf{c}^\top (\mathbf{A}_i + \sum_{k=1}^m \mathbf{B}_{k,i}) =: \mathbf{c}^\top \tilde{\mathbf{A}}_i \prec \mathbf{0} \forall i = 1, \dots, N$ if and only if $\tilde{\mathbf{A}}_s (\tilde{\mathbf{A}}_1, \dots, \tilde{\mathbf{A}}_N)$ is Hurwitz for all $s \in \mathcal{S}_{n,N}$.

Proof See Theorem 3.2 on page 41.

6.1.3 Wireless communications

Having laid out some necessary mathematical groundwork, let us now present a model of the wireless communications and later the famous Foschini-Miljanic power control algorithm.

Channel model

We consider a network in which the links are unidirectional and each node is supported by an omnidirectional antenna. The link quality is measured by the Signal-to-Interference-and-Noise-Ratio (SINR). Let \mathcal{S} and \mathcal{R} denote all transmitters and receivers in the network, respectively. In a network with n communication pairs ($n = |\mathcal{S}| = |\mathcal{R}|$), the channel gain on the link between transmitter $i \in \mathcal{S}$ and receiver $j \in \mathcal{R}$ is denoted by $g^{(ij)}$ and incorporates the mean path-loss as a function of distance, shadowing and fading, as well as cross-correlations between signature sequences. All the $g^{(ij)}$ are positive (since all nodes are equipped with omnidirectional antennae) and can take values in the range $(0, 1]$. Without loss of generality, we assume that the intended receiver of transmitter i is also indexed by i . The power level used by transmitter i is denoted by $p^{(i)}$, and $\nu^{(i)}$ denotes the variance of thermal noise at the receiver i , which is assumed to be an additive Gaussian noise.

The interference power at the i th receiver consists of both the interference caused by all the other transmitters in the network $\sum_{j \neq i} g^{(ji)} p^{(j)}$ and the thermal noise $\nu^{(i)}$ in node i 's receiver. That means the SINR at the receiver i is

$$\text{SINR}^{(i)} = \frac{g^{(ii)} p^{(i)}}{\sum_{j \neq i} g^{(ji)} p^{(j)} + \nu^{(i)}} \quad (6.5)$$

Due to the unreliability of the wireless links, it is necessary to ensure Quality of Service (QoS) in terms of the SINR in wireless networks. That is, a transmission from transmitter i to its corresponding receiver is successful (error-free) if the SINR at the receiver with respect to that transmission is greater than or equal to the *capture ratio* $\gamma^{(i)}$, which depends




Figure 6.1: Illustration of a wireless ad-hoc network with 5 communication pairs. The channel gains for each pair $\{S_i \rightarrow R_i\}$ is shown as well as the interference caused by S_1 on the other four receivers.

on the modulation and coding characteristics of the radio. In other words, it is required that

$$\frac{g^{(ii)}p^{(i)}}{\sum_{j \neq i} g^{(ji)}p^{(j)} + \nu^{(i)}} \geq \gamma^{(i)} \quad (6.6)$$

Inequality (6.6) describes the QoS requirement of a communication pair (i, i) while a transmission takes place. After manipulation, (6.6) becomes

$$p_i \geq \gamma^{(i)} \left(\sum_{j \neq i} \frac{g^{(ji)}}{g^{(ii)}} p^{(j)} + \frac{\nu_i}{g^{(ii)}} \right) \quad (6.7)$$

In matrix form, for a network consisting of n communication pairs, this can be written as

$$\mathbf{p} \succeq \mathbf{\Gamma} \mathbf{Z} \mathbf{p} + \boldsymbol{\eta} \quad (6.8)$$

where we define $\mathbf{p}^\top = (p^{(1)} \dots p^{(n)})$; $\mathbf{\Gamma} = \text{diag}\{\gamma^{(i)}\}$; $z^{(ij)} = g^{(ji)}/g^{(ii)}$ if $i \neq j$, zero otherwise; and $\boldsymbol{\eta}^\top = (\eta^{(1)} \dots \eta^{(n)})$ with $\eta^{(i)} = \gamma^{(i)}\nu^{(i)}/g^{(ii)}$. Finally, letting $\mathbf{C} := \mathbf{\Gamma} \mathbf{Z}$,

$$(\mathbf{I} - \mathbf{C})\mathbf{p} \succeq \boldsymbol{\eta} \quad (6.9)$$

We note that \mathbf{C} has strictly positive off-diagonal elements which implies that it is irreducible. By the Perron-Frobenius Theorem (Horn and Johnson, 1985) we then have that the spectral radius of \mathbf{C} is a simple eigenvalue, while the corresponding eigenvector is positive elementwise. A necessary and sufficient condition for existence of a non-negative solution to inequality (6.9) for every positive vector $\boldsymbol{\eta}$ is that $(\mathbf{I} - \mathbf{C})^{-1}$ exists and is non-negative. However, $(\mathbf{I} - \mathbf{C})^{-1} \succeq \mathbf{0}$ if and only if the spectral radius $\rho(\mathbf{C}) < 1$, or, equivalently, $(\mathbf{C} - \mathbf{I})$ is Hurwitz (since $(\mathbf{C} - \mathbf{I})$ is Metzler), see Horn and Johnson (1991).

The Foschini-Miljanic power control algorithm

The Foschini-Miljanic (FM) algorithm is given by the following distributed power update formula Foschini and Miljanic (1993):

$$\frac{dp^{(i)}(t)}{dt} = \kappa^{(i)} \left[-p_i(t) + \gamma^{(i)} \left(\sum_{j \neq i} \frac{g^{(ji)}}{g^{(ii)}} p^{(j)}(t) + \frac{\nu^{(i)}}{g^{(ii)}} \right) \right] \quad (6.10)$$

where $\kappa^{(i)} > 0$ denote the proportionality constants and $\gamma^{(i)}$ denote the desired SINR. It is assumed that each node i has only knowledge of the interference at its own receiver.

In matrix form, for a given network configuration this yields

$$\dot{\mathbf{p}}(t) = -\mathbf{K}(\mathbf{I} - \mathbf{C})\mathbf{p}(t) + \boldsymbol{\eta} \quad (6.11)$$

Since the transmitter uses measurements from its intended receiver, delays are inevitably introduced into the system for a number of reasons such as processing time (coding/decoding), propagation delays and availability of the channel for transmission. Consequently, a realistic analysis of the algorithm must consider, time-varying delays:

$$\frac{dp^{(i)}(t)}{dt} = \kappa^{(i)} \left[-p^{(i)}(t) + \gamma^{(i)} \left(\sum_{j \neq i} \frac{g^{(ji)}}{g^{(ii)}} p^{(j)}(t - \tau^{(j)}(t)) + \frac{\nu^{(i)}}{g^{(ii)}} \right) \right] \quad (6.12)$$

where we assume that $\tau^{(i)}(t)$ satisfy Assumption 6.1. In matrix form this can be written as

$$\dot{\mathbf{p}}(t) = -\mathbf{K}\mathbf{p}(t) + \mathbf{K} \left(\sum_{k=1}^n \mathbf{B}_k \mathbf{p}(t - \tau_k(t)) + \boldsymbol{\eta} \right) \quad (6.13)$$

where $\mathbf{K} = \text{diag} \{ \kappa^{(i)} \}$, and $b_k^{(ij)}$ is zero if $j = k$ or $i \neq k$, or equal to $\gamma^{(k)} g^{(ji)} / g^{(kk)}$ otherwise. Note that $\sum_{k=1}^n \mathbf{B}_k = \mathbf{C}$.

Assuming feasibility of the solution, and defining $\mathbf{x}(t) = \mathbf{p}_* - \mathbf{p}(t)$ to describe the deviation from the desired power levels $\mathbf{p}_* = (\mathbf{I} - \mathbf{C})^{-1} \boldsymbol{\eta} \succ \mathbf{0}$ in order to satisfy (6.9), then the stability of (6.13) is equivalent to and can be assessed by study of the following system:

$$\dot{\mathbf{x}}(t) = -\mathbf{K}\mathbf{x}(t) + \sum_{k=1}^n \mathbf{K}\mathbf{B}_k \mathbf{x}(t - \tau_k(t)) \quad (6.14)$$

for which it is easy to see that the origin is the equilibrium. If its initial condition is non-negative (which can be guaranteed by starting from all zero power levels) then (6.14) defines a positive system as the diagonal matrix $-\mathbf{K}$ is Metzler and the $\mathbf{K}\mathbf{B}_k$ are non-negative.

6.1.4 Main results

Our main result states the following: In some situations all the possible variations in the gain matrix may be known *a priori*, and thus there is a finite number of configurations

that characterise the possible configuration of the system. In such situations, the next theorem provides a sufficient condition for stability of the Foschini-Miljanic algorithm under time-varying delays and when the topology changes arbitrarily among N different configurations.

Theorem 6.2 (Stability of the FM-Algorithm) _____

Consider a set of N different network configurations that are described by matrices $\mathbf{C}_i = \sum_{k=1}^n \mathbf{B}_{k,i}$, where $i = 1, \dots, N$, and let $\mathbf{A}_i := \mathbf{C}_i - \mathbf{I}$.

If the $\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$ are Hurwitz for all $s \in \mathcal{S}_{n,N}$, then the power control algorithm (6.13) is asymptotically stable under arbitrary switching (defined in all the real time axes and with $\inf_k (t_{k+1} - t_k) > 0$, where t_{k+1} and t_k are two consecutive switching instants), for any time-varying delays $\tau_k(t)$ satisfying Assumption 6.1, for any initial states $p_i(0) \geq 0$, and for any proportionality constants $\kappa^{(i)} > 0$. _____

Proof By construction, all \mathbf{A}_i are Metzler matrices. $\mathbf{A}_\sigma(\mathbf{A}_1, \dots, \mathbf{A}_N)$ being Hurwitz for all $s \in \mathcal{S}_{n,N}$ is a necessary and sufficient condition, according to Corollary 6.1 to say that there exists a positive vector $\mathbf{c} \succ \mathbf{0}$ such that $\mathbf{c}^\top (-\mathbf{I} + \sum_{k=1}^n \mathbf{B}_{k,i}) \prec \mathbf{0}$ for all i . This again also means that since \mathbf{K} is a diagonal matrix with strictly positive entries, then $\tilde{\mathbf{c}}^\top (-\mathbf{K} + \sum_{k=1}^n \mathbf{K} \mathbf{B}_{k,i}) \prec \mathbf{0}$ for all i , where $\tilde{\mathbf{c}}^\top = \mathbf{c}^\top \mathbf{K}^{-1} \succ \mathbf{0}$.

By Theorem 6.1, comparing (6.14) to (6.3), this is sufficient to guarantee stability. \square

Comment As we mentioned earlier, Theorem 6.2 may also be formulated in terms of feasibility of suitably defined linear programming problem. One such program might be for example: Find a vector $\mathbf{c} \succ \mathbf{0}$ such that $\mathbf{c}^\top [\mathbf{A}_1 \ \dots \ \mathbf{A}_N \ - \mathbf{I}] \prec \mathbf{0}$. /

6.1.5 Example

To illustrate the theoretical result presented by Theorem 6.2, we now consider a three dimensional model consisting of three modes such that the above stability condition is fulfilled. It is given by the following matrices

$$\mathbf{C}_1 = \begin{bmatrix} 0 & 0.18 & 0.23 \\ 0.31 & 0 & 0.04 \\ 0.22 & 0.12 & 0 \end{bmatrix}, \quad \mathbf{C}_2 = \begin{bmatrix} 0 & 0.35 & 0.15 \\ 0.40 & 0 & 0.45 \\ 0.37 & 0.53 & 0 \end{bmatrix}, \quad \mathbf{C}_3 = \begin{bmatrix} 0 & 0.36 & 0.61 \\ 0.47 & 0 & 0.28 \\ 0.71 & 0.26 & 0 \end{bmatrix} \quad (6.15)$$

From Theorem 6.2, if for all $s \in \mathcal{S}_{3,3}$ the matrices $\mathbf{C}_s(\mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3)$ have a spectral radius less than one then the power control algorithm (6.13) is asymptotically stable under arbitrary switching. In the example here, indeed the largest spectral radius over all matrices $\max_s \{\rho(\mathbf{C}_s(\mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3))\} \simeq 0.985 < 1$ (corresponding to the permutation $s = (3, 2, 3)$) and thus the resulting system would be asymptotically stable under arbitrary switching.




Figure 6.2: Simulation of the switched network represented by the matrices in (6.15). The plot shows the evolution of the deviation from the desired power levels $p_{*,\sigma(t)}^{(i)}$. The switching sequence $\sigma(t)$ is also shown with the dash-dotted line (that is, if $\sigma(t) = 1$ then the network is represented by matrix C_1 , and so on).

Figure 6.2 above confirms this. It shows the results from a simulation run, plotting the deviations from the desired power levels $\Delta p^{(i)}(t) := p_{*,\sigma(t)}^{(i)} - p^{(i)}(t)$ as a function of time for each of the three states, where $p_{*,k}^{(i)}$ denotes the desired power level of the i th state in the k th subsystem. The switched system used was based on the above matrices, where the time-varying delays have been simulated with different sinusoidal generators (of the type $\tau(t) = \alpha \sin(\beta t + \gamma) + \delta$) and the switching sequence has been chosen randomly (it is indicated with the grey dashed line in the plot). As suggested earlier, the system was initialised with zero power levels. It can be seen that indeed the deviations disappear asymptotically.

Note that if, for instance, the (1,2) element in the matrix C_2 was equal to 0.45 instead of 0.35, then its spectral radius $\rho([C_3^{(1)} \ C_2^{(2)} \ C_3^{(3)}]) \simeq 1.015$ which would violate the stability condition.

These examples conclude our first application that makes use of one of the main results from Chapter 3 in order to derive conditions under which the Foschini-Miljanic algorithm is asymptotically stable, in particular in the presence of time-varying delays and changing network topologies.

6.2 Emissions control in a fleet of Hybrid Vehicles

The second application is inspired by the motivating example from the first chapter. It has been submitted as a contribution to the *Joint 50th IEEE Conference on Decision and Control and the 2011 European Control Conference*, Knorn *et al.* (2011b).

6.2.1 Introduction

Reducing greenhouse gas emissions as well as emissions of directly harmful gases and particulates are one of the major challenges of the future. In the European Union for instance, see Spence *et al.* (2009), attempts to reduce emissions include schemes to encourage optimum driver behaviour (emissions reducing driving style for instance), more efficient use of the transport network (traffic management and smart navigation systems to reduce congestion, dedicated lanes for specific vehicle types, real-time information systems for locations of available parking spaces, *etc.*), or to modify the transport demand (improved logistics to reduce commercial traffic, better public transport, more low-polluting vehicles, *etc.*).

Contributions

In this section, we would like to make a contribution to these efforts by proposing a novel emissions control scheme that makes use of our cooperative control results from Chapter 4. In a fleet of *Plug-in Hybrid Vehicles* (PHEV) we intend to regulate the energy mix used by the cars (that is whether the car should rely more on electric or combustion based propulsion) in order to control the fleet-wide emission of greenhouse gases or harmful particulates.

Structure

In the following, we shall provide some background on the environmental issues that underline the need for better emissions control schemes and mention some of the recent technical developments that should make this possible. We shall then discuss the implementation of our proposed control scheme and finally give some simulation results that validate our vision.

6.2.2 Background

Attempts by large cities like London (Mayor of London, 2008) or Berlin (Schoenburg, 2008) to reduce emissions have received much public attention, particularly due to the direct impact they have on the public's mobility. They try to either discourage drivers to take their car into the city centre by charging a significant fee for doing so, or by strictly only allowing (certified) low-polluting vehicles to enter. While these attempts indeed succeed in somewhat diminishing the number of vehicles in the typically congested city centres,

they basically are open-loop schemes that do not use feedback to respond to the actual situation. Factors like the weather, the time of day, day of the week, or public holidays all have a significant impact on air quality and green house gas emissions. Another problem is that although cars become greener and greener, there are more and more cars in circulation so that the effects of more efficient and less polluting engines is offset by the ever growing number of cars, Mayor of London (2008).

Research and development in the field of electric vehicles has progressed significantly in recent years. *Hybrid electric vehicles* (HEV), which combine a conventional internal combustion engine (ICE) based propulsion system with an electric engine, were introduced to the mass market around the early 2000s, and, apart from their economic advantages in terms of fuel economy and their “green appeal”, a number of additional factors have led to fast growing sales, Gallagher and Muehlegger (2008). Just to name a few, strong tax incentives in most countries make a compelling argument for these low-emission vehicles; social preferences and awareness for environmental quality or energy security have increased; fuel prices can rise and already have risen sharply in the past, with a consistent upward trend over time; most major car manufacturers now offer hybrid cars in their portfolio, broadening the range of available models from small city cars to big SUVs and even vans. Nonetheless, consumer adopting rates could still be improved upon, Lane and Potter (2007).

A new generation of hybrid vehicles are so-called *Plug-in Hybrid Vehicles* (PHEV). These cars have a much larger battery than traditional hybrids and are designed to be charged not only while driving (through regenerative braking for instance), but more importantly by means of “plugging” into an external power supply such as a wall socket when the car is parked. At the current state of the art, this allows the car to drive several tens of kilometres purely on electric power, hence producing zero *local* emissions. The electrical energy, however, still has to be produced somewhere: This can either happen in a “clean” fashion (such as wind, solar, water or nuclear power based) or a “dirty” fashion (traditional fossil fuel based power plants). But while the latter also pollute the air and produce greenhouse gases, the overall emissions and harmful particulates may be filtered more effectively and, since power plants are usually located far away from urbanised zones, their pollution does not accumulate in the cities as is the case with traditional, fossil fuel based transport. Thus, the air quality in densely populated areas — which pose major health concerns (Friends of the Earth Trust, 1999; Gorham, 2001) — will be improved either way.

Unfortunately, market adoption of PHEVs is still somewhat slow, mainly due to economical reasons and technical limitations of the current battery technology. In short, it appears that battery technology still needs to improve in order for this class of vehicle to be economically viable, Axsen *et al.* (2008). Additionally, very few vehicles currently can drive farther than 100km in purely electrical mode, and this figure drastically reduces in cold weather conditions. For that reason, the combustion engine currently serves mainly

as a range extender, allowing the car to run (as commonly expected) several *hundreds* of kilometres — but at the expense of local air pollution.




Figure 6.3: Illustration of a simple parallel drive train configuration in hybrid electric vehicles.

6.2.3 Controlling emissions, maximising driving distance

Hybrid electric vehicles clearly offer many new and exciting possibilities for urban mobility. For the first time, cars can be truly context-aware. In principle, it is possible to combine GPS and engine management unit to enable vehicles to choose where to be most polluting. For example, it makes eminent sense for a hybrid vehicle to switch to full electric mode in the neighbourhood of schools or hospitals. In the following application we explore, at a very high level, a fleet-wide notion of such context awareness. We wish to, in a manner that is fair, adjust the behaviour of the hybrid electric vehicles such that city-wide pollution and/or emissions are regulated. Before proceeding, we give a few words on hybrid electric vehicle fundamentals.

