

Lanczos Tau Methods for Time-Delay Systems

Approximation of System and its H^2 -norm

Evert Provoost

Dissertation presented in partial fulfilment
of the requirements for the degree of
Doctor of Engineering Science (PhD):
Computer Science

Supervisor:
Prof. dr. ir. W. Michiels

October 2026

LANCZOS TAU METHODS FOR TIME-DELAY SYSTEMS

APPROXIMATION OF SYSTEM AND ITS H^2 -NORM

Evert PROVOOST

Supervisor:

Prof. dr. ir. W. Michiels

Members of the

Examination Committee:

Em. prof. dr. ir. O. Van der Biest

Prof. dr. ir. J. Suykens

Prof. dr. R. Vandebril

Dr. F. Scarabel

(University of Leeds)

Prof. dr. J. Sieber

(University of Exeter)

Dissertation presented in partial
fulfilment of the requirements for
the degree of Doctor of Engineering
Science (PhD): Computer Science

October 2026

© 2026 Evert Provoost

Uitgegeven in eigen beheer, Evert Provoost, Liersesteenweg 174, 2590 Berlaar (Belgium)

Alle rechten voorbehouden. Niets uit deze uitgave mag worden vermenigvuldigd en/of openbaar gemaakt worden door middel van druk, fotokopie, microfilm, elektronisch of op welke andere wijze ook zonder voorafgaande schriftelijke toestemming van de uitgever.

All rights reserved. No part of the publication may be reproduced in any form by print, photoprint, microfilm, electronic or any other means without written permission from the publisher.

Abstract

In various modelling and control tasks, one is often faced with underlying transport phenomena for which it is often unnecessary, if not infeasible, to model them in their entirety. In many cases, the right middle ground, and the primary topic of this text, is to model the transport as a simple delay in time resulting in a so-called time-delay system.

We will further develop a number of computational aspects of linear, time-invariant, time-delay systems. Due to a time-delay system's dependence on the past, its true state is given by a function, i.e. is infinite dimensional. As computers are inherently finite in space and time, a discretization method is generally needed for their automated analysis. A first aim is to further develop the understanding of one such discretization scheme: the Lanczos tau method of Ito and Teglas (1986, *SIAM J. Sci. Comput.*, 24/4: 737–59). A second aim is to approximate the H^2 -norm of these systems, a measure used to quantify disturbance rejection in robust control and as an error measure in model approximation.

More concretely, we introduce an operator based formulation of the Lanczos tau method, which simplifies theoretical derivations. We use this formulation to show connections to other discretization schemes and to derive a highly sparse matrix realization. In the frequency domain we prove a connection to rational approximation, which allows us, among other things, to answer questions about the preservation of stability and convergence of the H^2 -norm.

Next, we take a closer look at a method to approximate the H^2 -norm of a time-delay system introduced by Vanbiervliet et al. (2011, *Int. J. Control*, 84/2: 228–41). In this approach we use a system's discretization as proxy, i.e. we use the exact H^2 -norm of the discretized system, which we know how to compute, as approximation of the norm of the original. As mentioned, we prove this approach to converge. Furthermore, we demonstrate that the choice of discretization has a large impact on the convergence rate of the H^2 -norm. In particular, we illustrate that a good choice attains geometric convergence for single delay systems instead of the algebraic rate that is assumed in the existing literature. Whilst we have no proof of these rates, we do provide a number of clues and intuitions as to why this might happen. Additionally, these clues hint at

the use of splines instead of a single polynomial for systems with multiple discrete delays. We illustrate that this use of splines indeed improves convergence rates for such systems.

A further extension upon the literature, is an adaption of this method to delay differential algebraic systems, which is a significantly larger model class. Notably, such models allow for so-called neutral type systems and systems with feedthrough, which make approximating the H^2 -norm much harder. Under suitable conditions we prove that discretization does not worsen these issues and that convergence is maintained. Finally, we derive efficient formulas for the gradient of the square of the H^2 -norm with respect to system matrices and delays, which allows us to design H^2 -optimal controllers and approximate models.

Beknopte samenvatting

Veel problemen in wiskundige modellering en regeltechniek omvatten een onderliggend transportverschijnsel waarvan het vaak onnodig, zo niet onmogelijk, is om dit in zijn volledigheid te modelleren. In veel gevallen is de juiste middenweg om dit verschijnsel te benaderen als een afhankelijkheid van de toestand op een eerder tijdstip.

In deze thesis ontwikkelen we enkele rekenkundige aspecten van lineaire, tijdsinvariante systemen met tijdsvertraging. Gezien zulk een systeem afhankelijk is van zijn verleden, is zijn toestand gegeven door een functie en is het dus oneindigdimensionaal. Omdat computers fundamenteel eindig zijn in tijd en ruimte, is een discretisatiemethode doorgaans nodig voor de geautomatiseerde analyse van zulk een systeem. Een eerste doel is om de kennis over één zulke discretisatie uit te breiden: de Lanczos-tau-methode van Ito en Teglas (1986, *SIAM J. Sci. Comput.*, 24/4: 737–59). Een tweede doel is om de H^2 -norm van zulke systemen te benaderen. Deze norm wordt gebruikt om de gevoeligheid aan ruis te kwantificeren in robuuste regeltechniek en als een foutmaat in modelbenadering.

Specifiek introduceren we een formulering van de Lanczos-tau-methode gebaseerd op operatoren, wat theoretische afleidingen vereenvoudigt. Met behulp van deze formulering tonen we connecties aan met andere discretisatieschema's en leiden we een zeer ijle matrixvorm af. In het frequentiedomein bewijzen we een overeenkomst met rationale benadering, wat ons toelaat om, onder andere, vragen te beantwoorden over het behoud van stabiliteit en de convergentie van de H^2 -norm.

Vervolgens bestuderen we de methode van Vanbiervliet et al. (2011, *Int. J. Control*, 84/2: 228–41) om de H^2 -norm van een systeem met tijdsvertraging te benaderen. In hun aanpak gebruiken we de exacte H^2 -norm van een discretisatie, waarvoor efficiënte methoden bestaan, als benadering voor de norm van het oorspronkelijke systeem. Zoals eerder vermeld tonen we aan dat deze methode convergeert. Verder demonstreren we dat de keuze van discretisatie een grote invloed heeft op de convergentiesnelheid. In het bijzonder zien we dat een goede keuze geometrische convergentie geeft voor een systeem met een enkele tijdsvertraging in tegenstelling tot de derde orde algebraïsche convergentie die men veronderstelt in de literatuur. Hoewel we dit niet kunnen bewijzen,

bieden we enkele vermoedens en intuïties voor waarom dit het geval zou kunnen zijn. Deze doen verder vermoeden dat het gebruik van een spline, in plaats van een enkele veelterm, voordelen biedt voor systemen met meerdere tijdsvertragingen. We tonen aan dat zulke splines wel degelijk de convergentiesnelheid voor deze systemen verbetert.

Een verdere uitbreiding op de literatuur is een variatie op deze methode voor systemen met zowel tijdsvertraging als algebraïsche vergelijkingen. Dit vormt een significant grotere klasse van systemen en omvat ook zogenaamde neutrale systemen en systemen met een directe verbinding tussen in- en uitgang. Beide maken het benaderen van de H^2 -norm aanzienlijk ingewikkelder. Onder redelijke voorwaarden kunnen we bewijzen dat een Lanczos-tau-methode deze situaties niet verergert en dat convergentie behouden blijft. Finaal leiden we efficiënte formules af voor de gradiënt van het kwadraat van de H^2 -norm als functie van de systeemmatrices en tijdsvertragingen. Deze laten ons toe om H^2 -optimale regelaars en benaderende modellen te ontwerpen.

Preface

This thesis is the result of my doctoral programme, an endeavour one, luckily, does not face alone. I would thus like to thank a number of people who made mine a richer and more memorable experience.

There is some friendly rivalry between PhD students about who has the best supervisor. Rest assured, one would be hard pressed to find a metric in which the research team of mine does not prevail. Whilst the official terminology is supervisor, Wim Michiels is a proper mentor, caring not only for his team's research output, but also for their well-being and future careers. He rarely acts like a boss, but extends the trust and support of a patron. Still very much an active and well-regarded researcher, he contributes to both the content and writing of articles as a true coauthor. For me specifically, I am tremendously grateful for not only allowing me, but also supporting me in taking up roles and engagements somewhat outside the doctoral programme and trusting that I would be able to combine both (even if I must admit, I doubted that myself at times).

I would also like to extend my gratitude to the members of my jury: Omer Van der Biest, Johan Suykens, Raf Vandebril, Francesca Scarabel, and Jan Sieber, for diligently reading this work and providing valuable feedback. Additionally, I want to specifically thank Fancesca Scarabel and Jan Sieber for having made the long trip from the United Kingdom to Leuven.

Next to a supervisor, a research team also consists of a tightly knit group of PhD students and post-docs. I would like to thank, in order of appearance, Pieter, Adrian, Sanaz, Kang-Kang, Victor, Diego, Jasper, and Pournami for the many great chats about research, history, philosophy, ..., and their delightful company during the weekly team lunch.

I have had the opportunity to represent my colleagues at the university level during almost two years. Being this closely involved with the inner workings of an organization as complex as KU Leuven is a fascinating, if not a little daunting, and definitely exhausting, experience. Luckily, I could count and trust on the experience of Charlotte and Jana (my two colleague representatives) and the input of the wide network of

representatives throughout the university. I would also like to thank the other members of these bodies for giving the appropriate weight to our words and their dedication to making this institute the best it can be.

Whilst research can be exciting, this excitement sometimes gets stumped (often somewhere early in the third year) by a proof that just won't quite work or a problem for which you seem to have exhausted all strategies. When this happens, it is good to get out of the usual environment and work on a (slightly) different project. I am thus tremendously grateful to Francesca Scarabel for hosting me at the University of Leeds during August of last year and providing such a project.

Outside of research, teaching, and service there must of course also be some entertainment. I have had the pleasure to work with Arne, Astrid, Toon, Leonie, and Filippo as part of the oPPC (official Party Planning Committee) of NUMA. I would also like to specifically thank Vince as co-organizer and co-creator of the Wuitschoven (mock) conference. Committee planned parties, conferences, and coffee trains (for those unfamiliar with the phenomenon: *de rigueur* coffee breaks) are of course no fun without fun-loving attendees, and I am happy to have been part of a group like NUMA that has such people in abundance.