Hybrid vehicles come in several power-train configurations, the most common of which would be the *parallel* power-train configuration, illustrated in Figure 6.3 above. In this set-up, a combustion engine works in conjunction with an electric motor to provide extra torque, or, particularly in the case of plug-in hybrids, to extend the driving range. An interesting variation of this basic design idea is the so-called *power-split hybrid* configuration: It uses power-split devices (such as planetary gear sets combined with additional clutches) to allow a precise control over the different power paths from the engines to the wheel. One essentially attempts to decouple the power supply from the power demand by the driver. The end result is that the two methods of propulsion can either run exclusively or in conjunction (“blended mode”). In other words, it is possible to “mix” the power

sources and vary between emission-free, all-electric mode (but with very limited range) or emission-producing combustion-based mode (allowing for much larger driving ranges).

Let us now propose a scheme to manage this trade-off in order to cooperatively regulate CO₂ emissions³ in a fleet of n vehicles, while maximising their overall driving range for a given level of overall emissions. For that, we shall make the following assumptions:

- (i) The participating PHEVs have a parallel power-train configuration that allows arbitrary blending between the power output of the combustion engine and the electric motor.
- (ii) The drive train power mixing can be described by a convex combination, in other words the car can seamlessly interpolate between the two extremes (all-electric or all-combustion).
- (iii) The vehicles are equipped with some broadcast-based vehicle-to-vehicle communication system (such as the proposed 802.11p protocol for *Co-operative Awareness Messages*, Bilstrup *et al.*, 2008) that allows each car to broadcast its current emission level to other cars in the area. The emissions need not be measured in real-time but could be derived from offline measurements, taking into consideration the currently used power blend.
- (iv) Information about the aggregate CO₂ emissions are available to each car. They could either be measured externally and broadcast to the fleet (through the *Traffic Management Channel* for instance, TMC Forum, 2007), or the cars could collectively estimate them through some distributed averaging scheme such as discussed in Chapter 5.
- (v) The emissions control scheme should be fair in the sense that no car should be allowed to have higher emissions than others.

6.2.4 Implementation

Given these assumptions, this set-up can easily be cast into the framework presented in Chapter 4.

Let us begin by defining the *blending parameter* $r^{(i)} \in [0, 1]$ for each car i that describes the energy mix used for propulsion. By convention, let $r^{(i)} = 0$ if the car is in all-electric mode, and $r^{(i)} = 1$ if the car is propelled purely by the combustion engine. This blending parameter would be the “physical state” in our earlier terminology. With the assumption that the power blending can be described as a convex combination, the utility functions would then be linear functions that map the interval $[0, 1]$ of the blending-parameter $r^{(i)}$ into the corresponding range of emissions $t^{(i)}$ that vary between 0 (when

³ Note that we use CO₂ emissions here only as an example — our scheme can easily be applied to any other type of emissions.

in emissions-free all-electric mode) and $\bar{t}^{(i)}$, the nominal CO₂ emissions of the combustion engine. Specifically,

$$t^{(i)} = f^{(i)}(r^{(i)}) = r^{(i)}\bar{t}^{(i)} \quad (6.16)$$

as illustrated in Figure 6.4 below.




Figure 6.4: Illustration of the emissions as a function of the power blending parameter. On the left, the vehicle is in fully electric mode producing no emissions; on the right it relies completely on its combustion engine and produces the maximum amount of emissions.

The global function in this setting is simply the sum of all the CO₂ emissions, that is

$$g(\mathbf{r}) = \sum_{i=1}^n \bar{t}^{(i)} r^{(i)} \quad (6.17)$$

The overall objective is to maximise driving range for each car in fair way, given a certain “budget” of permissible aggregate emissions. Thus, in order to satisfy the fairness requirement, the emissions between the different cars must be equalised. At the same time, in order to maximise the driving distance, the cars should rely on their combustion engines as much as possible.

These two objectives can easily be achieved using Algorithm 2 on page 62. In order to implement the following slightly simplified version of the control law (4.23)

$$r_{k+1}^{(i)} = r_k^{(i)} + \bar{\eta}_k^{(i)} \sum_{j \in \mathcal{N}_k^{(i)}} (t_k^{(j)} - t_k^{(i)}) + \mu \sigma_k \quad (6.18)$$

we need to calculate suitable gains $\bar{\eta}_k^{(i)}$ and μ . Given the linear / multi-linear nature of the global- and utility functions involved here, this task is straightforward, and we shall briefly demonstrate this process for a small fleet with $n = 4$ cars. Assume the cars have

nominal emissions $\bar{t} = (100 \ 120 \ 140 \ 160)^\top$, measured in g CO₂/kg. Now, recall that the gains $\eta_k^{(ij)}$ had to satisfy (4.25) on page 63:

$$\eta_k^{(ij)} \geq \varepsilon_1 \quad \text{for } j \in \mathcal{N}_k^{(i)}, \quad \text{and} \quad \sum_{j \in \mathcal{N}_k^{(i)}} \eta_k^{(ij)} \leq \frac{1}{\bar{d}^{(i)}} - \varepsilon_2 \quad (6.19)$$

which, in the simplified case, means

$$\bar{\eta}_k^{(i)} \geq \varepsilon_1, \quad \text{and} \quad (n-1)\bar{\eta}_k^{(i)} \leq \frac{1}{\bar{d}^{(i)}} - \varepsilon_2 \quad (6.20)$$

for each $i = 1, 2, 3, 4$. Thanks to the linear / multi-linear nature of the global- and utility functions, we have $\underline{d}^{(i)} = \bar{d}^{(i)} = \underline{h}^{(i)} = \bar{h}^{(i)} = \bar{t}^{(i)}$. Picking $\varepsilon_1 = \varepsilon_2 = 1.5 \cdot 10^{-3}$ and setting the gains as

$$\bar{\eta}_k^{(i)} = \frac{1 - \varepsilon_2 \bar{t}^{(i)}}{|\mathcal{N}^{(i)}| \bar{t}^{(i)}} \quad \text{for } k = 0, 1, \dots \quad (6.21)$$

it is straightforward to check that both inequalities in (6.20) are satisfied for each i .

Next, we need to set small enough gains μ on the global term so that (4.26) is satisfied. Using (4.50) on page 65, this yields in the current setting

$$\mu = 2.1 \cdot 10^{-5} \quad (6.22)$$

With this example on how to actually implement the emissions control scheme, let us now present three simulations of this set-up.

6.2.5 Simulations

The following simulations we generated for fleets of $n = 4$ as well as $n = 50$. The former simulates the numerical example we just discussed; the latter uses a much larger fleet with cars whose emissions are realistically distributed among the different emissions classes based on currently available CO₂ emission statistics, Figure 6.5 on the following page.⁴

In each case, the topology of the communication graph was changed randomly in each time step (but so as to always guarantee strong connectedness). For each simulation run, the agents were initialised to use a 50/50 power mix, that is $r_{k=0}^{(i)} = 0.5$ for each $i = 1, \dots, n$. From then on, the blending-parameter was modified iteratively based on the update law presented earlier. In each case, the desired aggregate emissions level g_* was set to be 25 % lower than that at time $k = 0$, thus requiring all the cars in the network to adjust their energy mix in order to reduce overall emissions by 25 %.

In all the following figures, the first sub-plot shows the evolution over time of the overall emissions $g(\mathbf{r}_k)$ (with the desired level g_* indicated by the dashed line), the next sub-plot displays the corresponding evolution of the blending-parameters $r^{(i)}$, and the last sub-plot gives the evolution of the emissions $t^{(i)}$.

⁴ These statistics give the distribution of emissions produced by “regular” cars among the different EU emission classes. While the combustion engines found in PHEVs should have lower emission levels than regular cars we still used this data for lack of emissions statistics for PHEVs.

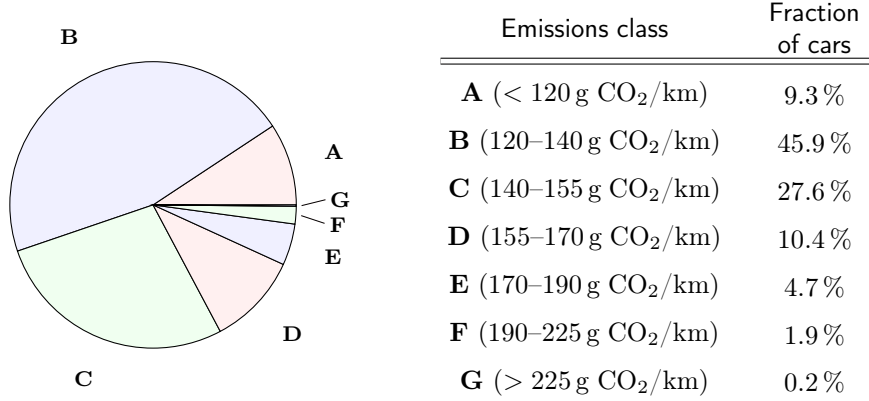


Figure 6.5: Distribution of the fleet's cars among the different EU emission classes, data based on statistics from the Department of Transport (2010).

Fleet of 4 cars

The control gains for the first simulation were set as derived above, and the results from the simulation run are shown in Figure 6.6 on the next page. The plots show that eventually the global emissions in the first subplot converge to the desired level, and that all cars indeed equalise their local emissions (third subplot). The detail view on the right shows the first 15 time steps during which consensus on the emissions is quickly reached. From then on, the agents jointly decrease their blending-parameters as to reduce the overall emissions to the desired level. The overall view on the left, however, shows that convergence ultimately can be considered rather slow, which is due to the conservative nature of our results.

To further illustrate this fact, we also ran a second simulation based on the same network and initial conditions, but this time setting the gain $\bar{\mu}^{(i)}$ 20 times higher than in the previous case. As shown in Figure 6.7, this resulted in an about 10 times faster convergence rate.

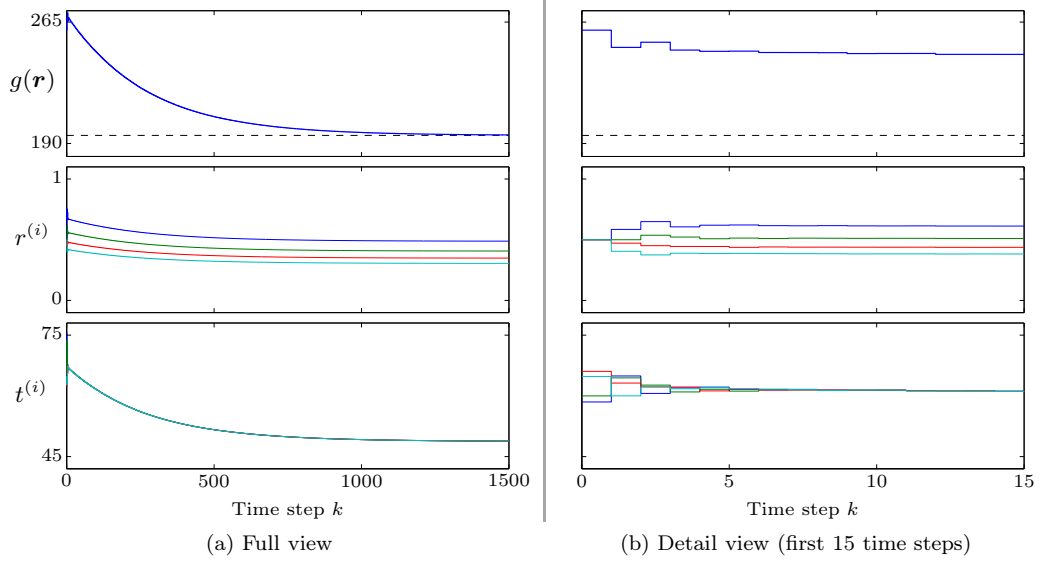


Figure 6.6: Simulation results for the fleet of 4 cars, gains set in accordance with Theorem 4.2. Left: Full view, right: detail view of the first 15 time steps.




Figure 6.7: Same network as in Figure 6.6, but the gains $\bar{\mu}^{(i)}$ were set 20 times larger than in the previous case.

Fleet of 50 cars

Simulating a larger fleet, Figure 6.8 on the following page shows the results for a fleet of $n = 50$ cars. The “jumps” in all the sub-plots at times k that are multiples of $n - 1 = 49$

are due to the inclusion of the global term in the update rule at those instants. For these simulations, again a larger gain $\bar{\mu}^{(i)}$ was used.

Note that in all simulations, as expected, some agents use a larger blending-parameter than others. These would be cars with overall less polluting engines, which means they are allowed to rely more on their combustion engines. This in turn means that these cars should be able to drive farther than others, so that their “eco-friendliness” is rewarded with extended range.




Figure 6.8: Fleet of $n = 50$ cars, gains $\mu_k^{(i)}$ set manually.

Comments In the simulations here the update law from Algorithm 2 was used only in its basic. In a real-world setting, however, one may be required to employ the two extensions for asynchronous state updates and limited access to the global term.

Also, the control scheme as presented in Chapter 4 requires the states (and utility values) to be defined for the *entire* field of real numbers. In the application presented here, however, both variables are only defined on compact intervals. We thus assume that, with the blending-parameters all initialised properly, the solution is feasible and does not drive the parameters beyond their domain of definition. If, however, this was the case, the blending-parameters would simply saturate in either fully electric or combustion mode.

Lastly, the CO₂ emissions of cars are typically strongly dependent on the driving speed as well as the individual driver’s behaviour — both of which is not taken into account here. We rather focus on the *average* emissions that would be produced in typical city traffic.

Furthermore, the frequency at which new aggregate emissions measurements are provided determines the rate at which the discrete updates occur. /

This concludes our first application of Algorithm 2 which aimed at cooperatively regulating CO₂ emissions in a fleet of plug-in hybrid electric vehicles. Before moving on, we would like to stress again that we used CO₂ emissions purely for illustration purposes, any kind of emissions (such as the directly harmful respirable dust produced by combustion engines) or combinations of different emissions may indeed be considered.

* * *

6.3 Real-world implementation of cooperative control

The last application that we would like to discuss is an actual real-world implementation of Algorithm 2 (Theorem 4.2 on page 62).

6.3.1 Introduction

All our earlier results were developed with real word implementations in mind, so in order to see whether indeed our *theory* can be put into *practice*, a test and validation program was jointly developed with Dr. Ronan Farrell and Mr. James Kinsella. Both are with the *Callan Institute* here at the *National University of Ireland Maynooth*, which has great expertise and resources in electronic, hardware and software systems as well as wireless communications.

Over the course of one year, thanks to the kind help of Mr. Kinsella a total of six wireless motes was developed, built and programmed in order to set up a small wireless network of autonomous agents in which to test our results. Five of the six motes that were built are shown in Figure 6.9 on the following page.

Contributions and structure

This section briefly describes a validation experiment of our theoretical contributions in Chapter 4 by developing an actual hardware/software implementation of Algorithm 2 in




Figure 6.9: The five “regular” motes.

the presence of real-world limitations and problems (in particular, communication failures and limited hardware capabilities).

In the following, we shall first describe the overall set-up as well as the hardware and software layout of the wireless units. We then present and analyse the results from two indicative experiment runs.

6.3.2 Overall setting

Among the six motes built, five were *regular nodes* that formed the actual multi-agent network. Those nodes were completely autonomous units, in that they were battery powered and only communicated wirelessly. The sixth mote acted as a *master node*: It “measured” and fed back the global term to the network. Additionally, it was used to start off the experiment and also collect debug information from each of the five regular nodes. This information consisted of data packets containing the node’s id, physical state and utility value, and this data was directly relayed to a PC so that the system’s state could be recorded and displayed in real-time.

The experiments themselves consisted of each node initialising itself with a fictitious physical state and utility function (as in the simulations in Chapter 4, those functions were either of linear or quadratic type, cf. Section 4.A.2 on page 81). Then, the consensus protocol given by Algorithm 2 was run: The nodes broadcast their utility value using their radios and receive similar broadcasts from neighbouring nodes as well as the master the global term from node in order to update their own state. The global function used in the master node was the mean function, that is the master node calculated and fed back the mean of the physical states of all the nodes. The desired global value was changed several

times over the course of each experiment in order to demonstrate the network's ability to react to and track such changes.

The controller gains were calculated in a similar procedure as presented in the previous section, but again manually increased by one order of magnitude in order to reduce the overall run time of the experiments.

Hardware

All the nodes were built on small printed circuit boards (PCB) with the following basic components:

- (i) *CPU: Microchip PIC 18LF4550 (8bit, 32KB Flash, 2K RAM, USB)*
- (ii) *Radio: Texas Instruments CC1100 (ultra low power, sub-1 GHz transceiver)*
- (iii) *Power: 3x AA batteries, regulated to 3.3V*
- (iv) *Interface: 2 status LEDs, 1 reset button, serial connector*

The master node additionally had a physical USB port wired to the CPU so that it could be connected to a PC for real-time monitoring of the network's behaviour.


Software

The software for the motes was written in C/C++ and consists of a simple firmware to initialise and control the hardware components as well as an algorithmic block which contains the actual implementation of the consensus algorithm.

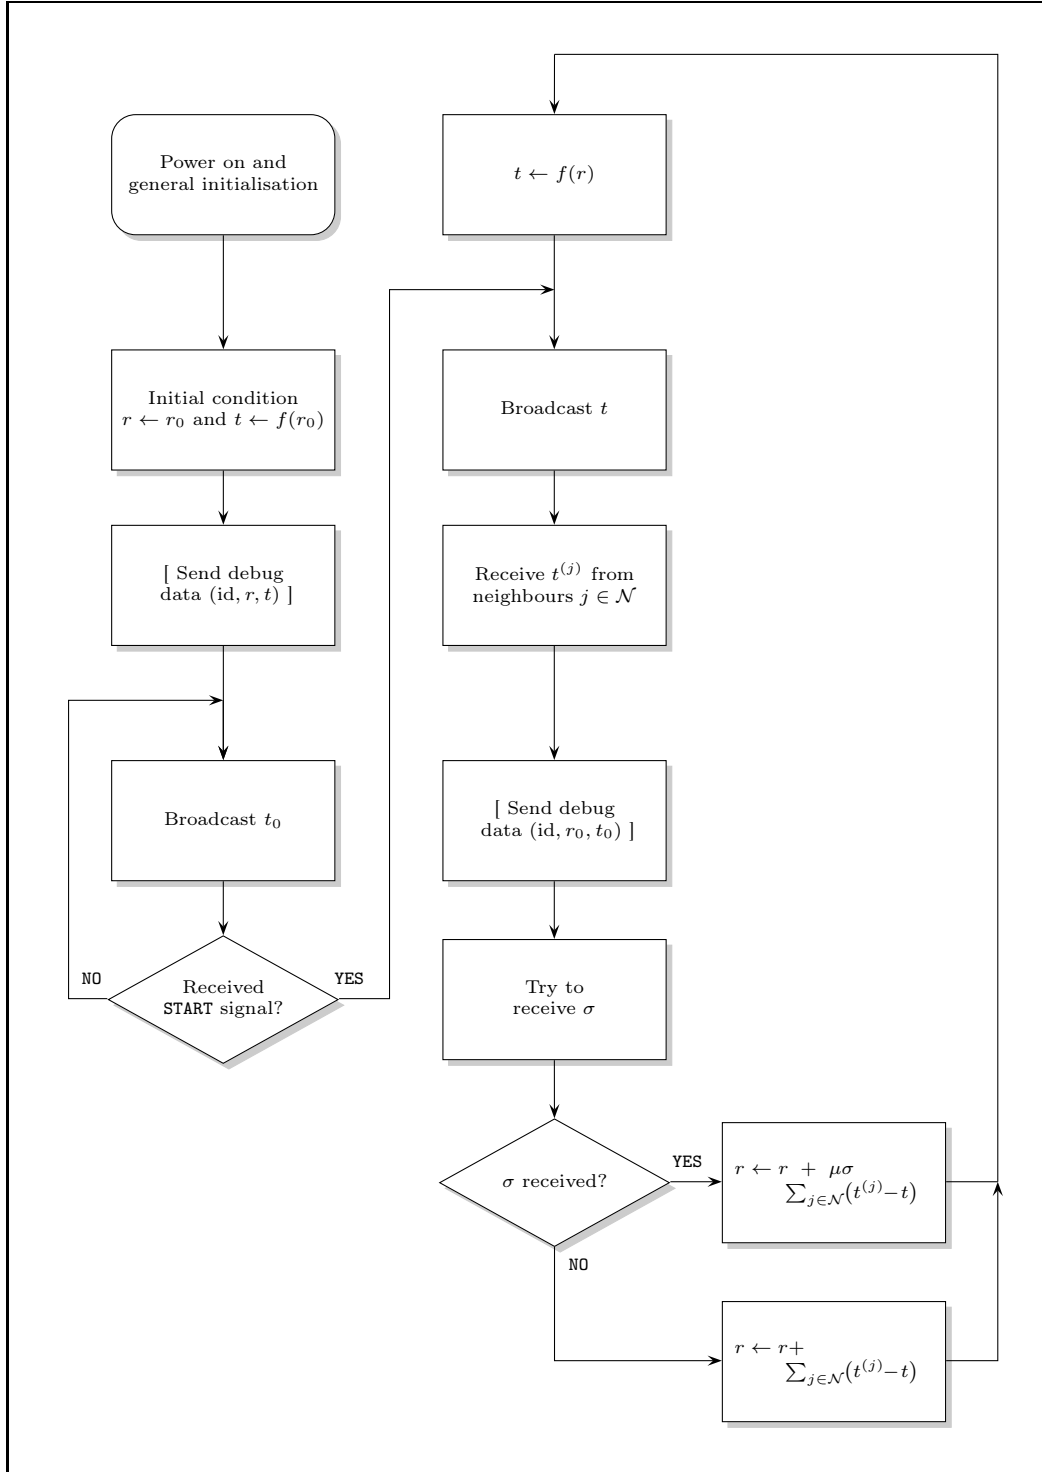
Roughly, the software set-up operates as follows (please also refer to the flow charts in the next two sections). First the master node had to be powered on, then the other nodes. When a node has finished booting up and is ready for the experiment the begin it was set to continuously broadcast its initial state to signal its readiness.

The experiment would start when the master node had received the initial states from all the nodes in the network. At that point, the master node would broadcast a trigger signal to start the consensus algorithm in each node. In each iteration, the nodes were programmed to exchange their utility states with each other and update their state accordingly. However, every $n - 1 = 4$ iterations they were additionally instructed to send out debugging information directly to the master node (containing their id, physical state and utility value). This information was required for two purposes: (i) to enable the master node to calculate the global value and subsequently feed the difference between desired and actual global value back to the nodes, (ii) to protocol the evolution of the experiment.

Software layout of master node



Software layout of regular node



6.3.3 Practical issues

For reasons of simplicity the wireless communications between the nodes were not realised using one of the established wireless communication protocols. Rather, they were implemented in a straightforward *round-robin* or *time division multiple access* (TDMA) fashion. Simply speaking, this means that the nodes take turns broadcasting. To coordinate this, first the master node would broadcast the global value. Relative to this broadcast the different nodes would broadcast with different, fixed delays so that the transmissions are “staggered out” and collisions are avoided. That way the $n - 1$ iterations between the global term updates were performed. At the end of these iterations the nodes would spend another round broadcasting their debug data packets to the master node. Having received this data, the master node would then broadcast the new global term and the cycle starts anew.

In terms of timing, each node’s broadcast window was about 200 ms wide so that the $n - 1$ iterations usually took around 4 seconds. Then, roughly another second was spent transmitting the debug data packets. Finally, the master node had 200 ms to broadcast the updated global term. Hence, all in all, the network would perform about 11 to 12 global term updates per minute, provided no data packets were lost (which cause some global term updates not to be performed).




Figure 6.10: The static communication networks used in the two experiments (master node not shown).

Due to technical reasons and limited space for the experiment, every node would pick up every other node’s broadcast. This would result in a rather uninteresting, complete graph. We thus manually added an exclusion list to each node which instructed it to

ignore broadcasts from certain nodes. With this method, we created two different network topologies,⁵ which are shown in Figure 6.10 on the facing page (master node not included).

Furthermore, due to the unreliable nature of the nodes and wireless communications, many data packets were dropped. This shall become evident in the somewhat “non-smooth” evolution of the states (as compared to the computer simulations from Chapter 4). However, this effect will always be encountered in real-world applications and thus allows us to demonstrate the robustness of our work to such communication problems.

Lastly, the microprocessors only used finite precision arithmetics and the states could only assume integer values. While we did not explicitly took this into account in the present work, it shows that our algorithm is also robust with regards to such perturbations.

6.3.4 Results from experiments

We shall now discuss the results from the two experiment runs. In the first experiment, the network was using utility functions of quadratic type. Initially, the target value for the global value was set to 750, which means that the network’s physical states had to be adjusted so that their mean would equal 750. The master node was instructed to automatically switch the desired global value to 400 once the networks states have converged (that is, when the precision of the arithmetic-logic-unit was reached). This occurred about 11 minutes into the experiment.

Figure 6.11 on the next page was generated using the data recorded during the first experiment, that is the debug-data received from the master node and the global value that it had calculated and broadcast to the nodes. The first sub-plot shows the difference between actual and desired global value, which starts off positive since the average value of the physical states (shown in the second sub-plot) is clearly below 750. As the physical states increase, the difference starts to disappear. At the same time, the utility values in the thirds subplot approach each other quickly and eventually converge to perfect consensus. Then, after about 11 minutes, the target global value was changed to a different, lower value, which meant that the physical states had to generally decrease in order to meet it — which can indeed be observed in the plots.

A similar picture is painted in Figure 6.12. This time, the utility functions were of linear type, and the global value changed three times throughout the experiment (from 200 to 600 to 400). Again, the network behaves as desired.

Comment As in the emissions control application earlier, gains were set higher than required by the theorem. This was particularly important in this application as the system was somewhat unstable, with nodes more or less randomly crashing. The gain on the global term was double from the first to the second experiment, which clearly resulted in faster convergence times. /

⁵ In order to limit the complexity of the code we did not change the network topology over the course of an experiment.




Figure 6.11: Evolution of the deviation from the desired global value (which changed at $t=10:50$ min), the physical states and the utility values of the five nodes in the network. Utility functions: quadratic-type, global function: mean of physical states.




Figure 6.12: Linear-type utility functions, desired global value changed at $t=1:40$ min and $t=6:40$ min.

This concludes our real-world validation of Algorithm 2. The results from the experiments certainly support our claims of the robustness of our proposed cooperative control scheme. We shall now close this application chapter and move on to draw some final conclusions of the work presented in this thesis.

Conclusion

In this last chapter we summarise the contributions of this thesis and suggest possible future directions for continued research in the relevant areas.

Chapter contents

- 7.1 Summary
 - 7.2 Future directions
-

7.1 Summary

In the first chapter we gave several examples to motivate the work carried out in this thesis. The first one concerned a transmit power control algorithm for wireless networks. The famous Foschini-Miljanic algorithm is a distributed control scheme that is known to be robust to various types of perturbations, in particular time-varying time-delays of the states. We noted that switched positive systems theory can be used to explain this robustness and give conditions under which stability can be guaranteed. The second example suggested a municipal emissions control scheme for a fleet of cars. The idea was to regulate each participating car's driving speed in order to control on a global level the overall fleet emissions while also equalising the local emissions among cars (fairness). The final motivational example concerned a topology control problem in wireless sensor networks. The objective was to find a decentralised algorithm which regulates the broadcast power in each node so that a certain overall connectivity level was maintained in the network while, at the same time, balancing battery lifetimes among all the nodes in order to maximise the network's lifetime without node failures (due to insufficient battery power). All three examples or variations thereof were later revisited in Chapters 5 and 6.

Chapter 2 then reviewed literature from related fields of research, in particular the areas of switched positive systems, large-scale systems, decentralised control, and cooperation in networked multi-agent systems.

The third chapter was concerned with deriving easily verifiable stability conditions for switched positive linear systems, in particular by giving conditions for the existence of

common linear co-positive Lyapunov functions. We noted that these switched systems may represent a networked of interacting scalar systems which switches between different interaction topologies. As we noted in the literature review, existence of any type of *common* Lyapunov functions, in general, is sufficient for exponential stability of a switched linear system. In that context, the first result that we presented dealt with a switched positive linear system where the switching could not occur arbitrarily, but depended directly on the states: Given a state space covered with (possibly overlapping) convex cones each of which was associated with one of the constituent subsystems, the system was only allowed to switch to whichever subsystem(s) that was (were) associated with the cone(s) the system was in at that point. Our result then stated that existence of a common linear co-positive Lyapunov function is equivalent to the cone generated by all the columns of all the constituent system matrices not intersecting the positive orthant. As this condition is somewhat hard to test in general, we presented a reformulation of this result that permitted easy checking by running a simple feasibility test provided the cones encountered were polyhedral. Attention was then turned to the arbitrary switching case, for which a necessary and sufficient existence condition was found that consisted of an extended Hurwitz condition on all the system matrices. These results were complemented by remarks concerning the insights gained from the algebraic condition, their applicability to discrete time systems, and a number of possible applications for them.

The following chapter, Chapter 4, discussed a novel cooperative control paradigm for networked systems. To achieve this, a global feedback loop was added to the network, relating back the aggregate network behaviour to each agent. To formalise the discussion, we began by defining more concretely the oft-encountered system setting that we had already briefly presented in the first chapter. We then derived and proved convergence of three basic algorithms that allowed the network to cooperatively achieve a common goal given certain local and global constraints. In terms of the existing literature on consensus and coordination, our results can be interpreted as enabling an implicit and constrained consensus to be found in a fully decentralised setting, running on directed and time-changing communication network topologies. Of the three algorithms, the first one assumed perfect knowledge of all the system parameters, in particular (the inverse of) the utility functions; the second algorithm was much less demanding in that only bounds on the slopes of global- and utility functions were needed; the third algorithm was even more general, even allowing for dynamics to occur in the utility functions. Extensions to these algorithms additionally enabled them to be deployed in even more demanding settings, such as situations where synchronous state updates cannot be guaranteed, and where the global value may not be accessible to all the nodes. However, for all these results one key assumption was made: the global term had to be known by at least one node. This could be satisfied either by some external entity providing it to the nodes, or by the network measuring or estimating it itself.

One such situation where the global value can indeed be estimated in a decentralised fashion was discussed in the subsequent Chapter 5. Hence, in comparison to the previous chapters, an identification component was added to the problem. The main contributions of the chapter solved the problem posed in the second motivational example: A wireless networks where the level of connectivity of the communication network needs to be regulated, as there are several algorithms that evolve over such networks whose convergence rate directly depends on the algebraic connectivity level. As proxy for the connectivity level we used the second largest eigenvalue in magnitude of the stochastic normalisation of the adjacency matrix. This value is closely related to the traditionally used Fiedler eigenvalue of the graph Laplacian, but it has the advantage that it is also defined for directed graphs. Additionally, a fully decentralised scheme can be devised that allows this value to be estimated locally in each node — one of the main contributions of the chapter. Once obtained, this estimate was then used to inform a decentralised control scheme that locally adjusted the network topology to successfully regulate the overall connectivity to some desired level, even if the network can only assume a discrete number of different topologies and hence connectivity levels.

The sixth chapter then gave three further applications for our main results, in part following up on some of the motivational examples. The first application discussed concerned the Foschini-Miljanic power control algorithm for wireless networks. Our arbitrary switching result from Chapter 3 was used to provide sufficient conditions for the algorithm's stability under time-varying time-delays and arbitrarily changing network topologies. Next, we suggested an emissions control scheme for a fleet of plug-in hybrid electrical vehicles that was based on our second cooperative control algorithm proposed in the fourth chapter. This application is similar to the third motivational example as it proposes a decentralised scheme to regulate (in a fair way) the total emissions of the cars participating in the scheme. Lastly, the third application we considered was a real-world implementation of the same cooperative control scheme, validating our claims that the algorithm can indeed be implemented, even with all the imperfections and limitations that are inherent in real-world applications.

7.2 Future directions

In closing, let us suggest a few different directions that the present work may be extended in.

Switched systems in general are very hard to analyse, as reflected by the fact that the vast majority of results in this area only concern linear systems. It is thus not surprising that there are still many fundamental questions that remain unanswered in the non-linear case. Similarly, our contributions from Chapter 3 also hold only for *linear* positive systems. Unfortunately, the linearity assumption in relation to positive systems in particular is

somewhat problematic: Most real world systems are non-linear and the classical approach of linearising these system would inevitably destroy any positivity properties (as linearisation yields states that describe the *deviation* from the operating point and these error coordinates may thus assume negative values). Hence, there is a clear need for non-linear results in the field of positive systems and positive switched systems in particular.

Nonetheless, in some cases the linearity assumption may be justified; one example for this was encountered in the Foschini-Miljanic application we discussed earlier. In this application, the need for results covering switched positive systems with time-delays became apparent. Consequently, further work further investigating the effects of delays on the system's stability would certainly be of benefit.

While the results of Chapter 4 are certainly promising, a number of open questions remain and should be the subject of further investigations. For instance, the gain μ in the second and third algorithm may become very small in larger networks, and there is much experimental evidence that the bounds presented here tend to be rather conservative. This can be explained, in part, by the fact that for sufficiently connected graphs (and not pathological worst-case scenarios such as, for example, directed n -cycles) significantly less than $n - 1$ multiplications in (4.36) on page 64 would be required to produce strictly positive \bar{S} matrices — which in turn means that the corresponding s_{\min} value in (4.46) on page 65 would be much larger and ultimately allows μ to be increased. One possible future extension of our work that accounts for unlikely topological effects is *via* a stochastic formulation of this problem. Here, expected quantities are controlled rather than deterministic ones.

Also, in the second and third algorithm, the nodes incorporate the global term in their state updates every M steps. A number of simulations and tests have shown that the system may well include the global term in every time step and thus achieve faster convergence. In the future, it would be interesting to see if a proof can be developed that allows the inclusion of the global term in *every* time step, as this may speed up overall convergence.

Other open questions in a problem setting as encountered here concern the effect of communication delays, quantisation effects when using finite precision arithmetics (for instance when implementing our algorithms on digital processors without floating point precision, as was the case in Section 6.3) or the effect of nodes joining or leaving the network. These issues may possibly be addressed using ideas from previous (unconstrained consensus) literature such as Kashyap *et al.* (2007); Nedić and Ozdaglar (2010).

Furthermore, the present work only considers a single physical state and single utility value associated with each node; in a more general setting, nodes could have two or more states associated with them. This may lead to a MIMO-like formulation of our problem.

Concerning the graph topology control problem, although examples are presented to illustrate the efficacy and promises of this approach, there are also a number of open questions that remain to be resolved. The most important of these concerns the fact that

the relationship between the network states and the eigenvalue locus is not known exactly *a priori* (and thus the required bounds κ for Theorems 5.2 and 5.3 on page 105). However, this should not be a problem for most practical applications where the graph setup is roughly known in advance, since then estimates of those parameters could be obtained off-line using simulations of typical graphs. Another approach would be to attempt to estimate them in an adaptive fashion as the consensus algorithm evolves.

Furthermore, another interesting problem to study would be to attempt to reproduce our results without the assumption of separation of time scales between the estimation- and control parts in the overall scheme (Section 5.4 on page 97). Indeed, some preliminary experiments have shown that estimation- and control iterations may be interlaced (that is, individual estimation and control updates may simply be alternated), without affecting the system's stability or convergence to the correct solution.

Another extension to the chapter's work may be to not consider "binary" adjacency matrices (where the entries are either 1 or 0) but rather matrices where those elements would transition smoothly from 1 to 0 as nodes get further apart from each other and lose their communication link. In that case, we suspect that the eigenvalue locus will become a smooth function of the nodes' connection radii.

Lastly, as in the previous chapter, further investigations may also consider communication delays, quantisation effects or the impact of using only finite precision arithmetics on the control scheme, as well as the quantitative effect on the eigenvalue locus when agents join or leave the network.

Moving on to the applications chapter, future directions for the Foschini-Miljanic application may include finding additional stability conditions for constantly changing network topologies where it is not possible to identify finitely many different configurations. Further, a comparison of our results with the stability conditions of the undelayed Foschini-Miljanic algorithm could lead to a better understanding of the impact of delays on the overall algorithm. Lastly, on a more abstract level, it would be of great interest to determine whether the delay-independent stability properties exhibited by positive systems are due to their monotonicity or positive properties.

In the work on fleet-wide emissions control, future studies should consider the effect of nodes joining and leaving the network, how effects like saturation of the states could be incorporated directly into the mathematical framework, and ideally derive tighter bounds on the maximum permissible gain for the global term (as the bounds presented here are only sufficient for stability, and we have shown in the simulations that they can be increased significantly without compromising stability). Also, it would be interesting to attempt a real-life implementation of our suggested application.

Our last application using the purpose built wireless motes also raised a range of questions. For instance, what is the effect of quantisation in the states on the cooperative control algorithms, as caused by real-world, finite precision arithmetics? We suspect that

quantisation will not be able to destabilise the system; similar to the quantisation effects encountered in the graph connectivity problem, one cannot expect to converge asymptotically to the theoretical solution, but rather only to a neighbourhood of it.

In overall conclusion, the cooperative control algorithms should present a new paradigm for cooperative control. However, with the gains set strictly according to the rather conservative, theoretical limits, convergence rates are much too slow for actual applications. Hence, further work deriving tighter bounds for these gains is imperative.

Notation

n	Scalars; lowercase letters
\mathbf{x}	Vectors; lowercase bold letters (always supposed to be column vectors, if not transposed), or elementwise in parentheses: $(x^{(1)} \ x^{(2)} \ \dots \ x^{(n)})^\top$
\mathbf{A}	Matrices; uppercase bold letters, or elementwise in brackets: $\begin{bmatrix} a^{(11)} & a^{(12)} \\ a^{(21)} & a^{(22)} \end{bmatrix}$
\mathcal{S}	Sets; uppercase calligraphic letters
\bar{x}	Typically an upper bound to the variable x
\underline{x}	Typically a lower bound to the variable x
x_*	Typically the “desired” value of x
x_k	The value of variable x at time k , sometimes also denoted $x(k)$
\mathbf{I}	Identity matrix of suitable dimensions
\mathbf{e}_n	The n th unit vector of suitable dimension
$\mathbf{0}$	Zero matrix of suitable dimensions
\mathbb{R}	The field of real numbers
\mathbb{R}^n	The n -dimensional Euclidean space
$\mathbb{R}_{\geq 0}^n$	The closed positive orthant of the \mathbb{R}^n
$\mathbb{R}_{> 0}^n$	The open positive orthant of the \mathbb{R}^n
$\mathbb{R}^{n \times n}$	The space of $n \times n$ matrices with real entries

Bibliography

To facilitate access to all the referenced material, I tried to either provide *Digital Object Identifiers* (DOI) or give short web links for as many entries as possible.