Throughout the past four years a number of these people I have seen more often than anyone else. A warm thank you to my (former) office mates Nick, Nikhil, Yiqing, Marlies, and Leonie for the chats about research ideas, jokes, drawings, reminders of delicacies in the cafeteria, ..., and being each other's support group through the trials of the PhD.

Niet enkel doorheen het doctoraat, maar ook de jaren ervoor mocht ik altijd op de steun, het advies en de nodige ludieke afleiding van vrienden en familie rekenen. Te veel namen om op te noemen, maar jullie weten wie jullie zijn en ik ben jullie voor dat alles enorm dankbaar.

Als laatste zijn er nog mijn ouders. Bedankt voor de nagenoeg onuitputtelijke bron van warmte, begrip en steun die jullie al bijna 27 jaar vormen. Zonder had dit niet mogelijk geweest.

Evert Provoost
Leuven, September 2026

Contents

Preface	v
Contents	vii
List of figures	ix
List of abbreviations	xi
List of symbols	xiii
1 Introduction	1
1.1 Overview of chapters	2
1.2 Visualization of complex functions	4
2 Time-delay systems	7
2.1 Time-delay systems of retarded type	8
2.2 Exponential stability and control	12
2.3 Elements of C_0 -semigroup theory	14
2.4 Neutral systems and algebraic constraints	16
2.5 The H^∞ - and H^2 -norm	19
2.6 Summary	22
3 The Lanczos tau method	25
3.1 A short history of discretization techniques	26
3.2 A review of orthogonal polynomials	28
3.3 The Lanczos tau method	30
3.4 Implicit rational approximation	32
3.5 Connection to pseudospectral collocation	36
3.6 Sparse, nested discretizations	39
3.7 Relation to Padé approximation	44
3.8 Preservation of stability	48
3.9 Summary	49

4	Computing the H^2-norm of a single delay system	53
4.1	Proof of convergence	54
4.2	The H^2 -norm of a delay-free system	57
4.3	An unexpected improvement to geometric rate	58
4.4	Connection to the delay-Lyapunov equations	62
4.5	A case with super convergence	64
4.6	Distributed and multiple delays	67
4.7	Summary	69
5	Computing the H^2-norm of a DDAE	71
5.1	Finiteness of the strong H^2 -norm of a DDAE	73
5.2	Lanczos tau methods for DDAEs	76
5.3	Convergence	80
5.4	Summary	87
6	Accelerating H^2-norm convergence using splines	89
6.1	A spline based Lanczos tau method	90
6.2	Properties in the frequency domain	91
6.3	Convergence of the H^2 -norm	94
6.4	Further acceleration using additional knots	97
6.5	Summary	100
7	H^2-optimal synthesis	101
7.1	Computing the gradient	102
7.2	Numerical examples	106
7.3	Summary	109
8	Concluding remarks	111
	Acknowledgements	115
	Code availability	117
	Curriculum vitae	119
	Bibliography	123

List of figures

1.1	Domain-colouring of $z \mapsto e^{-z}$	5
1.2	Domain-colouring of $z \mapsto \frac{z-1}{(z+1)^2}$	5
1.3	Comparison of the typical colour map used to represent phase and the one used in this text.	6
2.1	Illustration of the history function.	9
2.2	Characteristic roots of a simple example RDDE.	11
2.3	Characteristic roots of a simple example NDDE.	16
2.4	Example of an NDDE that is sensitive to infinitesimal delay perturbations.	18
2.5	The combined system used in model approximation.	22
3.1	Domain-colourings of the rational approximation of the exponential underlying a Lanczos tau method based on Legendre polynomials.	40
3.2	Comparison of the true roots of a time-delay system of retarded type to those obtained from two low degree Lanczos tau approximations.	40
3.3	Convergence rate of the rightmost root of fig. 3.2.	41
3.4	The non-zeroes of the Lanczos tau discretization using an ultraspherical matrix realization based on Chebyshev polynomials.	42
3.5	The non-zeroes of the Lanczos tau discretization using an ultraspherical matrix realization based on Legendre and the corresponding generalized Gegenbauer polynomials.	43
3.6	Domain-colouring of a rational approximation of the exponential with poles on the right half-plane.	51
4.1	Comparison of the convergence of the H^2 -norm of a single delay system for pseudospectral collocation in Chebyshev extremal nodes and a symmetric Lanczos tau method.	59
4.2	Comparison of the convergence of the H^2 -norm of a single delay system for a nonsymmetric Lanczos tau method and a symmetric one.	59