Ian F. Akyıldız, Weilian Su, Yogesh Sankarasubramaniam, and Erdal Çayırıcı. Wireless sensor networks: A survey. *Computer Networks*, 38(4):393–422, Mar. 2002.

DOI: 10.1016/S1389-1286(01)00302-4

Hugo Alonso and Paula Rocha. A general stability test for switched positive systems based on a multidimensional system analysis. *IEEE Transactions on Automatic Control*, 55(11):2660–2664, Nov. 2010.

DOI: 10.1109/TAC.2010.2070730

Masahiko Amano, Aleksandar I. Zečević, and Dragoslav D. Šiljak. An improved block-parallel newton method via epsilon decompositions for load-flow calculations. *IEEE Transactions on Power Systems*, 11(3):1519–1527, Aug. 1996.

DOI: 10.1109/59.535693

Brian D. O. Anderson. Time delays in large-scale systems. In *CDC'79: Proceedings of the 18th IEEE Conference on Decision and Control, including the Symposium on Adaptive Processes*, volume 18, pages 655–660, Dec. 1979.

DOI: 10.1109/CDC.1979.270267

Mark R. Anderson and Andrew C. Robbins. Formation flight as a cooperative game. In *Proceedings of the AIAA Guidance, Navigation, and Control Conferences*, volume 1, pages 244–251, Boston, MA, USA, Aug. 1998.

<http://goo.gl/pxKUUb>

Anthanasios C. Antoulas and Dan C. Sorensen. Approximation of large-scale dynamical systems: An overview. *International Journal of Applied Mathematics and Computer Science*, 11(5):1093–1121, Dec. 2001.

<http://goo.gl/hz0VB>

Anthanasios C. Antoulas, Dan C. Sorensen, and Serkan Gugercin. A survey of model reduction methods for large-scale systems. In *Structured Matrices in Mathematics, Computer Science, and Engineering I*, volume 280 of *Contemporary Mathematics*, pages 193–219. American Mathematical Society, Dec. 1999.

<http://goo.gl/v5L9g>

Athanasios C. Antoulas. *Approximation of Large-Scale Dynamical Systems*, volume 6 of *Advances in design and control*. Cambridge University Press, New York, NY, USA, 2005.

<http://books.google.ie/books?id=IYfRIJIA1mkC>

- Mituhiko Araki. Application of M -matrices to the stability problems of composite dynamical systems. *Journal of Mathematical Analysis and Applications*, 52(2):309–321, 1975.
DOI: 10.1016/0022-247X(75)90099-2
- Mituhiko Araki. Input-output stability of composite feedback systems. *IEEE Transactions on Automatic Control*, 21(2):254–259, Apr. 1976.
DOI: 10.1109/TAC.1976.1101187
- Mituhiko Araki. Stability of large-scale nonlinear systems: Quadratic-order theory of composite-system method using M -matrices. *IEEE Transactions on Automatic Control*, 23(2):129–142, Apr. 1978.
DOI: 10.1109/TAC.1978.1101728
- Mituhiko Araki and Bunji Kondo. Stability and transient behavior of composite nonlinear systems. *IEEE Transactions on Automatic Control*, 17(4):537–541, Aug. 1972.
DOI: 10.1109/TAC.1972.1100042
- Mituhiko Araki, Kazuaki Ando, and Bunji Kondo. Stability of sampled-data composite systems with many nonlinearities. *IEEE Transactions on Automatic Control*, 16(1):22–27, Feb. 1971.
DOI: 10.1109/TAC.1971.1099623
- Murat Arcak and Eduardo D. Sontag. Diagonal stability of a class of cyclic systems and its connection with the secant criterion. *Automatica*, 42(9):1531–1537, 2006.
DOI: 10.1016/j.automatica.2006.04.009
- Karl J. Åström and Björn Wittenmark. *Computer-controlled systems: Theory and Design*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 3rd edition, 1997.
<http://books.google.ie/books?id=jUTIAAAACAAJ>
- Mark A. Aizerman and Feliks R. Gantmakher. *Die absolute Stabilität von Regelsystemen*. Beihefte zur Zeitschrift “Regelungstechnik”. R. Oldenbourg, München, Germany, 1965.
<http://books.google.ie/books?id=1-7mAAAAMAAJ>
- Jonn Axsen, Andrew F. Burke, and Kenneth S. Kurani. Batteries for plug-in hybrid electric vehicles (PHEVs): Goals and the state of technology circa 2008. Technical Report UCD-ITS-RR-08-14, Institute of Transportation Studies, University of California, Davis, CA, USA, May 2008.
<http://goo.gl/WoqFD>
- Fredric N. Bailey. The application of Lyapunov’s second method to interconnected systems. *Journal of the Society for Industrial and Applied Mathematics, Series A: Control*, 3(3):443–462, 1965.
DOI: 10.1137/0303030
- Lubomír Bakule. On the overlapping decomposition of dynamic systems. *Annual Review in Automatic Programming*, 12(1):263–266, 1985.
DOI: 10.1016/0066-4138(85)90039-4
- Lubomír Bakule. Decentralized control: An overview. *Annual Reviews in Control*, 32(1):87–98, Apr. 2008.
DOI: 10.1016/j.arcontrol.2008.03.004
- Tucker Balch and Ronald C. Arkin. Behavior-based formation control for multirobot teams. *IEEE Transactions on Robotics and Automation*, 14(6):926–939, Dec. 1998.
DOI: 10.1109/70.736776
- Nikita E. Barabanov. Lyapunov indicators of discrete inclusions I–III. *Automation and Remote Control*, 49(2):152–157, 1988.

- Nikita E. Barabanov. Method for the computation of the Lyapunov exponent of a differential inclusion. *Automation and Remote Control*, 50:475–479, 1989.
- David J. Bartholomew, Andrew F. Forbes, and Sally I. McClean. *Statistical techniques for manpower planning*. Wiley series in probability and mathematical statistics: Applied probability and statistics. Wiley, 2nd edition, 1991. http://books.google.ie/books?id=e_-qAAAAIAAJ
- Randal W. Beard and Vahram Stepanyan. Information consensus in distributed multiple vehicle coordinated control. In *CDC'03: Proceedings of the 42nd IEEE Conference on Decision and Control*, volume 2, pages 2029–2034, Maui, HI, USA, Dec. 2003. DOI: 10.1109/CDC.2003.1272913
- Randal W. Beard, Jonathan R. T. Lawton, and Fred Y. Hadaegh. A feedback architecture for formation control. In *ACC'00: Proceedings of the 2000 American Control Conference*, volume 6, pages 4087–4091, Chicago, IL, USA, Jun. 2000. DOI: 10.1109/ACC.2000.876990
- Randal W. Beard, Al W. Beard, Jonathan Lawton, and Fred Y. Hadaegh. A coordination architecture for spacecraft formation control. *IEEE Transactions on Control Systems Technology*, 9(6):777–790, Nov. 2001. DOI: 10.1109/87.960341
- Richard E. Bellman. Vector Lyapunov functions. *SIAM Journal of Control*, 1(1):32–34, Jan. 1962. DOI: 10.1137/0301003
- Jon A. Benediktsson and Philip H. Swain. Consensus theoretic classification methods. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(4):688–704, Jul./Aug. 1992. DOI: 10.1109/21.156582
- Luca Benvenuti and Lorenzo Farina. On the class of linear filters attainable with charge routing networks. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 43(8):618–622, Aug. 1996. DOI: 10.1109/82.532010
- Abraham Berman and Robert J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. Computer science and applied mathematics. Academic Press, New York, NY, USA, 1979. http://books.google.ie/books?id=MRB7SUc_u6Yc
- Abraham Berman, Michael Neumann, and Ronald J. Stern. *Nonnegative Matrices in Dynamic Systems*. Pure and applied mathematics. John Wiley & Sons, Inc., New York, NY, USA, 1989. <http://books.google.ie/books?id=VQrvAAAAMAAJ>
- Dimitri P. Bertsekas, Angelia Nedić, and Asuman E. Ozdaglar. *Convex Analysis and Optimization*. Athena Scientific, Belmont, MA, USA, Apr. 2003. <http://goo.gl/YTbqf>
- Amit Bhaya and Francisco das Mota Chagas. Equivalence of stability concepts for discrete time-varying systems. *International Journal of Robust and Nonlinear Control*, 4(6):725–740, 1994. DOI: 10.1002/rnc.4590040603
- Katrin Bilstrup, Elisabeth Uhlemann, Erik G. Ström, and Urban Bilstrup. Evaluation of the IEEE 802.11p MAC method for vehicle-to-vehicle communication. In *VTC'08-Fall: The 68th IEEE Vehicular Technology Conference*, pages 1–5, Calgary, Canada, Sept. 2008. DOI: 10.1109/VETECF.2008.446
- Franco Blanchini. Nonquadratic Lyapunov functions for robust control. *Automatica*, 31(3):451–461, March 1995. DOI: 10.1016/0005-1098(94)00133-4

- Vincent D. Blondel, Julien M. Hendrickx, Alex Olshevsky, and John N. Tsitsiklis. Convergence in multiagent coordination, consensus, and flocking. In *CDC-ECC '05: Proceedings of the Joint 44th IEEE Conference on Decision and Control and the 2005 European Control Conference*, pages 2996–3000, Seville, Spain, Dec. 2005.
- Vivek Borkar and Pravin P. Varaiya. Asymptotic agreement in distributed estimation. *IEEE Transactions on Automatic Control*, 27(3):650–655, Jun. 1982. DOI: 10.1109/TAC.1982.1102982
- El-Kébir Boukas. *Stochastic switching systems: Analysis and design*. Control engineering. Birkhäuser, Boston, MA, USA, 2006. <http://books.google.ie/books?id=9CuKN5rCrRiC>
- Stephen Boyd, Persi Diaconis, and Lin Xiao. Fastest mixing markov chain on a graph. *SIAM Review*, 46(4):667–689, Apr. 2004. DOI: 10.1137/S0036144503423264
- Stephen P. Boyd, Laurent El Ghaoui, Eric Feron, and Venkataramanan Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. SIAM studies in applied mathematics. SIAM Press, 1994. <http://books.google.ie/books?id=AO6mvmmlKUkC>
- Michael S. Branicky. Stability of switched and hybrid systems. In *CDC'94: Proceedings of the 33th IEEE Conference on Decision and Control*, volume 4, pages 3498–3503, Lake Buena Vista, FL, USA, Dec. 1994. DOI: 10.1109/CDC.1994.411688
- Michael S. Branicky. Multiple Lyapunov functions and other analysis tools for switched and hybrid systems. *IEEE Transactions on Automatic Control*, 43(4):475–482, Apr. 1998. DOI: 10.1109/9.664150
- Robert K. Brayton and Christopher H. Tong. Stability of dynamical systems: A constructive approach. *IEEE Transactions on Circuits and Systems*, 26(4):224–234, Apr. 1979. DOI: 10.1109/TCS.1979.1084637
- Robert K. Brayton and Christopher H. Tong. Constructive stability and asymptotic stability of dynamical systems. *IEEE Transactions on Circuits and Systems*, 27(11):1121–1130, Nov. 1980. DOI: 10.1109/TCS.1980.1084749
- Richard A. Brualdi and Bolian Liu. Fully indecomposable exponents of primitive matrices. *Proceedings of the American Mathematical Society*, 112(4):1193–1201, Aug. 1991. DOI: 10.1090/S0002-9939-1991-1065941-8
- Stefan Bundfuss and Mirjam Dür. Copositive Lyapunov functions for switched systems over cones. *Systems & Control Letters*, 58(5):342–345, 2009. DOI: 10.1016/j.sysconle.2008.12.006
- João B. D. Cabrera, Ram Ramanatha, Carlos Gutiérrez, and Raman K. Mehra. Stable topology for mobile ad-hoc networks. *IEEE Communications Letters*, 11(7):574–576, Jul. 2007. DOI: 10.1109/LCOMM.2007.070256
- Louis Caccetta and Ventsi G. Rumchev. A survey of reachability and controllability for positive linear systems. *Annals of Operations Research*, 98(1):101–122, 2000. DOI: 10.1023/A:1019244121533
- Frank M. Callier, Wan Chan, and Charles A. Desoer. Input-output stability theory of interconnected systems using decomposition techniques. *IEEE Transactions on Circuits and Systems*, 23(12):714–729, Dec. 1976. DOI: 10.1109/TCS.1976.1084167

- Frank M. Callier, Wan Chan, and Charles A. Desoer. Input-output stability of interconnected systems using decompositions: An improved formulation. *IEEE Transactions on Automatic Control*, 23(2):150–163, Apr. 1978. DOI: 10.1109/TAC.1978.1101725
- Ming Cao, A. Stephen Morse, and Brian D. O. Anderson. Agreeing asynchronously: Announcement of results. In *CDC'06: Proceedings of the 45th IEEE Conference on Decision and Control*, pages 4301–4306, San Diego, CA, USA, Dec. 2006. DOI: 10.1109/CDC.2006.376812
- Ruggero Carli, Alessandro Chiuso, Luca Schenato, and Sandro Zampieri. A PI consensus controller for networked clocks synchronization. In *IFAC-WC'08: Proceedings of the 17th IFAC World Congress*, volume 17-1, Seoul, Korea, Jul. 2008. DOI: 10.3182/20080706-5-KR-1001.01741
- M. Kanat Çamlıbel and Hans Schumacher. Copositive Lyapunov functions. In Vincent D. Blondel and Alexandre Megretski, editors, *Unsolved Problems in Mathematical Systems and Control Theory*, chapter 6.1, pages 189–193. Princeton University Press, Princeton, NJ, USA, 2004. <http://goo.gl/4Lb99>
- S. Chae and Z. Bien. Techniques for decentralized control for interconnected systems. In Cornelius T. Leondes, editor, *Analysis and control system techniques for electric power systems*, volume 41 of *Control and dynamic systems: Advances in theory and applications*, pages 273–315. Academic Press, Boston, MA, USA, 1991. <http://books.google.ie/books?id=iUZJAQAAIAAJ>
- Mou-Hsiung Chang. Stability of interconnected stochastic delay systems. *Applied Mathematics and Computation*, 16(4):277–295, 1985. DOI: 10.1016/0096-3003(85)90011-6
- Themistoklis Charalambous, Ioannis Kanellakopoulos, and Glenn Vinnicombe. On the stability of the Foschini-Miljanic algorithm with time-delays. In *CDC'08: 47th IEEE Conference on Decision and Control*, pages 2991–2996, Cancun, Mexico, Dec. 2008. DOI: 10.1109/CDC.2008.4739288
- Ben M. Chen, Zongli Lin, and Yacov Shamash. *Linear systems theory: A structural decomposition approach*. Control engineering. Birkhäuser, Boston, MA, USA, 2004. <http://books.google.ie/books?id=wH-H2Z9QIAYC>
- Qin Chen and J.Y.S. Luh. Coordination and control of a group of small mobile robots. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, volume 3, pages 2315–2320, San Diego, CA, USA, May 1994. DOI: 10.1109/ROBOT.1994.350940
- Frank J. Christophersen and Manfred Morari. Further results on “Infinity norms as Lyapunov functions for linear systems”. *IEEE Transactions on Automatic Control*, 52(3):547–553, Mar. 2007. DOI: 10.1109/TAC.2007.892378
- Delin Chu and Dragoslav D. Šiljak. A canonical form for the inclusion principle of dynamic systems. *SIAM Journal on Control and Optimization*, 44(3):969–990, 2005. DOI: 10.1137/040609616
- Fan R. K. Chung. *Spectral Graph Theory*. Number 92 in CBMS Regional Conference Series in Mathematics. American Mathematical Society, Providence, RI, USA, 1997. <http://goo.gl/fN90N>
- Roberto Conti. Quelques propriétés de l’opérateur d’évolution. *Colloquium Mathematicum*, pages 73–75, 1967.

- Peter A. Cook. On the stability of interconnected systems. *International Journal of Control*, 20(3):407–415, Sept. 1974. DOI: 10.1080/00207177408932751
- William A. Coppel. *Dichotomies in stability theory*, volume 629 of *Lecture notes in mathematics*. Springer-Verlag, New York, NY, USA, 1978. <http://books.google.ie/books?id=hTvAAAAMAAJ>
- Edward J. Davison and Amir G. Aghdam. *Decentralized Control of Large-Scale Systems*. Springer-Verlag New York, Inc., New York, NY, USA, Sept. 2011. <http://books.google.ie/books?id=2XjOQgAACAAJ>
- Wijesuriya P. Dayawansa and Clyde F. Martin. A converse Lyapunov theorem for a class of dynamical systems which undergo switching. *IEEE Transactions on Automatic Control*, 44(4):751–760, Apr. 1999. DOI: 10.1109/9.754812
- Hidde de Jong, Jean-Luc Gouzé, Céline Hernandez, Michel Page, Tewfik Sari, and Johannes Geiselmann. Qualitative simulation of genetic regulatory networks using piecewise-linear models. *Bulletin of Mathematical Biology*, 66(2):301–340, Mar. 2004. DOI: 10.1016/j.bulm.2003.08.010
- Cristobald de Kerchove and Paul Van Dooren. Reputation systems and nonnegativity. In POSTA'06, pages 3–16. DOI: 10.1007/978-3-642-02894-6_1
- Elena De Santis, Maria D. Di Benedetto, and Giordano Pola. Can linear stabilizability analysis be generalized to switching systems. In *MTNS'04: Proceedings of the 16th International Symposium on Mathematical Theory of Network and Systems*, Leuven, Belgium, Jul. 2004. <http://goo.gl/NT1ko>
- Raymond A. Decarlo, Michael S. Branicky, Stefan Pettersson, and Bengt Lennartson. Perspectives and results on the stability and stabilizability of hybrid systems. *Proceedings of the IEEE: Special Issue on Hybrid Systems*, 88(7):1069–1082, Jul. 2000. DOI: 10.1109/5.871309
- Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345):118–121, Mar. 1974. <http://www.jstor.org/stable/2285509>
- Department of Transport. Irish bulletin of vehicle of driver statistics 2009. Dublin, Ireland, Sept. 2010. <http://goo.gl/00VPA>
- Charles A. Desoer. Slowly varying system $\dot{x}(t) = a(t)x$. *IEEE Transactions on Automatic Control*, 14(6):780–781, Dec. 1969. DOI: 10.1109/TAC.1969.1099336
- Charles A. Desoer and Mathukumalli Vidyasagar. *Feedback Systems: Input-Output Properties*, volume 55 of *Classics in Applied Mathematics*. SIAM Press, Philadelphia, PA, USA, 2009. http://books.google.ie/books?id=k_7bRS1uSBQC
- Odo Diekmann, Stephan A. van Gils, Sjoerd M.V. Lunel, and Hans-Otto Walther. *Delay equations: functional-, complex-, and nonlinear analysis*, volume 110 of *Applied mathematical sciences*. Springer-Verlag, New York, NY, USA, 1995. <http://books.google.ie/books?id=NqnHuv1FCZ8C>
- Dimos V. Dimarogonas and Kostas J. Kyriakopoulos. Connectedness preserving distributed swarm aggregation for multiple kinematic robots. *IEEE Transactions on Robotics*, 24(5):1213–1223, October 2008. DOI: 10.1109/TRO.2008.2002313