4.3	Comparison of the convergence of the H^2 -norm of a single delay system for a Lanczos tau method using a Chebyshev basis and one using a Legendre basis.	61
4.4	Domain-colouring of the rational approximant of the exponential function underlying a Lanczos tau discretization based on Chebyshev polynomials of the second kind.	61
4.5	Scalar example displaying super convergence of the H^2 -norm when using a symmetric Lanczos tau method.	65
4.6	The transfer function of the system of fig. 4.5 and its first order delay-free approximation.	65
4.7	Convergence rates of the H^2 -norm for an RDDE with distributed delay an one with two discrete delays.	68
5.1	Convergence rates of the H^2 -norm for RDDEs with one and multiple delays.	86
5.2	Convergence rates of the H^2 -norm for NDDEs with one and multiple delays.	87
5.3	Illustration of the effect of the position of neutral terms on the convergence rate of the H^2 -norm of systems with multiple delays.	88
6.1	Domain-colourings of the rational approximations of the exponential underlying a spline Lanczos tau method based on Legendre bases.	92
6.2	Similar illustration as fig. 6.1, but for equidistant delays. . . .	92
6.3	Comparison of convergence rates of polynomial and spline based discretizations for and RDDE.	95
6.4	Comparison of convergence rates of polynomial and spline based discretizations for an NDDE.	96
6.5	The experiment of fig. 5.3 but using a spline discretization. . .	96
6.6	The experiments of figs. 6.3 and 6.4 with additional knots at the first set of delay combinations.	98
6.7	The same experiment as fig. 6.5 but with equidistant delays. . .	98

List of abbreviations

ODE	Ordinary Differential Equation
PDE	Partial Differential Equation
FDE	Functional Differential Equation
DDE	Delay Differential Equation
RDDE	Retarded Delay Differential Equation
NDDE	Neutral Delay Differential Equation
DAE	Differential Algebraic Equation
DDAE	Delay Differential Algebraic Equation

List of symbols

SETS

$\mathbb{N}$	set of natural numbers, i.e. $\{0, 1, 2, \dots\}$
$\mathbb{Z}$	set of integers
$\mathbb{R}$	set of real numbers
$\mathbb{R}_+$	set of nonnegative real numbers, i.e. $\{x \in \mathbb{R}: x \geq 0\}$
$\mathbb{C}$	set of complex numbers
$\mathbb{C}_{>0}$	open right half-plane, i.e. $\{z \in \mathbb{C}: \operatorname{Re}(z) > 0\}$
$\mathbb{C}_{\geq 0}$	closed right half-plane, i.e. $\{z \in \mathbb{C}: \operatorname{Re}(z) \geq 0\}$
$\mathbb{P}_N$	set of complex valued polynomials of degree at most N
A^n	natural n dimensional extension of the set A
$\mathbb{D}$	unit disc, i.e. $\{z \in \mathbb{C}: z < 1\}$
$B_\delta(\boldsymbol{\tau})$	open ball around the vector $\boldsymbol{\tau}$ of radius δ

LINEAR ALGEBRA

Generally, scalars are denoted by italic lowercase letters, vectors by boldface lowercase letters, and matrices by uppercase letters.

A^T	transpose of matrix A
A^*	conjugate transpose of matrix A
$\operatorname{tr}(A)$	trace of matrix A
$\det(A)$	determinant of matrix A
$\ker(A)$	kernel of matrix A
$\operatorname{span}(A)$	span of matrix A
$\rho(A)$	spectral radius of matrix A
$A \otimes B$	Kronecker product of matrices A and B
I_n	the $n \times n$ identity matrix
$\mathbf{0}_{n \times m}$	the $n \times m$ zero matrix

NORMS AND SPACES

In what follows $\mathbb{K}$ can stand for either $\mathbb{R}$ or $\mathbb{C}$, p is a positive real number or infinity, and k a nonnegative integer.

$\ \varphi\ $	norm of φ (specified by context)
$\langle \varphi, \psi \rangle$	inner product of φ and ψ (specified by context)
$\ \varphi\ _2$	Euclidean norm of φ
$\ \varphi\ _\infty$	supremum norm of φ
$\ A\ _p$	induced p -norm of matrix A
$\ A\ _F$	Frobenius norm of matrix A
H^p	Hardy p space of (matrix-valued) functions on $\mathbb{C}_{>0}$
$\ G\ _{H^p}$	H^p -norm of (matrix-valued) function G
$\ G\ _{H^p}^s$	strong H^p -norm of (matrix-valued) function G
$C([a, b]; \mathbb{K}^n)$	space of continuous functions
$C^k([a, b]; \mathbb{K}^n)$	space of k times continuously differentiable functions
$L^p([a, b]; \mathbb{K}^n)$	Lebesgue p space
$W^{k,p}([a, b]; \mathbb{K}^n)$	Sobolev k, p space

OPERATORS AND FUNCTIONALS

The relevant space or basis, if needed, is always specified by context.

I	identity operator
ε_θ	evaluation functional in θ
$\mathcal{D}$	derivative operator
$\mathcal{T}_{N-1}$	truncation to a degree $N - 1$ polynomial
$\mathbf{0}$	constant map to the zero element
$[A]$	coordinate expression of $\mathcal{A}$

SYSTEMS

$\mathbf{x}(t)$	internal state at time t , a complex vector of length n
$\mathbf{u}(t)$	input at time t , a complex vector of length p
$\mathbf{y}(t)$	output at time t , a complex vector of length q
$\dot{\varphi}$	time derivative of φ
τ_m	maximal delay (also denoted by τ if there is only one)
X	state space $C([-\tau_m, 0]; \mathbb{C}^n)$
Z	state space $\mathbb{C}^n \times L^2([-\tau_m, 0]; \mathbb{C}^n)$

ORTHOGONAL POLYNOMIALS

Throughout this text we use classical orthogonal polynomials shifted to the interval between the maximal delay and zero, i.e. $[-\tau_m, 0]$. To lighten notation, we shall denote these shifted polynomials by their usual names in the literature for the interval $[-1, 1]$, unless otherwise noted.

$P_k^{(\alpha, \beta)}$	shifted k th Jacobi polynomial with parameters α and β
C_k^λ	shifted k th Gegenbauer polynomial with parameter λ
P_k	shifted k th Legendre polynomial
T_k	shifted k th Chebyshev polynomial of the first kind
U_k	shifted k th Chebyshev polynomial of the second kind

A review of these polynomials can be found in section 3.2.

MISCELLANEOUS

$O(g)$	Landau's big O , i.e. functions of order (at most) g
$o(g)$	Landau's little o , i.e. functions dominated by g
${}_1F_1(a; b; z)$	confluent hypergeometric function
δ_{jk}	Kronecker delta

Essentially, all models are wrong,
but some are useful.

George E. P. Box

1 *Introduction*

In various modelling and control tasks, one is often faced with underlying transport phenomena. Be it water flowing through a hose, the teeth of a milling tool rotating towards the material about to be removed, the prepubescent development of an individual within a population, the spread of a virus within a host during an outbreak of the flu, or even the strips of ocean and atmosphere involved in the El Niño–Southern Oscillation (see e.g. Erneux 2009; Glass et al. 2021; Keane et al. 2017; Munoa et al. 2016; Orosz and Molnár 2025; Sipahi et al. 2011).

Something all these examples share is that is often unnecessary, if not infeasible, to model this transport phenomenon in its entirety. In many cases, the right middle ground, and the primary topic of this text, is to model the transport as a simple delay in time. Thus, instead of solving the Navier–Stokes equations within the hose, we simply assume the fluid to take a set number of seconds to move from the one end to the other. Instead of modelling the spread and multiplication of the influenza virus within the body, we assume there to be a typical latent period, or distribution thereof, between infection and becoming infectious.

In this text we have two main aims. Due to a time-delay system’s dependence on the past, its true state is given by a function, i.e. is infinite dimensional. As computers are inherently finite in space and time, a discretization method is generally needed for their automated analysis. A first aim is to further develop the understanding of one such discretization scheme: the Lanczos tau method of Ito and Teglas (1986). A second aim is to approximate the H^2 -norm of these systems, a measure used to quantify disturbance rejection in robust control and as an error measure in model approximation. Naturally, our developments of the Lanczos tau method will feature in the latter.

More concretely, we will introduce an operator based formulation of the Lanczos tau method, which, as we shall see, simplifies theoretical derivations. We will use this formulation to show connections to other discretization schemes and to derive a highly sparse matrix realization. In the frequency domain we shall prove a connection to rational approximation, which will allow us, among other things, to answer questions about the preservation of stability and convergence of the H^2 -norm.

This naturally leads us back to our second aim. We will take a closer look at a method to approximate the H^2 -norm of a time-delay system introduced by Vanbiervliet, Michiels, et al. (2011). In this approach we use a system's discretization as proxy, i.e. we use the H^2 -norm of the discretized system, which we know how to compute, as approximation of the norm of the original. As mentioned, we will prove this approach to converge, to our knowledge, for the first time. Furthermore, we will see that the choice of discretization has a large impact on the convergence rate of the H^2 -norm. In particular, we illustrate that a good choice attains geometric convergence for single delay systems instead of the algebraic rate that is assumed in the existing literature. Whilst we have no proof of these rates, we do provide a number of clues and intuitions as to why this might happen. Additionally, these clues hint at the use of splines instead of a single polynomial for systems with multiple discrete delays. We shall see that this use of splines indeed improves convergence rates for such systems.

Another way in which we extend upon the literature, is by showing how to apply this method to delay differential algebraic systems, which is a significantly larger model class. Notably, such models allow for so-called neutral type systems and systems with feedthrough, which make approximating the H^2 -norm much harder. Under suitable conditions we prove that discretization does not worsen these issues and that convergence is maintained. Finally, we derive efficient formulas for the gradient of the square of the H^2 -norm with respect to system matrices and delays, which allows us to design H^2 -optimal controllers and approximate models.

Throughout, we will limit ourselves to the case of linear, time-invariant systems. The astute reader might object that this means our results are not, or at least not directly, applicable to some of our motivating examples. Epidemiological models, for one, are almost universally nonlinear. However, as we shall see, the classes we do consider are already rather rich. Furthermore, many analysis tasks of nonlinear systems can be solved, at least partially, by applying the linear techniques locally, i.e. to a linearization of the system around the point of interest. Hence, the reach of the presented results is somewhat wider than this focus might suggest.

1.1 OVERVIEW OF CHAPTERS

In chapter 2 we specify the classes of time-delay systems we will study in more detail and introduce some elementary results known about such systems and their control. In particular, we shall see how the introduction of delay makes these systems infinite dimensional, which, as mentioned, is the source of most complexity in their analysis. We show how we can

attain a very broad class of systems from a rather compact formulation if we allow algebraic equations, leading to so-called delay differential algebraic equations (DDAEs). Additionally, we will introduce the H^2 -norm with its primary applications and how it fits within the broader picture of system norms.