- Xiuyong Ding and Lan Shu. Stability analysis of linear switched systems with mixed delays. *International Journal of Engineering and Applied Sciences*, 6(3):160–166, 2010.
<http://goo.gl/LgA96>
- Zahia Drici. New directions in the method of vector Lyapunov functions. *Journal of Mathematical Analysis and Applications*, 184(2):317–325, 1994.
DOI: 10.1006/jmaa.1994.1202
- Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, Mar. 1959.
<http://www.jstor.org/stable/2237130>
- Deborah Estrin, Lewis D. Girod, Gregory J. Pottie, and Mani Srivastava. Instrumenting the world with wireless sensor networks. In *ICASSP '01: Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 4, pages 2033–2036, Salt Lake City, UT, USA, May 2001.
DOI: 10.1109/ICASSP.2001.940390
- Lior Fainshil, Michael Margaliot, and Pavel Chigansky. On the stability of positive linear switched systems under arbitrary switching laws. *IEEE Transactions on Automatic Control*, 54(4):897–899, Apr. 2009.
DOI: 10.1109/TAC.2008.2010974
- Shaun M. Fallat and Steve Kirkland. Extremizing algebraic connectivity subject to graph theoretic constraints. *Electronic Journal of Linear Algebra*, 3:48–74, Apr. 1998.
<http://goo.gl/PqMca>
- Lei Fang, Panos J. Antsaklis, and Anastasis Tzimas. Asynchronous consensus protocols: Preliminary results, simulations and open questions. In *CDC-ECC '05: Proceedings of the Joint 44th IEEE Conference on Decision and Control and the 2005 European Control Conference*, pages 2194–2199, Seville, Spain, Dec. 2005.
DOI: 10.1109/CDC.2005.1582487
- Lorenzo Farina and Sergio Rinaldi. *Positive Linear Systems*, volume 50 of *Pure and Applied Mathematics: A Wiley Series of Texts, Monographs and Tracts*. John Wiley & Sons, Inc., New York, NY, USA, 2000.
<http://books.google.ie/books?id=6EjK1SoQfgMC>
- J. Alexander Fax and Richard M. Murray. Information flow and cooperative control of vehicle formations. *IEEE Transactions on Automatic Control*, 49(9):1465–1476, Sept. 2004.
DOI: 10.1109/TAC.2004.834433
- Gang Feng. Stability analysis of piecewise discrete-time linear systems. *IEEE Transactions on Automatic Control*, 47(7):1108–1112, Jul. 2002.
DOI: 10.1109/TAC.2002.800666
- Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(98):298–305, 1973.
<http://goo.gl/Ch6wG>
- Ettore Fornasini and Maria Elena Valcher. Controllability and reachability of 2-D positive systems: A graph theoretic approach. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 52(3):576–585, Mar. 2005.
DOI: 10.1109/TCSI.2004.842872
- Ettore Fornasini and Maria Elena Valcher. Linear copositive Lyapunov functions for continuous-time positive switched systems. *IEEE Transactions on Automatic Control*, 55(8):1933–1937, Aug. 2010.
DOI: 10.1109/TAC.2010.2049918
- Ettore Fornasini and Maria Elena Valcher. Stability and stabilizability criteria for discrete-time positive switched systems. Preprint submitted to *IEEE Transaction on Automatic Control*, Feb. 2011.

- Gerard J. Foschini and Zoran Miljanic. A simple distributed autonomous power control algorithm and its convergence. *IEEE Transactions on Vehicular Technology*, 42(4):641–646, Nov. 1993.
DOI: 10.1109/25.260747
- Friends of the Earth Trust. Road transport, air pollution and health. Briefing, London, UK, Mar. 1999.
<http://goo.gl/dZ9oh>
- Kelly S. Gallagher and Erich J. Muehlegger. Giving green to get green: Incentives and consumer adoption of hybrid vehicle technology. Technical Report RWP08-009, John F. Kennedy School of Government, Harvard University, Cambridge, MA, USA, Feb. 2008.
<http://ssrn.com/abstract=1083716>
- Shuzhi S. Ge and Zhendong Sun. *Switched linear systems: Control and design*. Communications and control engineering series. Springer, London, UK, 2005.
http://books.google.ie/books?id=F_aVFriAdroC
- Maria C. De Gennaro and Ali Jadbabaie. Decentralized control of connectivity for multi-agent systems. In *CDC'06: Proceedings of the 45th IEEE Conference on Decision and Control*, pages 3628–3633, San Diego, CA, USA, Dec. 2006.
DOI: 10.1109/CDC.2006.377041
- José C. Geromel and Patrizio Colaneri. Stability and stabilization of discrete time switched systems. *International Journal of Control*, 79(7):719–728, Jul. 2006a.
DOI: 10.1080/00207170600645974
- José C. Geromel and Patrizio Colaneri. Stability and stabilization of continuous-time switched linear systems. *SIAM Journal on Control and Optimization*, 45(5):1915–1930, Nov. 2006b.
DOI: 10.1137/050646366
- Laurent El Ghaoui and Silviu-Iulian Niculescu, editors. *Advances in Linear Matrix Inequality Methods in Control*. Advances in design and control. SIAM Press, Philadelphia, PA, USA, 2000.
<http://books.google.ie/books?id=15TfsbFw7WUC>
- Arpita Ghosh and Stephen P. Boyd. Growing well-connected graphs. In *CDC'06: Proceedings of the 45th IEEE Conference on Decision and Control*, pages 6605–6611, San Diego, CA, USA, Dec. 2006.
DOI: 10.1109/CDC.2006.377282
- Keith Godfrey. *Compartmental Models and their Applications*. Academic Press, New York, NY, USA, 1983.
<http://books.google.ie/books?id=xlGNAAAIAAJ>
- Roger Gorham. Air pollution from ground transportation: An assessment of causes, strategies and tactics, and proposed actions for the international community. In *Proc. WSSD-CSD-9*. United Nations Global Initiative on Transport Emissions, Commission on Sustainable Development, Apr. 2001.
<http://goo.gl/EA5dw>
- H. A. John Green. *Aggregation in Economic Analysis*. Princeton University Press, Princeton, NJ, USA, 1964.
- Ljubomir T. Grujić and Dragoslav D. Šiljak. Asymptotic stability and instability of large-scale systems. *IEEE Transactions on Automatic Control*, 18(6):636–645, Dec. 1973.
DOI: 10.1109/TAC.1973.1100422

- Ljubomir T. Grujić, Anatoliĭ A. Martynčuk, and Mania Ribbens-Pavella. *Large Scale Systems Stability under Structural and Singular Perturbations*. Lecture notes in control and information sciences. Springer-Verlag, 1987. <http://books.google.ie/books?id=XARRAAAMAAJ>
- Ljubomir T. Grujić and Dragoslav D. Šiljak. On stability of discrete composite systems. *IEEE Transactions on Automatic Control*, 18(5):522–524, Oct. 1973. DOI: 10.1109/TAC.1973.1100377
- Di Guo and Wilson J. Rugh. A stability result for linear parameter-varying systems. *Systems & Control Letters*, 24(1):1–5, Jan. 1995. DOI: 10.1016/0167-6911(94)00013-L
- Vijay Gupta. *Distributed Estimation and Control in Networked Systems*. Ph.D. thesis, Control and Dynamical Systems, California Institute of Technology, Pasadena, CA, USA, Jun. 2006. <http://goo.gl/NTtQ5>
- Leonid Gurvits, Robert N. Shorten, and Oliver Mason. On the stability of switched positive linear systems. *IEEE Transactions on Automatic Control*, 52(6):1099–1103, Jun. 2007. DOI: 10.1109/TAC.2007.899057
- Robert L. Gutmann and Anthony N. Michel. Second-order stochastic input-output stability and instability of interconnected systems. *Zeitschrift für Angewandte Mathematik und Mechanik*, 59(11):599–610, 1979a. DOI: 10.1002/zamm.19790591103
- Robert L. Gutmann and Anthony N. Michel. Stability of interconnected systems described by stochastic nonlinear Volterra integral equations. *SIAM Journal on Mathematical Analysis*, 10(2):217–236, 1979b. DOI: 10.1137/0510021
- Wassim M. Haddad and VijaySekhar Chellaboina. Stability theory for nonnegative and compartmental dynamical systems with time delay. *Systems & Control Letters*, 51(5):355–361, 2004. DOI: 10.1016/j.sysconle.2003.09.006
- Wassim M. Haddad, Qing Hui, Vijaysekhar Chellaboina, and Sergey Nersesov. Vector dissipativity theory for discrete-time large-scale nonlinear dynamical systems. *Advances in Difference Equations*, 2004(1):37–66, 2004. DOI: 10.1155/S1687183904310071
- Wassim M. Haddad, VijaySekhar Chellaboina, and Qing Hui. *Nonnegative and compartmental dynamical systems*. Princeton University Press, Princeton, NJ, USA, 2010. <http://books.google.ie/books?id=9otBECWXcWoC>
- Jack K. Hale and Sjoerd M. Verduyn Lunel. *Introduction to Functional Differential Equations*, volume 99 of *Applied Mathematical Sciences*. Springer-Verlag New York, Inc., 1993. <http://books.google.ie/books?id=c9wDSaKZaY0C>
- Frank Harary. *Graph Theory*. Westview Press, Inc., Boulder, CO, USA, 1969. <http://books.google.ie/books?id=U8XjzfrngT0C>
- Darald J. Hartfiel. *Markov Set-Chains*, volume 1695 of *Lecture Notes in Mathematics*. Springer Berlin / Heidelberg, Germany, 1998. <http://books.google.ie/books?id=79wZAQAIAAJ>
- Darald J. Hartfiel. *Nonhomogeneous Matrix Products*. World Scientific Publishing Co. Pte. Ltd., Singapore, 2002. <http://books.google.ie/books?id=QjNoQgAACAAJ>

- Arash Hassibi and Stephen Boyd. Quadratic stabilization and control of piecewise-linear systems. In *ACC'98: Proceedings of the 1998 American Control Conference*, volume 6, pages 3659–3664, Philadelphia, PA, USA, Jun. 1998. DOI: 10.1109/ACC.1998.703296
- Yuko Hatano and Mehran Mesbahi. Agreement over random networks. *IEEE Transactions on Automatic Control*, 50(11):1867–1872, Nov. 2005. DOI: 10.1109/TAC.2005.858670
- Simon S. Haykin. *Adaptive Filter Theory*. Prentice-Hall information and system sciences series. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2002.
<http://books.google.ie/books?id=eMcZAQAIAAJ>
- Jean-Claude Hennet and Sophie Tarbouriech. Stability and stabilization of delay differential systems. *Automatica*, 33(3):347–354, Mar. 1997. DOI: 10.1016/S0005-1098(96)00185-9
- João P. Hespanha. Stabilization through hybrid control. In *Encyclopedia of Life Support Systems (EOLSS)*, volume 15 of *Control Systems, Robotics, and Automation*. Developed under the Auspices of the UNESCO, EOLSS Publishers, Oxford, UK, May 2004a. <http://goo.gl/rWVx0>
- João P. Hespanha. Uniform stability of switched linear systems: Extensions of LaSalle's invariance principle. *IEEE Transactions on Automatic Control*, 49(4):470–482, Apr. 2004b. DOI: 10.1109/TAC.2004.825641
- João P. Hespanha and A. Stephen Morse. Stability of switched systems with average dwell-time. In *CDC'98: Proceedings of the 37th IEEE Conference on Decision and Control*, volume 3, pages 2655–2660, Phoenix, AZ, USA, Dec. 1999. DOI: 10.1109/CDC.1999.831330
- Robert C. Hilborn. *Chaos and Nonlinear Dynamics*. Oxford University Press, New York, NY, USA, 1994. <http://books.google.ie/books?id=fwybfh-nlyEC>
- David J. Hill and Peter J. Moylan. The stability of nonlinear dissipative systems. *IEEE Transactions on Automatic Control*, 21(5):708–711, Oct. 1976. DOI: 10.1109/TAC.1976.1101352
- David J. Hill and Peter J. Moylan. Dissipative dynamical systems: Basic input-output and state properties. *Journal of the Franklin Institute*, 309(5):327–357, 1980. <http://goo.gl/Ae8vF>
- David M. Himmelblau, editor. *Decomposition of Large-Scale Problems*. North-Holland Publishing Co., Amsterdam, The Netherlands, 1973. <http://books.google.ie/books?id=1jm4uqUrIYYC>
- Diederich Hinrichsen and Anthony J. Pritchard. Real and complex stability radii: A survey. In *Control of uncertain systems: Proceedings of an international workshop*, volume 6 of *Progress in systems and control theory*, pages 119–162, Bremen, Germany, Jun. 1989.
- Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Fundamentals of convex analysis*. Grundlehren text editions, Lecture Notes in Artificial Intelligence. Springer-Verlag Berlin / Heidelberg, Germany, 2001. http://books.google.ie/books?id=Ben6nm_yapMC
- Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, UK, 1985. <http://books.google.ie/books?id=9wTacOjHE6IC>
- Roger A. Horn and Charles R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, Cambridge, UK, 1991. <http://books.google.ie/books?id=LeuNXB2bl5EC>

- James E. Humphreys. *Introduction To Lie Algebras and Representation Theory*. Number 9 in Graduate Texts in Mathematics. Springer-Verlag, New York, NY, USA, 1972.
<http://books.google.ie/books?id=mSD4DvUFa6QC>
- Philipp Hövel. *Control of Complex Nonlinear Systems with Delay*. Springer Theses. Springer, 2010.
<http://books.google.ie/books?id=2HqQY-5J5TIC>
- Salim Ibrir. Stability and robust stabilization of discrete-time switched systems with time-delays: LMI approach. *Applied Mathematics and Computation*, 206(2):570–578, 2008.
 DOI: 10.1016/j.amc.2008.05.149
- IFAC TC 5.4. Website of the IFAC Technical Committee for Large Scale Complex Systems, 2010.
<http://tc.ifac-control.org/5/4>
- Masao Ikeda. Decentralized control of large scale systems. In Hendrik Nijmeijer, Jan C. Willems, and Johannes M. Schumacher, editors, *Three decades of mathematical system theory: A collection of surveys at the occasion of the 50th birthday of Jan C. Willems*, pages 219–242. Springer-Verlag New York, Inc., New York, NY, USA, 1989. http://books.google.ie/books?id=_PxQAAAAMAAJ
- Masao Ikeda, Dragoslav D. Šiljak, and D. E. White. Decentralized control with overlapping information sets. *Journal of Optimization Theory and Applications*, 34(2):279–310, 1981.
 10.1007/BF00935477
- Achim Ilchmann, David H. Owens, and Dieter Prätzel-Wolters. Sufficient conditions for stability of linear time-varying systems. *Systems & Control Letters*, 9(2):157–163, August 1987.
 DOI: 10.1016/0167-6911(87)90022-3
- John A. Jacquez and Carl P. Simon. Qualitative theory of compartmental systems. *SIAM Review*, 35(1):43–79, 1993.
 DOI: 10.1137/1035003
- John A. Jacquez and Carl P. Simon. Qualitative theory of compartmental systems with lags. *Mathematical Biosciences*, 180(1-2):329–362, Nov. 2002. DOI: 10.1016/S0025-5564(02)00131-1
- Ali Jadbabaie, Jie Lin, and A. Stephen Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on Automatic Control*, 48(6):988–1001, Jun. 2003.
 DOI: 10.1109/TAC.2003.812781
- Mohammad Jamshidi. *Large-Scale Systems: Modeling and Control*. North-Holland Publishing Co., New York, NY, USA, 1983.
<http://books.google.ie/books?id=5H9RAAAAMAAJ>
- Meng Ji and Magnus Egerstedt. Connectedness preserving distributed coordination control over dynamic graphs. In *ACC'05: Proceedings of the 2005 American Control Conference*, pages 93–98, Portland, OR, USA, June 2005.
 DOI: 10.1109/ACC.2005.1469914
- Meng Ji, Abubakr Muhammad, and Magnus Egerstedt. Leader-based multi-agent coordination: Controllability and optimal control. In *ACC'06: Proceedings of the 2006 American Control Conference*, pages 1358–1363, Minneapolis, MN, USA, Jun. 2006. DOI: 10.1109/ACC.2006.1656406
- Zhong-Ping Jiang. Decentralized control for large-scale nonlinear systems: A review of recent results. *Dynamics of Continuous, Discrete and Impulsive Systems, Series B: Algorithms & Applications*, 11(4-5b):537–552, 2004.
<http://goo.gl/ivJal>

- Mikael Johansson and Anders Rantzer. Computation of piecewise quadratic Lyapunov functions for hybrid systems. *IEEE Transactions on Automatic Control*, 43(4):555–559, Apr. 1998.
DOI: 10.1109/9.664157
- Mikael Johansson, Anders Rantzer, and Karl-Erik Arzén. Piecewise quadratic stability of fuzzy systems. *IEEE Transactions on Fuzzy Systems*, 7(6):713–722, Dec. 1999. DOI: 10.1109/91.811241
- Charles R. Johnson. Sufficient conditions for D-stability. *Journal of Economic Theory*, 9(1):53–62, Sept. 1974.
DOI: 10.1016/0022-0531(74)90074-X
- Charles R. Johnson, Volker Mehrmann, and D. Dale Olesky. Sign controllability of a nonnegative matrix and a positive vector. *SIAM Journal on Matrix Analysis and Applications*, 14(2):398–407, Apr. 1993. ISSN 0895-4798.
DOI: 10.1137/0614028
- Tadeusz Kaczorek. *Positive 1D and 2D systems*. Communications and control engineering. Springer-Verlag, London, UK, 2002.
http://books.google.ie/books?id=Y_wOjqrQZ4MC
- Akshay Kashyap, Tamer Basar, and Ramakrishnan Srikant. Quantized consensus. *Automatica*, 43(7):1192–1203, Jul. 2007.
DOI: 10.1016/j.automatica.2007.01.002
- Hassan K. Khalil. *Nonlinear Systems*. Macmillan Publishing Company, New York, NY, USA, 1992.
<http://books.google.ie/books?id=RVHvAAAAMAAJ>
- Christopher King and Michael Nathanson. On the existence of a common quadratic Lyapunov function for a rank one difference. *Linear Algebra and its Applications*, 419(2–3):400–416, Dec. 2006.
DOI: 10.1016/j.laa.2006.05.010
- Christopher King and Robert N. Shorten. A singularity test for the existence of common quadratic Lyapunov functions for pairs of stable LTI systems. In *ACC'04: Proceedings of the 2004 American Control Conference*, volume 4, pages 3881–3884, Jun. 2004.
DOI: 10.1109/ACC.2004.182893
- Christopher King and Robert N. Shorten. Singularity conditions for the non-existence of a common quadratic Lyapunov function for pairs of third order linear time invariant dynamic systems. *Linear Algebra and its Applications*, 413(1):24–35, Feb. 2006.
DOI: 10.1016/j.laa.2005.07.023
- Florian Knorn and Douglas J. Leith. Adaptive Kalman filtering for anomaly detection in software appliances. In *INFOCOM-ANM'08: Proceedings of the 1st IEEE Workshop on Automated Network Management*, pages 1–6, Phoenix, AZ, USA, Apr. 2008.
DOI: 10.1109/INFOCOM.2008.4544581
- Florian Knorn, Oliver Mason, and Robert N. Shorten. On linear co-positive Lyapunov functions for sets of linear positive systems. *Automatica*, 45(8):1943–1947, Aug. 2009a.
DOI: 10.1016/j.automatica.2009.04.013
- Florian Knorn, Oliver Mason, and Robert N. Shorten. Applications of linear co-positive Lyapunov functions for switched linear positive systems. In *POSTA'09*, pages 331–338.
DOI: 10.1007/978-3-642-02894-6_32
- Florian Knorn, Rade Stanojević, Martin Corless, and Robert N. Shorten. A problem in positive systems stability arising in topology control. In *POSTA'09*, pages 339–347.
DOI: 10.1007/978-3-642-02894-6_33

- Florian Knorn, Rade Stanojević, Martin J. Corless, and Robert N. Shorten. A framework for decentralised feedback connectivity control with application to sensor networks. *International Journal of Control*, 82(11):2095–2114, Nov. 2009d. DOI: 10.1080/00207170902912056
- Florian Knorn, Martin J. Corless, and Robert N. Shorten. Results in cooperative control and implicit consensus. *International Journal of Control*, 84(3):476–495, Mar. 2011a. DOI: 10.1080/00207179.2011.561442
- Florian Knorn, Martin J. Corless, and Robert N. Shorten. Results and application of cooperative control using implicit and constrained consensus. Submitted to *CDC-ECC '11: Joint 50th IEEE Conference on Decision and Control and the 2011 European Control Conference*, Mar. 2011b.
- Gabriel Kron. *Diakoptics: The Piecewise Solution of Large-Scale Systems*, volume 2. MacDonald, London, UK, 1963. <http://books.google.ie/books?id=1N9XAAAAMAAJ>
- Yang Kuang. *Delay differential equations: With applications in population dynamics*, volume 191 of *Mathematics in Science and Engineering*. Academic Press, San Diego, CA, USA, 1993. <http://books.google.ie/books?id=YWsz1IXkvy0C>
- Gangabam S. Ladde and Dragoslav D. Šiljak. Connective stability of large-scale stochastic systems. *International Journal of Systems Science*, 6(8):713–721, 1975. DOI: 10.1080/00207727508941855
- Vangipuram Lakshmikantham and Srinivasa Leela. *Differential and integral inequalities: Theory and applications.*, volume 55 of *Mathematics in Science and Engineering*. Academic Press, New York, NY, USA, 1969. http://books.google.ie/books?id=_sZ_SILHtiQC
- Vangipuram Lakshmikantham, Vladimir M. Matrosov, and Seenith Sivasundaram. *Vector Lyapunov Functions and Stability Analysis of Nonlinear Systems*. Mathematics and its Applications. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1991. <http://books.google.ie/books?id=Q0LP3mXiuFkC>
- Ben Lane and Stephen Potter. The adoption of cleaner vehicles in the UK: Exploring the consumer attitude-action gap. *Journal of Cleaner Production*, 15(11-012):1085–1092, 2007. DOI: 10.1016/j.jclepro.2006.05.026
- Eric L. Lasley and Anthony N. Michel. Input-output stability of interconnected systems. *IEEE Transactions on Automatic Control*, 21(1):84–89, Feb. 1976. DOI: 10.1109/TAC.1976.1101140
- Jonathan R. T. Lawton, Randal W. Beard, and Brett J. Young. A decentralized approach to formation maneuvers. *IEEE Transactions on Robotics and Automation*, 19(6):933–941, Dec. 2003. DOI: 10.1109/TRA.2003.819598
- Winston P. Ledet and David M. Himmelblau. Decomposition procedures for the solving of large-scale systems. *Advances in Chemical Engineering*, 8:185–254, 1970. DOI: 10.1016/S0065-2377(08)60185-4
- Ji-Woong Lee. Separation in stability analysis of piecewise linear systems in discrete time. In Magnus Egerstedt and Bud Mishra, editors, *Hybrid Systems: Computation and Control*, volume 4981 of *Lecture Notes in Computer Science*, pages 626–629. Springer Berlin / Heidelberg, 2008. DOI: 10.1007/978-3-540-78929-1_50

- George Leitmann. On one approach to the control of uncertain systems. *Journal of Dynamic Systems, Measurement, and Control*, 115(2B):373–380, 1993. DOI: 10.1115/1.2899077
- Bernard C. Levy, David A. Castañón, George C. Verghese, and Alan S. Willsky. A scattering framework for decentralized estimation problems. *Automatica*, 19(4):373–384, 1983. DOI: 10.1016/0005-1098(83)90051-1
- M. Anthony Lewis and Kar-Han Tan. High precision formation control of mobile robots using virtual structures. *Autonomous Robots*, 4(4):387–403, Oct. 1997. DOI: 10.1023/A:1008814708459
- Richard M. Lewis and Brian D. O. Anderson. Necessary and sufficient conditions for delay-independent stability of linear autonomous systems. *IEEE Transactions on Automatic Control*, 25(4):735–739, Aug. 1980. DOI: 10.1109/TAC.1980.1102420
- Ping Li, Shou-Ming Zhong, and Jin-Zhong Cui. Stability analysis of linear switching systems with time delays. *Chaos, Solitons & Fractals*, 40(1):474–480, Apr. 2009. DOI: 10.1016/j.chaos.2007.07.087
- Zhengguo Li, Yengchai Soh, and Changyun Wen. *Switched and impulsive systems: Analysis, design and applications*, volume 313 of *Lecture notes in control and information sciences*. Springer, Berlin, Germany, 2005. <http://books.google.ie/books?id=V5MKyVvdLjAC>
- Daniel Liberzon. *Switching in Systems and Control*. Systems & Control: Foundations & Applications. Birkhäuser, Boston, MA, USA, 2003. <http://books.google.ie/books?id=WN-P5njyf28C>
- Daniel Liberzon and A. Stephen Morse. Basic problems in stability and design of switched systems. *IEEE Control Systems Magazine*, 19(5):59–70, Oct. 1999. DOI: 10.1109/37.793443
- Daniel Liberzon and Roberto Tempo. Common lyapunov functions and gradient algorithms. *IEEE Transactions on Automatic Control*, 49(6):990–994, Jun. 2004. DOI: 10.1109/TAC.2004.829632
- Ching-Tai Lin. Structural controllability. *IEEE Transactions on Automatic Control*, 19(3):201–208, Jun. 1974. DOI: 10.1109/TAC.1974.1100557
- Hai Lin and Panos J. Antsaklis. A converse Lyapunov theorem for uncertain switched linear systems. In *CDC-ECC '05: Proceedings of the Joint 44th IEEE Conference on Decision and Control and the 2005 European Control Conference*, pages 3291–3296, Seville, Spain, Dec. 2005a. DOI: 10.1109/CDC.2005.1582669
- Hai Lin and Panos J. Antsaklis. Stability and persistent disturbance attenuation properties for a class of networked control systems: Switched system approach. *International Journal of Control*, 78(18):1447–1458, Dec. 2005b. DOI: 10.1080/00207170500329182
- Hai Lin and Panos J. Antsaklis. Stability and stabilizability of switched linear systems: A survey of recent results. *IEEE Transactions on Automatic Control*, 54(2):308–322, Feb. 2009. DOI: 10.1109/TAC.2008.2012009
- Hai Lin, Guisheng Zhai, and Panos J. Antsaklis. Robust stability and disturbance attenuation analysis of a class of networked control systems. In *CDC'03: Proceedings of the 42 IEEE Conference on Decision and Control*, volume 2, pages 1182–1187, Maui, HI, USA, Dec. 2003. DOI: 10.1109/CDC.2003.1272768

- Duyu Liu, Xinzhi Liu, and Shouming Zhong. Delay-dependent robust stability and control synthesis for uncertain switched neutral systems with mixed delays. *Applied Mathematics and Computation*, 202(2):828–839, Aug. 2008. DOI: 10.1016/j.amc.2008.03.028
- Kai Liu and Frank L. Lewis. An improved result on the stability analysis of nonlinear systems. *IEEE Transactions on Automatic Control*, 37(9):1425–1431, Sept. 1992. DOI: 10.1109/9.159586
- Xingwen Liu. Stability analysis of switched positive systems: A switched linear copositive Lyapunov function method. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 56(5):414–418, May 2009. DOI: 10.1109/TCSII.2009.2019326
- David G. Luenberger. *Introduction to Dynamic Systems. Theory, Models & Applications*. John Wiley & Sons, Inc., New York, NY, USA, 1979. <http://books.google.ie/books?id=mvlQAAAAMAAJ>
- Jan Lunze. Stability analysis of large-scale systems composed of strongly coupled similar subsystems. *Automatica*, 25(4):561–570, 1989. DOI: 10.1016/0005-1098(89)90098-8
- Jan Lunze. *Feedback Control of Large-Scale Systems*. Prentice-Hall International Series in Systems and Control Engineering. Prentice-Hall, London, UK, 1992. <http://books.google.ie/books?id=U91SAAAAMAAJ>
- Jan Lunze. *Regelungstechnik 1: Systemtheoretische Grundlagen, Analyse und Entwurf Einerschleifiger Regelungen*. Springer-Lehrbuch. Springer-Verlag Berlin / Heidelberg, Germany, 7th edition, 2008. DOI: 10.1007/978-3-540-68909-6
- Ren C. Luo and Michael G. Kay. Multisensor integration and fusion in intelligent systems. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(5):901–931, Sep./Oct. 1989. DOI: 10.1109/21.44007
- Aleksandr M. Lyapunov and Anthony T. Fuller. *The general problem of the stability of motion*. Control Theory and Applications Series. Taylor & Francis, London, UK, 1992. English translation of the original published 1892 in Russian by the Mathematical Society of Khrakov. http://books.google.ie/books?id=4tmAvU3_SCoC
- Magdi S. Mahmoud. *Switched Time-Delay Systems: Stability and Control*. Springer, New York, NY, USA, 2010. <http://books.google.ie/books?id=I3UmLoISQ0MC>
- A. D. Maĭzel'. On stability of solutions of systems of differential equations. *Trudy Ural'sk. Politekhn. Inst., Ser. Mat.*, 51:20–50, 1954.
- José-Luis Mancilla-Aguilar and Rafael A. García. A converse Lyapunov theorem for nonlinear switched systems. *Systems & Control Letters*, 41(1):67–71, 2000. DOI: 10.1016/S0167-6911(00)00040-2
- José-Luis Mancilla-Aguilar and Rafael A. García. On converse Lyapunov theorems for ISS and iISS switched nonlinear systems. *Systems & Control Letters*, 42(1):47–53, Jan. 2001. ISSN 0167-6911. DOI: 10.1016/S0167-6911(00)00079-7
- Michael Margaliot and Michael S. Branicky. Stability analysis of positive linear switched systems: A variational approach. Preprint submitted to *Automatica*, Jun. 2009. <http://goo.gl/3SXNz>

- Michael Margaliot and Christos A. Yfoulis. Absolute stability of third-order systems: A numerical algorithm. *Automatica*, 42(10):1705–1711, 2006. DOI: 10.1016/j.automatica.2006.04.025
- Anatoliĭ A. Martyniuk. *Stability by Liapunov's matrix function method with applications*. Monographs and textbooks in pure and applied mathematics. Marcel Dekker, 1998.
<http://books.google.ie/books?id=87HVPjyH6uIC>
- Anatoliĭ A. Martyniuk. *Qualitative methods in nonlinear dynamics: Novel approaches to Liapunov's matrix functions*. Monographs and textbooks in pure and applied mathematics. Marcel Dekker, New York, NY, USA, 2002.
<http://books.google.ie/books?id=SSWdhgfsOmMC>
- Anatoliĭ A. Martyniuk, V. Miladzhyanov, and M. Muminov. Stability of large-scale discrete systems under structural perturbations. *Ukrainian Mathematical Journal*, 48(10):1533–1545, 1996.
DOI: 10.1007/BF02377822
- Oliver Mason and Robert N. Shorten. Some results on the stability of positive switched linear systems. In *CDC'04: Proceedings of the 43th IEEE Conference on Decision and Control*, volume 5, pages 4601–4606, Atlantis, Paradise Island, Bahamas, Dec. 2004. DOI: 10.1109/CDC.2004.1429509
- Oliver Mason and Robert N. Shorten. On linear copositive Lyapunov functions and the stability of switched positive linear systems. *IEEE Transactions on Automatic Control*, 52(7):1346–1349, Jul. 2007.
DOI: 10.1109/TAC.2007.900857
- Oliver Mason and Mark Verwoerd. Observations on the stability properties of cooperative systems. *Systems & Control Letters*, 58(6):461–467, Jun. 2009.
DOI: 10.1016/j.sysconle.2009.02.009
- Oliver Mason, Vahid Bokharaie, and Robert N. Shorten. Stability and D-stability for switched positive systems. In *POSTA'09*, pages 101–109.
DOI: 10.1007/978-3-642-02894-6_10
- Vladimir M. Matrosov. Method of Lyapunov-vector functions in feedback systems. *Automation and Remote Control*, 33(9):1458–1469, Sept. 1972.
- Vladimir M. Matrosov. The method of vector Lyapunov functions in analysis of composite systems with distributed parameters. *Automation and Remote Control*, 34(1):1–6, Jan. 1973.
- Mayor of London. Central London Congestion Charging, Impacts Monitoring — Sixth Annual Report. Transport for London., Jul. 2008.
<http://goo.gl/xAsfa>
- Mehran Mesbahi and Fred Y. Hadaegh. Formation flying control of multiple spacecraft via graphs, matrix inequalities, and switching. In *CCA'99: Proceedings of the 1999 IEEE International Conference on Control Applications*, volume 2, pages 1211–1216, Kohala Coast-Island of Hawai'i, HI, USA, Aug. 1999.
DOI: 10.1109/CCA.1999.801145
- Lloyd A. Metzler. Stability of multiple markets: The Hicks conditions. *Econometrica*, 13(4): 277–292, Oct. 1945.
<http://www.jstor.org/stable/1906922>
- Sean P. Meyn. *Control Techniques for Complex Networks*. Cambridge University Press, New York, NY, USA, 2008.
- Anthony N. Michel. Stability analysis of stochastic large-scale systems. *Zeitschrift für Angewandte Mathematik und Mechanik*, 55(2):113–123, 1975a.
DOI: 10.1002/zamm.19750550207

- Anthony N. Michel. Stability analysis of stochastic composite systems. *IEEE Transactions on Automatic Control*, 20(2):246–250, Apr. 1975b. DOI: 10.1109/TAC.1975.1100903
- Anthony N. Michel. Scalar vs. vector Lyapunov functions in stability analysis of large scale systems: Raprochement. In *CDC'77: Proceedings of the 16th IEEE Conference on Decision and Control, including the 16th Symposium on Adaptive Processes and A Special Symposium*, volume 16, pages 1262–1266, New Orleans, LA, USA, Dec. 1977. DOI: 10.1109/CDC.1977.271763
- Anthony N. Michel. On the status of stability of interconnected systems. *IEEE Transactions on Automatic Control*, 28(6):639–653, Jun. 1983. DOI: 10.1109/TAC.1983.1103292
- Anthony N. Michel. Recent trends in the stability analysis of hybrid dynamical systems. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 46(1):120–134, Jan. 1999. DOI: 10.1109/81.739260
- Anthony N. Michel and Richard K. Miller. *Qualitative Analysis of Large Scale Dynamical Systems*, volume 134 of *Mathematics in Science and Engineering*. Academic Press, New York, NY, USA, 1977. <http://books.google.ie/books?id=-Md2CcjIGWgC>
- Anthony N. Michel and Richard K. Miller. Qualitative analysis of interconnected systems described on Banach spaces: Well posedness, instability and lagrange stability. *Zeitschrift für Angewandte Mathematik und Mechanik*, 58(5):289–300, 1978. DOI: 10.1002/zamm.19780580507
- Anthony N. Michel and David W. Porter. Analysis of discontinuous large-scale systems: Stability, transient behaviour and trajectory bounds. *International Journal of Systems Science*, 2(1):77–95, Jul. 1971. DOI: 10.1080/00207727108920179
- Anthony N. Michel and David W. Porter. Stability analysis of composite systems. *IEEE Transactions on Automatic Control*, 17(2):222–226, Apr. 1972. DOI: 10.1109/TAC.1972.1099952
- Anthony N. Michel, Richard K. Miller, and Wang Tang. Lyapunov stability of interconnected systems: Decomposition into strongly connected subsystems. *IEEE Transactions on Circuits and Systems*, 25(9):799–809, Sept. 1978. DOI: 10.1109/TCS.1978.1084537
- Anthony N. Michel, Richard K. Miller, and Boo Hee Nam. Stability analysis of interconnected systems using computer generated lyapunov functions. *IEEE Transactions on Circuits and Systems*, 29(7):431–440, Jul. 1982. DOI: 10.1109/TCS.1982.1085181
- Richard K. Miller and Anthony N. Michel. *Ordinary Differential Equations*. Dover Publications, New York, NY, USA, 2007. <http://books.google.ie/books?id=YznLGAAACAAJ>
- Alexander P. Molchanov and Evgenii S. Pyatnitskii. Criteria of asymptotic stability of differential and difference inclusions encountered in control theory. *Systems & Control Letters*, 13(1):59–64, Jul. 1989. DOI: 10.1016/0167-6911(89)90021-2
- Luc Moreau. Stability of multiagent systems with time-dependent communication links. *IEEE Transactions on Automatic Control*, 50(2):169–182, Feb. 2005. DOI: 10.1109/TAC.2004.841888
- Takehiro Mori, Norio Fukuma, and Michiyoshi Kuwahara. Simple stability criteria for single and composite linear systems with time delays. *International Journal of Control*, 34(6):1175–1184, Dec. 1981. DOI: 10.1080/00207178108922590

- Yoshihiro Mori, Takehiro Mori, and Yasuaki Kuroe. A solution to the common Lyapunov function problem for continuous-time systems. In *CDC'97: Proceedings of the 36th IEEE Conference on Decision and Control*, volume 4, pages 3530–3531, Dec. 1997. DOI: 10.1109/CDC.1997.652397
- A. Stephen Morse. Supervisory control of families of linear set-point controllers—Part I: Exact matching. *IEEE Transactions on Automatic Control*, 41(10):1413–1431, Oct. 1996. DOI: 10.1109/9.539424
- Peter J. Moylan. Matrices with positive principal minors. *Linear Algebra and its Applications*, 17(1):53–58, 1977. DOI: 10.1016/0024-3795(77)90040-4
- Peter J. Moylan and David J. Hill. Stability criteria for large-scale systems. *IEEE Transactions on Automatic Control*, 23(2):143–149, Apr. 1978. DOI: 10.1109/TAC.1978.1101721
- Richard M. Murray. Recent research in cooperative control of multivehicle systems. *Journal of Dynamic Systems, Measurement, and Control*, 129(5):571–583, Sept. 2007. DOI: 10.1115/1.2766721
- Roderick Murray-Smith and Robert N. Shorten, editors, Roderick Murray-Smith and Robert N. Shorten. *Switching and Learning in Feedback Systems: European Summer School on Multi-Agent Control*, Maynooth, Ireland, Sept. 2003. Springer, Berlin, Germany. <http://books.google.ie/books?id=dcXiy-2R4pAC>
- Arthur G. O. Mutambara. *Decentralized Estimation and Control for Multisensor Systems*. CRC Press, Inc., Boca Raton, FL, USA, 1998. <http://books.google.ie/books?id=Z1YfUGkG8poC>
- Max Müller. Über das Fundamentaltheorem in der Theorie der gewöhnlichen Differentialgleichungen. *Mathematische Zeitschrift*, 26:619–645, 1926. <http://goo.gl/5rxrb>
- Kumpati S. Narendra and Jeyendran Balakrishnan. A common Lyapunov function for stable LTI systems with commuting A -matrices. *IEEE Transactions on Automatic Control*, 39(12):2469–2471, Dec. 1994. DOI: 10.1109/9.362846
- Kumpati S. Narendra and Robert N. Shorten. Hurwitz stability of metzler matrices. *IEEE Transactions on Automatic Control*, 55(6):1484–1487, Jun. 2010. DOI: 10.1109/TAC.2010.2045694
- Kumpati S. Narendra and James H. Taylor. *Frequency domain criteria for absolute stability*. Electrical science series. Academic Press, New York, NY, USA, 1973. <http://books.google.ie/books?id=J-9qQgAACAAJ>
- Angelia Nedić and Asuman Ozdaglar. Convergence rate for consensus with delays. *Journal of Global Optimization*, 47(3):437–456, Jul. 2010. DOI: 10.1007/s10898-008-9370-2
- Sergey G. Nersesov and Wassim M. Haddad. On the stability and control of nonlinear dynamical systems via vector Lyapunov functions. *IEEE Transactions on Automatic Control*, 51(2):203–215, Feb. 2006. DOI: 10.1109/TAC.2005.863496
- Torsten Norvig. Consensus of subjective probabilities: A convergence theorem. *The Annals of Mathematical Statistics*, 38(1):221–225, Feb. 1967. <http://www.jstor.org/stable/2238885>