With chapter 3 we move our focus to computation. We introduce an operator based description of the Lanczos tau method and discuss how it can equivalently be interpreted as rational approximation in the frequency domain; a connection that will be useful to understand the effects seen in later chapters. We show how Lanczos tau methods relate to other discretization schemes, which will allow us to reuse some preexisting results on the convergence of the characteristic roots. Next, we present one of the biggest advantages of Lanczos tau methods from a computational perspective, namely that they make it easy to construct sparse and numerically beneficial matrix realizations. Finally, we return to the frequency domain, where, for the specific Lanczos tau method using Legendre polynomials, a link with Padé approximation allows us to prove preservation of stability.

In the remaining chapters we apply the Lanczos tau discretization to the approximation of the H^2 -norm, where we shall use the exact H^2 -norm of the discretization as an approximation of the norm of the original system. In chapter 4, we start with the simpler case of a system of retarded type with only a single delay. We prove, under mild assumptions, for the first time that this method converges and show that these conditions hold for a method using Legendre polynomials. We shall see a somewhat unexpected improvement in convergence rate, in light of the existing literature on this approach, if the discretization is chosen well. Whilst a full understanding of this accelerated convergence remains an open question, we provide two interpretations that partially explain why such an improvement might occur. One is through the magnitude of the underlying rational approximation of the exponential on the imaginary axis, the other is through an interesting link to the delay-Lyapunov equation. Finally, we are able to use these properties to prove a case of super-convergence, i.e. an exact result that is attained at a finite discretization degree.

In chapter 5 we extend this approach to multiple discrete delays and neutral systems using the aforementioned DDAE formulation. As we shall see, this class of systems can be sensitive to infinitesimal delay perturbations, requiring us to use a robust system norm, the so-called strong H^2 -norm. We show that, within this robust framework, discretization does not introduce feedthrough—something that would render the H^2 -norm infinite. Using a slightly extended list of assumptions compared to the single delay case, we show the method to converge and we prove that

for a single delay system—that can now include a neutral term—these assumptions are again guaranteed for the discretization using Legendre polynomials.

In the final two main chapters we extend upon this method in two ways. In chapter 6 we see how using a spline approximation with the knots at the delays and combinations thereof—instead of the single polynomial covering the entire delay interval of the preceding chapters—allows us to recover some of the accelerated convergence of chapter 4 in the multiple delay case. In chapter 7 we derive a method to efficiently compute the gradient of the H^2 -norm with respect system parameters and delays and illustrate how this can be used for H^2 -norm optimal synthesis of controllers and stable reduced-order models.

To guide the reader, each chapter comes with an overview paragraph in its introduction and an extensive summary section at the end. Finally, we provide some broader concluding remarks and possible avenues for future research in chapter 8.

As mentioned, a lot of our developments will relate to rational approximation in the frequency domain. In this setting, it is often useful to visualize these approximations on the complex plane. We conclude this introductory chapter with a short note on the interpretation of domain-colourings, a technique that allows us to display the four dimensional graph of such a function within the confines of the printed page.

1.2 VISUALIZATION OF COMPLEX FUNCTIONS

Most figures in this text will be standard, however, there is one type of illustration that requires some introduction. Domain-colouring is a visualization technique for complex functions proposed by Farris (1998) and later popularized by the textbook of Wegert (2012). The idea is to map, for each point of the input plane, hue to phase (red corresponding to phase 0, green to $\frac{2}{3}\pi$, and blue to $\frac{4}{3}\pi$) and lightness to magnitude (black being zero and white infinity) as seen on the domain-colouring of $z \mapsto e^{-z}$ on fig. 1.1.

These visualizations can help illuminate many of the more exotic properties of complex functions. In the remainder of this text, however, we shall only ever encounter rational approximations of the exponential (for a sneak peak, compare figs. 1.1 and 3.1). Of such functions it is useful to be able to quickly identify the poles and zeroes. On fig. 1.2 we have an example of each. Of course, due to the colouring of the magnitude, a pole appears as a white dot and the zero as a black one. However, more strikingly, the phase reverses orientation between them. Indeed, from the polar form we directly see that $\frac{1}{re^{\theta i}} = \frac{1}{r}e^{-\theta i}$. Hence, we can quickly identify a zero as a singularity in the phase where the hue goes,




Figure 1.1: Domain-colouring of $z \mapsto e^{-z}$. Note how—as expected from Euler’s formula—the phase oscillates along the imaginary direction and how, along the real direction, the magnitude shrinks to zero.




Figure 1.2: Domain-colouring of $z \mapsto \frac{z-1}{(z+1)^2}$. Note how the phase around the zero at 1 goes from red, over green, to blue, how this order flips around the pole at -1 , and how the number of times these colours are repeated corresponds to the multiplicity and order of the zero and pole, respectively.

anticlockwise, from red, over green, to blue, and a pole as one where the hue goes from red, over blue, to green. Furthermore, the number of times these colours are repeated corresponds to the order and multiplicity of the pole and zero, respectively.

Note that the complex conjugate also reverse the phase, hence, in general, the magnitude still needs to be taken into account. However, for the functions appearing in this text, this ordering of the phase always holds.