- Yuzo Ohta. Qualitative analysis of nonlinear quasi-monotone dynamical systems described by functional-differential equations. *IEEE Transactions on Circuits and Systems*, 28(2):138–144, Feb. 1981. DOI: 10.1109/TCS.1981.1084959
- Reza Olfati-Saber. Flocking for multi-agent dynamic systems: Algorithms and theory. *IEEE Transactions on Automatic Control*, 51(3):401–420, Mar. 2006. DOI: 10.1109/TAC.2005.864190
- Reza Olfati-Saber. Distributed kalman filtering for sensor networks. In *CDC'07: Proceedings of the 46th IEEE Conference on Decision and Control*, pages 5492–5498, New Orleans, LA, USA, Dec. 2007. DOI: 10.1109/CDC.2007.4434303
- Reza Olfati-Saber and Richard M. Murray. Consensus protocols for networks of dynamic agents. In *ACC'03: Proceedings of the 2003 American Control Conference*, volume 2, pages 951–956, Denver, CO, USA, Jun. 2003. <http://www.cds.caltech.edu/~murray/preprints/om03-acc.pdf>
- Reza Olfati-Saber and Richard M. Murray. Consensus problems in networks of agents with switching topology and time-delays. *IEEE Transactions on Automatic Control*, 49(9):1520–1533, Sept. 2004. DOI: 10.1109/TAC.2004.834113
- Reza Olfati-Saber and Jeff S. Shamma. Consensus filters for sensor networks and distributed sensor fusion. In *CDC-ECC '05: Proceedings of the Joint 44th IEEE Conference on Decision and Control and the 2005 European Control Conference*, pages 6698–6703, Dec. 2005. DOI: 10.1109/CDC.2005.1583238
- Reza Olfati-Saber, J. Alex Fax, and Richard M. Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, Jan. 2007. DOI: 10.1109/JPROC.2006.887293
- Tatsushi Ooba and Yasuyuki Funahashi. Stability robustness for linear state space models: A Lyapunov mapping approach. *Systems & Control Letters*, 29(4):191–196, January 1997a. DOI: 10.1016/S0167-6911(96)00068-0
- Tatsushi Ooba and Yasuyuki Funahashi. Two conditions concerning common quadratic Lyapunov functions for linear systems. *IEEE Transactions on Automatic Control*, 42(5):719–722, May 1997b. DOI: 10.1109/9.580888
- Tatsushi Ooba and Yasuyuki Funahashi. On a common quadratic lyapunov function for widely distant systems. *IEEE Transactions on Automatic Control*, 42(12):1697–1699, Dec. 1997c. DOI: 10.1109/9.650019
- Tatsushi Ooba and Yasuyuki Funahashi. On the simultaneous diagonal stability of linear discrete-time systems. *Systems & Control Letters*, 36(3):175–180, Mar. 1999. DOI: 10.1016/S0167-6911(98)00082-6
- Alan V. Oppenheim, Alan S. Willsky, and S. Hamid Nawab. *Signals & Systems*. Prentice-Hall signal processing series. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2nd edition, 1996. <http://books.google.ie/books?id=QuNgQgAACAAJ>
- Antonis Papachristodoulou and Stephen Prajna. Robust stability analysis of nonlinear hybrid systems. *IEEE Transactions on Automatic Control*, 54(5):1035–1041, May 2009. DOI: 10.1109/TAC.2009.2017155

- Lynne E. Parker. ALLIANCE: an architecture for fault tolerant multirobot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, Apr. 1998. DOI: 10.1109/70.681242
- Philippos Peleties and Raymond DeCarlo. Asymptotic stability of m -switched systems using Lyapunov-like functions. In *ACC'91: Proceedings of the 1991 American Control Conference*, pages 1679–1684, Boston, MA, USA, Jun. 1991. <http://goo.gl/W4hiT>
- Philippos Peleties and Raymond DeCarlo. Asymptotic stability of m -switched systems using Lyapunov functions. In *CDC'92: Proceedings of the 31st IEEE Conference on Decision and Control*, volume 4, pages 3438–3439, Dec. 1992. DOI: 10.1109/CDC.1992.371213
- Mathew Penrose. *Random Geometric Graphs*, volume 5 of *Oxford Studies in Probability*. Oxford University Press, New York, NY, USA, Jun. 2003. <http://books.google.ie/books?id=M38e7nPGSCsC>
- Oskar Perron. Die Stabilitätsfrage bei Differentialgleichungen. *Mathematische Zeitschrift*, 32:703–728, 1930. DOI: 10.1007/BF01194662
- Karel Perutka. A survey of decentralized adaptive control. In Meng Joo Er, editor, *New Trends in Technologies: Control, Management, Computational Intelligence and Network Systems*, chapter 1. InTech, Rijeka, Croatia, Nov. 2010. <http://goo.gl/ccgeb>
- Stefan Pettersson and Bengt Lennartson. Stability and robustness for hybrid systems. In *CDC'96: Proceedings of the 35th IEEE Conference on Decision and Control*, volume 2, pages 1202–1207, Kobe, Japan, Dec. 1996. DOI: 10.1109/CDC.1996.572653
- Stefan Pettersson and Bengt Lennartson. Lmi for stability and robustness of hybrid systems. In *ACC'97: Proceedings of the 1997 American Control Conference*, volume 3, pages 1714–1718, Albuquerque, NM, USA, Jun. 1997. DOI: 10.1109/ACC.1997.610877
- Stefan Pettersson and Bengt Lennartson. Hybrid system stability and robustness verification using linear matrix inequalities. *International Journal of Control*, 75(16–17):1335–1355, 2002. DOI: 10.1080/0020717021000023762
- A. A. Piontkovskii and L. D. Rutkovskaya. Investigation of stability theory problems by the vector Lyapunov function method. *Automation and Remote Control*, 10:1422–1429, 1967.
- Giordano Pola, Jan Willem Polderman, and Maria D. Di Benedetto. Balancing dwell times for switching linear systems. In *MTNS'04: Proceedings of the 16th International Symposium on Mathematical Theory of Network and Systems*, Leuven, Belgium, Jul. 2004. <http://goo.gl/vpYm0>
- Andrzej Polański. On infinity norms as Lyapunov functions for linear systems. *IEEE Transactions on Automatic Control*, 40(7):1270–1274, Jul. 1995. DOI: 10.1109/9.400479
- Andrzej Polański. Lyapunov function construction by linear programming. *IEEE Transactions on Automatic Control*, 42(7):1013–1016, Jul. 1997. DOI: 10.1109/9.599986
- Andrzej Polański. On absolute stability analysis by polyhedral Lyapunov functions. *Automatica*, 36(4):573–578, 2000. DOI: 10.1016/S0005-1098(99)00180-6
- David W. Porter and Anthony N. Michel. Input-output stability of time-varying nonlinear multiloop feedback systems. *IEEE Transactions on Automatic Control*, 19(4):422–427, Aug. 1974. DOI: 10.1109/TAC.1974.1100582

POSTA'03: Proceedings of the first Multidisciplinary International Symposium on Positive Systems: Theory and Applications, volume 294 of *Lecture Notes in Control and Information Sciences*, Rome, Italy, Aug. 2003. Springer-Verlag Berlin / Heidelberg.

<http://books.google.ie/books?id=jMDXlc-hC4lC>

POSTA'06: Proceedings of the first Multidisciplinary International Symposium on Positive Systems: Theory and Applications, volume 341 of *Lecture Notes in Control and Information Sciences*, Grenoble, France, Aug. 2006. Springer-Verlag Berlin / Heidelberg.

<http://books.google.ie/books?id=jMDXlc-hC4lC>

POSTA'09: Proceedings of the third Multidisciplinary International Symposium on Positive Systems: Theory and Applications, volume 389 of *Lecture Notes in Control and Information Sciences*, Valencia, Spain, Sept. 2009. Springer-Verlag Berlin / Heidelberg.

<http://books.google.ie/books?id=pykXyhh1S3kC>

Stephen Prajna and Antonis Papachristodoulou. Analysis of switched and hybrid systems - beyond piecewise quadratic methods. In *ACC'03: Proceedings of the 2003 American Control Conference*, volume 4, pages 2779–2784, Denver, CO, USA, Jun. 2003. DOI: 10.1109/ACC.2003.1243743

Ram Ramanathan and Regina Rosales-Hain. Topology control of multihop wireless networks using transmit power adjustment. In *INFOCOM'00: Proceedings of the 19th Annual Joint Conference of the IEEE Computer and Communications Societies.*, volume 2, pages 404–413, Tel Aviv, Israel, Mar. 2000. DOI: 10.1109/INFCOM.2000.832213

M. Ait Rami. Stability analysis and synthesis for linear positive systems with time-varying delays. In *POSTA'09*, pages 205–215. DOI: 10.1007/978-3-642-02894-6_20

M. Ait Rami and Fernando Tadeo Rico. Controller synthesis for positive linear systems with bounded controls. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 54(2):151–155, Feb. 2007. DOI: 10.1109/TCSII.2006.886888

Robert D. Rasmussen and Anthony N. Michel. On vector Lyapunov functions for stochastic dynamical systems. *IEEE Transactions on Automatic Control*, 21(2):250–254, Apr. 1976a. DOI: 10.1109/TAC.1976.1101172

Robert D. Rasmussen and Anthony N. Michel. Stability of interconnected dynamical systems described on Banach spaces. *IEEE Transactions on Automatic Control*, 21(4):464–471, Aug. 1976b. DOI: 10.1109/TAC.1976.1101302

Wei Ren and Randal W. Beard. Formation feedback control for multiple spacecraft via virtual structures. *IEEE Proceedings of Control Theory and Applications*, 151(3):357–368, May 2004. DOI: 10.1049/ip-cta:20040484

Wei Ren and Randal W. Beard. Consensus seeking in multiagent systems under dynamically changing interaction topologies. *IEEE Transactions on Automatic Control*, 50(5):655–661, May 2005. DOI: 10.1109/TAC.2005.846556

Wei Ren, Randal W. Beard, and Ella M. Atkins. A survey of consensus problems in multi-agent coordination. In *ACC'05: Proceedings of the 2005 American Control Conference*, volume 3, pages 1859–1864, Jun. 2005. DOI: 10.1109/ACC.2005.1470239

Craig Reynolds. Boids (flocks, herds, and schools: a distributed behavioral model), Sept. 2001. Website. <http://www.red3d.com/cwr/boids/>

- Craig W. Reynolds. Flocks, herds, and schools: A distributed behavioral model. In *SIG-GRAPH'87: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, pages 25–34, Anaheim, CA, USA, Jul. 1987. <http://goo.gl/BFfxO>
- R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, NJ, USA, 1970. http://books.google.ie/books?id=wj4Fh4h_V7QC
- Howard H. Rosenbrock. A Lyapunov function with applications to some nonlinear physical systems. *Automatica*, 1(1):31–53, 1963. DOI: 10.1016/0005-1098(63)90005-0
- Shigui Ruan. Connective stability of discontinuous large scale systems. *Journal of Mathematical Analysis and Applications*, 160(2):480–484, 1991. DOI: 10.1016/0022-247X(91)90320-Y
- Walter Rudin. *Principles of Mathematical Analysis*. International series in pure and applied mathematics. McGraw-Hill, New York, NY, USA, third, revised edition, 1976. <http://books.google.ie/books?id=kwqzPAAACAAJ>
- Björn S. Rüffer, Christopher M. Kellett, and Peter Dower. On copositive Lyapunov functions for a class of monotone systems. In *MTNS'10: Proceedings of the 19th International Symposium on Mathematical Theory of Network and Systems*, Budapest, Hungary, Jul. 2010. <http://goo.gl/IMxw2>
- Andrew P. Sage. *Methodology for Large-Scale Systems*. McGraw-Hill, New York, NY, USA, 1977. <http://books.google.ie/books?id=Om5RAAAAMAAJ>
- Irwin W. Sandberg. A frequency-domain condition for the stability of feedback systems containing a single time-varying nonlinear element. *Bell System Technical Journal*, 43(4):1601–1608, May 1964. <http://goo.gl/ltlxq>
- Irwin W. Sandberg. On the stability of interconnected systems. In *ISCAS'78: Proceedings of the 1978 IEEE Symposium on Circuits and Systems*, pages 228–231, New York, NY, USA, May 1978.
- Nils R. Sandell, Jr., Pravin P. Varaiya, Michael Athans, and Michael G. Safonov. Survey of decentralized control methods for large scale systems. *IEEE Transactions on Automatic Control*, 23(2):108–128, Apr. 1978. DOI: 10.1109/TAC.1978.1101704
- Paolo Santi. *Topology Control in Wireless Ad Hoc and Sensor Networks*. John Wiley & Sons, Inc., Chichester, UK, Sept. 2005. <http://books.google.ie/books?id=ZAtbAAAACAAJ>
- Roger W. H. Sargent and Arthur W. Westerberg. SPEED-UP (simulation programme for the economic evaluation and design of unsteady-state processes) in chemical engineering design. *Transactions of the Institution of Chemical Engineers*, 42:190–197, 1964.
- Luca Schenato and Giovanni Gamba. A distributed consensus protocol for clock synchronization in wireless sensor network. In *CDC'07: Proceedings of the 46th IEEE Conference on Decision and Control*, pages 2289–2294, New Orleans, LA, USA, Dec. 2007. DOI: 10.1109/CDC.2007.4434671
- Jen Schoenburg. Germany introduces environmental zones. *Frontiers in Ecology and the Environment*, 6(1):5, Feb. 2008. DOI: 10.1890/1540-9295(2008)6[4:D]2.0.CO;2

- Henrik Schulze and Christian Lüders. *Theory and applications of OFDM and CDMA: Wideband wireless communications*. John Wiley & Sons, Ltd., Chichester, UK, 2005.
<http://books.google.ie/books?id=PzWSwn9gHWEc>
- George R. Sell. Stability theory and Lyapunov's second method. *Archive for Rational Mechanics and Analysis*, 14:108–126, Jan. 1963.
 DOI: 10.1007/BF00250695
- M. Erol Sezer and Dragoslav D. Šiljak. Nested ϵ -decompositions and clustering of complex systems. *Automatica*, 22(3):321–331, May 1986.
 DOI: 10.1016/0005-1098(86)90030-0
- M. Erol Sezer and Dragoslav D. Šiljak. Nested ϵ -decompositions of linear systems: Weakly coupled and overlapping blocks. *SIAM Journal on Matrix Analysis and Applications*, 12(3):521–533, Jul. 1991.
 DOI: 10.1137/0612037
- M. Erol Sezer and Dragoslav D. Šiljak. Decentralized control. In William S. Levine, editor, *The Control Handbook*, chapter 49. CRC Press, Boca Raton, FL, USA, 1996.
<http://books.google.ie/books?id=2WQP5JGaJOGC>
- Hong Shi, Long Wang, and Tianguang Chu. Virtual leader approach to coordinated control of multiple mobile agents with asymmetric interactions. *Physica D: Nonlinear Phenomena*, 213(1):51–65, 2006.
 DOI: 10.1016/j.physd.2005.10.012
- Robert N. Shorten and Fiacre Ó. Cairbre. On the stability of pairwise triangularisable and related switching systems. In *ACC'01: Proceedings of the 2001 American Control Conference*, volume 3, pages 1882–1883, Jun. 2001a.
 DOI: 10.1109/ACC.2001.946011
- Robert N. Shorten and Fiacre Ó. Cairbre. A proof of global attractivity for a class of switching systems using a non-quadratic Lyapunov approach. *IMA Journal of Mathematical Control and Information*, 18(3):341–353, Sept. 2001b.
 DOI: 10.1093/imamci/18.3.341
- Robert N. Shorten and Fiacre Ó. Cairbre. A new methodology for the stability analysis of pairwise triangularizable and related switching systems. *IMA Journal of Applied Mathematics*, 67(5):441–457, Oct. 2002.
 DOI: 10.1093/imamat/67.5.441
- Robert N. Shorten and Kumpati S. Narendra. On the stability and existence of common Lyapunov functions for stable linear switching systems. In *CDC'98: Proceedings of the 37th IEEE Conference on Decision and Control*, volume 4, pages 3723–3724, Dec. 1998.
 DOI: 10.1109/CDC.1998.761788
- Robert N. Shorten and Kumpati S. Narendra. Necessary and sufficient conditions for the existence of a common quadratic Lyapunov function for M stable second order linear time-invariant systems. In *ACC'00: Proceedings of the 2000 American Control Conference*, volume 1, pages 359–363, Chicago, IL, USA, Jun. 2000.
 DOI: 10.1109/ACC.2000.878913
- Robert N. Shorten and Kumpati S. Narendra. Necessary and sufficient conditions for the existence of a common quadratic Lyapunov function for a finite number of stable second order linear time-invariant systems. *International Journal of Adaptive Control and Signal Processing*, 16(10):709–728, Dec. 2002.
 DOI: 10.1002/acs.719
- Robert N. Shorten and Kumpati S. Narendra. On common quadratic Lyapunov functions for pairs of stable LTI systems whose system matrices are in companion form. *IEEE Transactions on Automatic Control*, 48(4):618–621, Apr. 2003.
 DOI: 10.1109/TAC.2003.809795

- Robert N. Shorten, Fabian Wirth, and Douglas J. Leith. A positive systems model of tcp-like congestion control: asymptotic results. *IEEE/ACM Transactions on Networking*, 14(3):616–629, Jun. 2006. DOI: 10.1109/TNET.2006.876178
- Robert N. Shorten, Fabian Wirth, Oliver Mason, Kai Wulff, and Christopher King. Stability criteria for switched and hybrid systems. *SIAM Review*, 49(4):545–592, Nov. 2007. DOI: 10.1137/05063516X
- Dragoslav D. Šiljak. Stability of large-scale systems under structural perturbations. *IEEE Transactions on Systems, Man, and Cybernetics*, 2(5):657–663, Nov. 1972. DOI: 10.1109/TSMC.1972.4309194
- Dragoslav D. Šiljak. *Large-Scale Dynamic Systems: Stability and Structure*. North-Holland Publishing Co., New York, NY, USA, 1978. http://books.google.ie/books?id=O_tQAAAAAMAAJ
- Dragoslav D. Šiljak. Complex dynamical systems: Dimensionality, structure and uncertainty. *Large Scale Systems In Information And Decision Technologies*, 4(3):279–294, Jun. 1983. <http://goo.gl/sCM4D>
- Dragoslav D. Šiljak. *Decentralized Control of Complex Systems*, volume 184 of *Mathematics in science and engineering*. Academic Press, San Diego, CA, USA, 1991. <http://books.google.ie/books?id=8LQUfp7ns40C>
- Dragoslav D. Šiljak. Decentralized control and computations: Status and prospects. *Annual Reviews in Control*, 20:131–141, May 1996. DOI: 10.1016/S1367-5788(97)00011-4
- Dragoslav D. Šiljak and Dusan M. Stipanović. Robust stabilization of nonlinear systems: The LMI approach. *Mathematical Problems in Engineering*, 6(5):461–493, 2000. DOI: 10.1155/S1024123X00001435
- Dragoslav D. Šiljak and Aleksandar I. Zečević. Large-scale and decentralized systems. In *Wiley Encyclopedia of Electrical and Electronics Engineering*. John Wiley & Sons, Inc., 1999. DOI: 10.1002/047134608X.W1021
- Dragoslav D. Šiljak and Aleksandar I. Zečević. Control of large-scale systems: Beyond decentralized feedback. *Annual Reviews in Control*, 29(2):169–179, Oct. 2005. DOI: 10.1016/j.arcontrol.2005.08.003
- Herbert A. Simon and Albert Ando. Aggregation of variables in dynamical systems. *Econometrica*, 29:111–138, 1961.
- Sahjendra N. Singh, Phil Chandler, Corey Schumacher, Siva Banda, and Meir Pachter. Nonlinear adaptive close formation control of unmanned aerial vehicles. *Dynamics and Control*, 10(2): 179–194, Apr. 2000. DOI: 10.1023/A:1008348025564
- Lesław Socha. The asymptotic stochastic stability in large of the composite stochastic systems. *Automatica*, 22(5):605–610, 1986. DOI: 10.1016/0005-1098(86)90071-3
- Selim Solmaz, Robert N. Shorten, and Fiacre Ó. Cairbre. A global attractivity result for a class of switching discrete-time systems. In *ACC'07: Proceedings of the 2007 American Control Conference*, pages 3462–3463, New York, NY, USA, Jul. 2007. DOI: 10.1109/ACC.2007.4282444

- Yoon Song, M. Seetharama Gowda, and Gomatam Ravindran. On some properties of P-matrix sets. *Linear Algebra and its Applications*, 290(1–3):237–246, Mar. 1999.
DOI: 10.1016/S0024-3795(98)10228-8
- Angela Spence, Siebe Turksma, Ab Schelling, Thomas Benz, Jean-Pierre Medevielle, Ian McCrae, Juhani Jaaskelainen, and Eva Boethius. Methodologies for assessing the impact of ITS applications on CO₂ emissions. Technical report, EC-METI Task Force, Mar. 2009. <http://goo.gl/L0roR>
- Rade Stanojević and Robert N. Shorten. Fully decentralized emulation of best-effort and processor sharing queues. In *SIGMETRICS'08: Proceedings of the 2008 ACM International Conference on Measurement and Modeling of Computer Systems*, pages 383–394, Annapolis, MD, USA, Jun. 2008. ACM.
DOI: 10.1145/1375457.1375501
- Rade Stanojević and Robert N. Shorten. Load balancing vs. distributed rate limiting: An unifying framework for cloud control. In *ICC'09: Proceedings for the 2009 IEEE International Conference on Communications*, pages 1091–1096, Dresden, Germany, Jun. 2009a.
DOI: 10.1109/ICC.2009.5199141
- Rade Stanojević and Robert N. Shorten. Generalized distributed rate limiting. In *IWQoS'09: Proceedings of 17th IEEE International Workshop on International Workshop on Quality of Service*, pages 1–9, Charleston, SC, USA, Jul. 2009b.
DOI: 10.1109/IWQoS.2009.5201389
- Donald V. Steward. On an approach to techniques for the analysis of the structure of large systems of equations. *SIAM Review*, 4(4):321–342, Oct. 1962.
DOI: 10.1137/1004088
- Donald V. Steward. Partitioning and tearing systems of equations. *SIAM Journal: Series B, Numerical Analysis*, 2(2):345–365, 1965.
DOI: 10.1137/0702028
- Gilbert W. Stewart. *Introduction to matrix computations*. Computer science and applied mathematics. Academic Press, New York, NY, USA, 1973.
<http://books.google.ie/books?id=awHvAAAAMAAJ>
- Dusan M. Stipanović and Dragoslav D. Šiljak. Connective stability of discontinuous interconnected systems via parameter dependent Lyapunov functions. In *ACC'01: Proceedings of the 2001 American Control Conference*, volume 6, pages 4189–4196, Arlington, VA, USA, Jun. 2001.
DOI: 10.1109/ACC.2001.945633
- Xi-Ming Sun, Wei Wang, Guo-Ping Liu, and Jun Zhao. Stability analysis for linear switched systems with time-varying delay. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 38(2):528–533, Apr. 2008.
DOI: 10.1109/TSMCB.2007.912078
- Malur K. Sundareshan and Mathukumalli Vidyasagar. L_2 -stability of large-scale dynamical systems: Criteria via positive operator theory. *IEEE Transactions on Automatic Control*, 22(3):396–399, Jun. 1977.
DOI: 10.1109/TAC.1977.1101536
- Adarsha Swarnakar, Horacio Jose Marquez, and Tongwen Chen. Robust stabilization of nonlinear interconnected systems with application to an industrial utility boiler. *Control Engineering Practice*, 15(6):639–654, 2007.
DOI: 10.1016/j.conengprac.2006.11.004
- Hiroyuki Tamura and Tsuneo Yoshikawa. *Large-Scale Systems Control and Decision Making*. Electrical and Computer Engineering. Marcel Dekker, New York, NY, USA, 1990.
<http://books.google.ie/books?id=x6SeW76GvhQC>

- Wang Tang, Anthony N. Michel, and Harry W. Hale. On structure and stability of interconnected dynamical systems. *IEEE Transactions on Circuits and Systems*, 27(5):391–405, May 1980. DOI: 10.1109/TCS.1980.1084828
- Herbert G. Tanner, Ali Jadbabaie, and George J. Pappas. Stable flocking of mobile agents part I: Fixed topology. In *CDC'03: Proceedings of the 42nd IEEE Conference on Decision and Control*, volume 2, pages 2010–2015, Maui, HI, USA, Dec. 2003a. DOI: 10.1109/CDC.2003.1272910
- Herbert G. Tanner, Ali Jadbabaie, and George J. Pappas. Stable flocking of mobile agents part II: Dynamic topology. In *CDC'03: Proceedings of the 42nd IEEE Conference on Decision and Control*, volume 2, pages 2016–2021, Maui, HI, USA, Dec. 2003b. DOI: 10.1109/CDC.2003.1272911
- Henri Theil. *Linear Aggregation of Economic Relations*. North-Holland Publishing Co., Amsterdam, The Netherlands, 1954. <http://books.google.ie/books?id=yco6AAAAIAAJ>
- Wiley E. Thompson. Exponential stability of interconnected systems. *IEEE Transactions on Automatic Control*, 15(4):504–506, Aug. 1970. DOI: 10.1109/TAC.1970.1099507
- TMC Forum. What is Traffic Message Channel (TMC)? Website, Aug. 2007. <http://goo.gl/Xfnvx>
- H. Tokumaru, N. Adachi, and T. Amemiya. On the input-output stability of interconnected systems. *Syst. Ctrl. (J. Japan Ass. Automat. Contr. Eng.)*, 17:121–125, 1973.
- John N. Tsitsiklis. *Problems in Decentralized Decision Making and Computation*. Ph.D. thesis, Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Cambridge, MA, USA, Nov. 1984. <http://goo.gl/BUReu>
- John N. Tsitsiklis, Dimitri P. Bertsekas, and Michael Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, Sept. 1986. DOI: 10.1109/TAC.1986.1104412
- Peter Uetz, Loic Giot, Gerard Cagney, Traci A. Mansfield, Richard S. Judson, James R. Knight, Daniel Lockshon, Vaibhav Narayan, Maithreyan Srinivasan, Pascale Pochart, Alia Qureshi-Emili, Ying Li, Brian Godwin, Diana Conover, Theodore Kalbfleisch, Govindan Vijayadamodar, Meijia Yang, Mark Johnston, Stanley Fields, and Jonathan M. Rothberg. A comprehensive analysis of protein–protein interactions in *Saccharomyces cerevisiae*. *Nature*, 403:623–627, Feb. 2000. DOI: 10.1038/35001009
- Frank Uhlig. A recurring theorem about pairs of quadratic forms and extensions: A survey. *Linear Algebra and its Applications*, 25:219–237, Jun. 1979. DOI: 10.1016/0024-3795(79)90020-X
- Maria Elena Valcher and Paolo Santesso. Reachability properties of single-input continuous-time positive switched systems. *IEEE Transactions on Automatic Control*, 55(5):1117–1130, May 2010. DOI: 10.1109/TAC.2010.2045442
- Manuela Veloso, Peter Stone, and Kwun Han. The CMUnited-97 robotic soccer team: Perception and multi-agent control. *Robotics and Autonomous Systems*, 29(2–3):133–143, Nov. 2000. DOI: 10.1016/S0921-8890(99)00048-2
- Tamás Vicsek, András Czirók, Eshel Ben-Jacob, Inon Cohen, and Ofer Shochet. Novel type of phase transition in a system of self-driven particles. *Physical Review Letters*, 75(6):1226–1229, Aug. 1995. DOI: 10.1103/PhysRevLett.75.1226

- Mathukumalli Vidyasagar. L_2 -stability of interconnected systems using a reformulation of the passivity theorem. *IEEE Transactions on Circuits and Systems*, 24(11):637–645, Nov. 1977.
DOI: 10.1109/TCS.1977.1084285
- Mathukumalli Vidyasagar. New passivity-type criteria for large-scale interconnected systems. *IEEE Transactions on Automatic Control*, 24(4):575–579, Aug. 1979.
DOI: 10.1109/TAC.1979.1102092
- Mathukumalli Vidyasagar. On the well-posedness of large-scale interconnected systems. *IEEE Transactions on Automatic Control*, 25(3):413–421, Jun. 1980. DOI: 10.1109/TAC.1980.1102345
- Mathukumalli Vidyasagar. *Input-Output Analysis of Large-Scale Interconnected Systems: Decomposition, Well-Posedness and Stability*, volume 29 of *Lecture notes in control and information sciences*. Springer-Verlag New York, Inc., New York, NY, USA, 1981.
<http://books.google.ie/books?id=Ct2nAAAAIAAJ>
- Mathukumalli Vidyasagar. *Nonlinear systems analysis*. Number 42 in Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, second edition, 2002.
http://books.google.ie/books?id=_JLrm1rRRUIC
- Elena Virnik. *Analysis of Positive Descriptor Systems*. Topics in Systems and Control Theory. VDM Verlag, Saarbrücken, Germany, 2008.
<http://books.google.ie/books?id=2M6TNwAACAAJ>
- Antonio Visioli. *Practical PID Control*. Advances in Industrial Control. Springer-Verlag, London, UK, 2006.
DOI: 10.1007/1-84628-586-0
- Muqiu Wang, Lian Wang, and Xuetang Du. On the decomposition problem of stability for Volterra integrodifferential equations. *Acta Mathematicae Applicatae Sinica (English Series)*, 8(1):82–96, 1992.
DOI: 10.1007/BF02006075
- Paul K. C. Wang. Navigation strategies for multiple autonomous mobile robots moving in formation. *Journal of Robotic Systems*, 8(2):177–195, Apr. 1991.
DOI: 10.1002/rob.4620080204
- Roger Wattenhofer, Li Li, Paramvir Bahl, and Yi-Min Wang. Distributed topology control for wireless multihop ad-hoc networks. In *INFOCOM'01: Proceedings of the 20th Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 3, pages 1388–1397, Anchorage, AK, USA, Apr. 2001.
DOI: 10.1109/INFCOM.2001.916634
- Stein Weissenberger. Piecewise-quadratic and piecewise-linear Lyapunov functions for discontinuous systems. *International Journal of Control*, 10(2):171–180, Aug. 1969.
DOI: 10.1080/00207176908905814
- Susan C. Weller and N. Clay Mann. Assessing rater performance without a “gold standard” using consensus theory. *Medical Decision Making*, 17(1):71–79, Feb. 1997.
DOI: 10.1177/0272989X9701700108
- Jan C. Willems. Dissipative dynamical systems part I: General theory. *Archive for Rational Mechanics and Analysis*, 45(5):321–351, 1972.
<http://goo.gl/VxS1t>
- Jan C. Willems. Qualitative behavior of interconnected systems. *Annals of Systems Research*, 3: 61–80, 73.

- Robert L. Winkler. The consensus of subjective probability distributions. *Management Science*, 15(2):B61–B75, Oct. 1968. DOI: 10.1287/mnsc.15.2.B61
- Fabian Wirth. A converse Lyapunov theorem for linear parameter-varying and linear switching systems. *SIAM Journal on Control and Optimization*, 44(1):210–239, 2005a. DOI: 10.1137/S0363012903434790
- Fabian Wirth. A converse Lyapunov theorem for switched linear systems with dwell times. In *CDC-ECC '05: Proceedings of the Joint 44th IEEE Conference on Decision and Control and the 2005 European Control Conference*, pages 4572–4577, Seville, Spain, Dec. 2005b. DOI: 10.1109/CDC.2005.1582883
- Bugong Xu. On delay-independent stability of large-scale systems with time delays. *IEEE Transactions on Automatic Control*, 40(5):930–933, May 1995. DOI: 10.1109/9.384233
- Peng Yang, Randy A. Freeman, Geoffrey J. Gordon, Kevin M. Lynch, and Siddhartha S. Srinivasa. Decentralized estimation and control of graph connectivity in mobile sensor networks. In *ACC'08: Proceedings of the 2008 American Control Conference*, pages 2678–2683, Seattle, WA, USA, Jun. 2008. DOI: 10.1109/ACC.2008.4586897
- Xiao-Song Yang and Guanrong Chen. Limit cycles and chaotic invariant sets in autonomous hybrid planar systems. *Nonlinear Analysis: Hybrid Systems*, 2(3):952–957, Aug. 2008. Special Issue Section: Analysis and Design of Hybrid Systems, Analysis and Design of Hybrid Systems. DOI: 10.1016/j.nahs.2008.03.004
- Hui Ye, A.N. Michel, and Ling Hou. Stability theory for hybrid dynamical systems. *IEEE Transactions on Automatic Control*, 43(4):461–474, Apr. 1998. DOI: 10.1109/9.664149
- Rama K. Yedavalli and Andrew Sparks. Ultimate boundedness control of linear switched systems using controlled dwell time approach. In *CDC'01: Proceedings of the 40th IEEE Conference on Decision and Control*, volume 3, pages 2490–2495, Orlando, FL, USA, Dec. 2001. DOI: 10.1109/.2001.980637
- Christos A. Yfoulis and Robert N. Shorten. A numerical technique for the stability analysis of linear switched systems. *International Journal of Control*, 77(11):1019–1039, Jul. 2004. DOI: 10.1080/002071704200026963
- Xu Yong, Li Jie, Tang Wansheng, Zhang Jianxiong, and Wei Jie. Stability analysis of piecewise linear Delta operator systems. In *ICAL'08: Proceedings of the 2008 IEEE International Conference on Automation and Logistics*, pages 1688–1692, Qingdao, China, Sept. 2008. DOI: 10.1109/ICAL.2008.4636426
- David M. Young. *Iterative Solution of Large Linear Systems*. Academic Press, New York, NY, USA, 1971. <http://books.google.ie/books?id=TWoboQnbZtsC>
- George Zames. On the input-output stability of time-varying nonlinear feedback systems—Part I: Conditions derived using concepts of loop gain, conicity, and positivity. *IEEE Transactions on Automatic Control*, 11(2):228–238, Apr. 1966. DOI: 10.1109/TAC.1966.1098316
- Jens Zander. Performance of optimum transmitter power control in cellular radio systems. *IEEE Transactions on Vehicular Technology*, 41(1):57–62, Feb. 1992. DOI: 10.1109/25.120145

- Annalisa Zappavigna, Themistoklis Charalambous, and Florian Knorn. Unconditional stability of the Foschini-Miljanic algorithm. To appear in *Automatica*, Mar. 2011. <http://goo.gl/o8rKm>
- Aleksandar I. Zečević and Dragoslav D. Šiljak. A block-parallel newton method via overlapping epsilon decompositions. *SIAM Journal on Matrix Analysis and Applications*, 15(3):824–844, Jul. 1994. DOI: 10.1137/S0895479892229115
- Aleksandar I. Zečević and Dragoslav D. Šiljak. A new approach to control design with overlapping information structure constraints. *Automatica*, 41(2):265–272, Feb. 2005a. DOI: 10.1016/j.automatica.2004.09.011
- Aleksandar I. Zečević and Dragoslav D. Šiljak. A decomposition-based control strategy for large, sparse dynamic systems. *Mathematical Problems in Engineering*, 2005(1):33–48, 2005b. DOI: 10.1155/MPE.2005.33
- Aleksandar I. Zečević and Dragoslav D. Šiljak. *Control of Large-Scale Systems Under Information Structure Constraints*. Springer-Verlag New York, Inc., New York, NY, USA, 2010. <http://books.google.ie/books?id=vqweTUAnjwYC>
- Guisheng Zhai and Hai Lin. Controller failure time analysis for symmetric \mathcal{H}_∞ control systems. *International Journal of Control*, 77(6):598–605, 2004. DOI: 10.1080/00207170410001703232
- Guisheng Zhai, Bo Hu, Kazunori Yasuda, and Anthony N. Michel. Disturbance attenuation properties of time-controlled switched systems. *Journal of the Franklin Institute*, 338(7):765–779, Nov. 2001. DOI: 10.1016/S0016-0032(01)00030-8
- Guisheng Zhai, Bo Hu, Kazunori Yasuda, and Anthony N. Michel. Qualitative analysis of discrete-time switched systems. In *ACC'02: Proceedings of the 2002 American Control Conference*, volume 3, pages 1880–1885, Anchorage, AK, USA, May 2002. DOI: 10.1109/ACC.2002.1023907
- Guisheng Zhai, Xuping Xu, Hai Lin, and Anthony N. Michel. Analysis and design of switched normal systems. *Nonlinear Analysis: Theory, Methods & Applications*, 65(12):2248–2259, Dec. 2006. DOI: 10.1016/j.na.2006.01.034
- Wei Zhang, Alessandro Abate, Jianghai Hu, and Michael P. Vitus. Exponential stabilization of discrete-time switched linear systems. *Automatica*, 45(11):2526–2536, Nov. 2009a. DOI: 10.1016/j.automatica.2009.07.018
- Wei Zhang, Alessandro Abate, Michael P. Vitus, and Jianghai Hu. On piecewise quadratic control-Lyapunov functions for switched linear systems. In *CDC-CCC'09: Proceedings of the Joint 48th IEEE Conference on Decision and Control and 28th Chinese Control Conference*, pages 1088–1093, Dec. 2009b. DOI: 10.1109/CDC.2009.5400642