A note on colour science

The domain-colourings throughout in this text were made using the `DomainColoring.jl` software package, developed by the author.¹ Notably, this package uses a colour wheel that is perceptually smooth and hence more readable than the commonly used sweep through the maximally saturated hues of the sRGB colour space. The latter was designed to mimic the colour response of old CRT monitors, not human perception, and the resulting colour wheel thus contains sharp kinks in both the perceived lightness and hue (IEC 61966-2-1:1999; Kovési 2015), as seen on fig. 1.3. For more information, see the accompanying documentation of the package.




Figure 1.3: Comparison of the colours used to represent phase at unit magnitude in `DomainColoring.jl` (top) and the more traditional sweep through the maximally saturated colours of the sRGB space (bottom). Phase runs from 0 to 2π from left to right, with a superimposed oscillation growing from bottom to top. Notice how much more detail of the oscillation is lost in the bottom image, in particular in the greens. Test images are modelled after Kovési (2015).

¹Available from <https://doi.org/10.5281/zenodo.20055018>.

2 *Time-delay systems*

In the introduction, we established that including delays in a model can be beneficial to abstract away the (possibly unknown) details of a transport phenomenon. Before turning to the numerical results of the following chapters, we provide a short review of the basic theory of linear time-delay systems with particular attention to the H^2 -norm, which will be the quantity we want to approximate starting chapter 4.

We limit our discussion mostly to theoretical properties we will need later on. Computational background material is deferred to the relevant chapter. More extensive introductions to the theory of time-delay systems, including nonlinear ones, can be found in the monographs of Hale (1977), Hale and Verduyn Lunel (1993), Diekmann et al. (1995), and Curtain and Zwart (1995). The theory of delay differential algebraic equations (DDAEs) is covered most extensively in the thesis of Unger (2020). A control oriented overview of computational methods, including DDAEs, is given in Michiels and Niculescu (2014). Finally, the classical reference on general numerical methods (with the notable omission of the Lanczos tau methods which are the topic of chapter 3) is Bellen and Zennaro (2003).

The first systematic study of time-delay systems is generally agreed to be in the work of Volterra (1928) in the early 1900s through the early 1930s, with applications to mathematical biology and viscoelasticity. In the decade that followed, thus the late 30s, early 40s, we find some of the first applications of delays in automatic control with the work of Minorsky (1942) on the stabilization and steering of ships. Again a decade later, now the late 40s, early 50s, the first monographs appear from Mishkis (1951), laying the groundwork for much of the general theory of linear systems, and Bellman and Danskin (1954), with the beginnings of stability theory. The latter was developed more extensively in the book of Bellman and Cooke (1963). By the late 70s and 90s much of the basic theory has stabilized and been systematically laid out in the volumes we already suggested. The interested reader may find a far more extensive history of the field in the (essentially identical) introductory chapters of Hale (1977) and Hale and Verduyn Lunel (1993).

Overview. The remainder of this chapter is structured as follows. We start our discussion with the simpler case of time-delay systems of retarded type (section 2.1) of which we discuss the basic existence and uniqueness results and spectral properties. In section 2.2 we make the connection to control by discussing the link between this spectrum and (exponential) stability. Many theoretical developments of these systems are based on the so called C_0 -semigroup formed by their solution operators and the associated infinitesimal generator, both of which are introduced in section 2.3. We then discuss how the aforementioned results change once neutral terms, i.e. delayed occurrences of the highest order derivative, are introduced and how we can include them using only algebraic constraints (section 2.4). Finally, we conclude this chapter with a discussion of the H^∞ - and H^2 -norm and their use in robust control and model order reduction in section 2.5.

2.1 TIME-DELAY SYSTEMS OF RETARDED TYPE

As mentioned in the introduction, we limit our study to linear, time-invariant, time-delay systems. Before discussing such systems in general, we start with about the simplest, autonomous, linear time-delay system imaginable, one with a single discrete delay:

$$\dot{\mathbf{x}}(t) = A_0 \mathbf{x}(t) - A_1 \mathbf{x}(t - \tau), \quad (2.1)$$

where $\mathbf{x}(t) \in \mathbb{C}^n$ is the state at time t , A_0 and A_1 are the system matrices, $\tau > 0$ is the delay, and $\dot{f} = \frac{d}{dt}f$ is the time derivative of f .

Note that, due to the presence of this delay, the information required to define a forward solution at $t = 0$ is not determined by $\mathbf{x}(0)$, but by the function segment $[-\tau, 0] \ni \theta \mapsto \mathbf{x}(\theta)$. More generally, the solution for all $t \geq t'$ is only defined given the solution for the time frame $[t' - \tau, t']$. Hence, the state at time t , in the natural meaning of the minimal information needed to determine the future evolution, corresponds to the function $\xi_t : [-\tau, 0] \rightarrow \mathbb{C}^n$, with $\xi_t(\theta) = \mathbf{x}(t + \theta)$, as illustrated on fig. 2.1. In other words, a delay differential equation (DDE) is a type of functional differential equation (FDE). This is why such systems are inherently infinite dimensional, their true state is a function.

This ξ_t , the so-called history function, allows us to provide the general formulation of a linear, time-invariant, delay differential equation of retarded type (RDDE):

$$\dot{\mathbf{x}}(t) = L(\xi_t), \quad (2.2)$$

with the Banach space

$$X := C([-\tau, 0]; \mathbb{C}^n),$$




Figure 2.1: Illustrative example of the history functions relevant to two different times. The black line on the left corresponds to the initial condition, i.e. the history function at $t = 0$, and the one on the right to $t = t'$. Notice how they always extend back to the (maximal) delay τ .

equipped with the supremum norm $\|f\|_\infty = \sup_{\theta \in [-\tau, 0]} f(\theta)$, the state space, i.e. $\xi_t \in X$,¹ and $L : X \rightarrow \mathbb{C}^n$ a bounded linear functional on X .² Note that this is still a fairly limited class of systems. Notably, we do not allow for a dependence on values of the highest order derivative in the past (then L would be unbounded). We will cover these so-called neutral type systems in section 2.4; it is, however, useful to start with the simpler case.

As mentioned, a forward solution of this system thus requires a function $\xi_0 \in X$ as initial state. For our choice of X , one can show that this solution is unique (Hale 1977: Theorem 6.1.1). In the delay-free case, i.e. (2.1) with $A_1 = \mathbf{0}$, we can decompose solutions in terms of eigenfunctions of the form

$$\mathbf{x}_s(t) = e^{st} \mathbf{p}_s(t),$$

where $\mathbf{p}_s(t)$ is a $\mathbb{C}^n$ -valued polynomial of degree smaller than the multiplicity of the eigenvalue s . The eigenvalues of this equation are those of A_0 and thus the solutions of the *characteristic equation* $\det(sI_n - A_0) = 0$. For RDDEs the connection is less direct, but exponential solutions will

¹For some applications, it is inconvenient to have a state space which does not form a Hilbert space, such as our X here. An alternative formulation, commonly used in the control setting, is introduced in section 2.3.

²In some texts, L is defined as $L(\phi) = \int_{-\tau}^0 [d\eta(\theta)] \phi(\theta)$, with η an $n \times n$ matrix valued function of bounded variation. In our case this representation is a direct consequence of Riesz' representation theorem for $C([- \tau, 0]; \mathbb{C}^n)$ (not to be confused with the related result for Hilbert spaces of the same name).

still be useful as they dominate solutions. Assume solutions of the same form as in the delay-free case. For simple roots we have $\mathbf{p}_s(t) = \mathbf{v}_s$, with $\mathbf{v}_s$ a constant vector, and find

$$s\mathbf{x}_s(t) = L(\theta \mapsto e^{(t+\theta)s}\mathbf{v}_s) = L(\theta \mapsto e^{\theta s}\mathbf{x}_s(t)),$$

where $\theta \mapsto f(\theta)$ is used to denote the univariate function f . Taking the natural extension of L to matrix-valued functions such that

$$L(F) = (L(\mathbf{f}_1) \ L(\mathbf{f}_2) \ \cdots \ L(\mathbf{f}_n)),$$

where $\mathbf{f}_k(s)$ is the k th column of $F(s)$, i.e. $\mathbf{f}_k(s) = (F(s))_k$, we can move $\mathbf{x}_s(t)$ out of the functional, giving

$$s\mathbf{x}_s(t) = L(\theta \mapsto I_n e^{\theta s})\mathbf{x}_s(t).$$

For non-simple roots the argument is more involved, but one can show that the same relation holds. The exponential solutions $\mathbf{x}_s(t)$ of (2.2) are thus those for which s is a solution of the characteristic equation $\det \Delta(s) = 0$, with

$$\Delta(s) = sI_n - L(\theta \mapsto I_n e^{\theta s})$$

the *characteristic matrix* of the system.³ We will refer to these s as *characteristic roots* and to the set of all these roots as the system's *spectrum* (the reason for this terminology will be apparent in section 2.3).

As mentioned before, these exponential solutions are relevant as they 'dominate' general solutions. The following classical result makes this precise, here in the formulation of Diekmann et al. (1995: Theorem I.5.4).

Theorem 2.1.1. *Let $\mathbf{x}(t)$ be a solution of the RDDE (2.2) for some initial condition ξ_0 . For any $\gamma \in \mathbb{R}$ such that $\det \Delta(\gamma + i\omega) \neq 0$ for all $\omega \in \mathbb{R}$, we have the asymptotic expansion*

$$\mathbf{x}(t) = \sum_{s \in C_\gamma} e^{st} \mathbf{p}_s(t) + o(e^{\gamma t}) \quad \text{as } t \rightarrow \infty,⁴$$

where $C_\gamma = \{s \in \mathbb{C} : \operatorname{Re} s > \gamma, \det \Delta(s) = 0\}$, i.e. the set of characteristic roots to the right of γ , and each $\mathbf{p}_s(t)$ is a $\mathbb{C}^n$ -valued polynomial of degree smaller than the multiplicity of s .

³In the formulation of footnote 2, we have $\Delta(s) = sI_n - \int_{-\tau}^0 e^{\theta s} d\eta(\theta)$.

⁴Remember that a function f is of order $o(g)$ if, and only if, f is dominated by g , i.e. $\lim_{x \rightarrow \infty} \frac{f(x)}{g(x)} = 0$. Later we will also use the related big O notation. A function f is of order $O(g)$ if, and only if, there is some constant M such that $|f(x)| \leq Mg(x)$ for all x in the domain of f . Note that generally some implicit asymptotic regime is assumed in the use of these notations.




Figure 2.2: The characteristic roots of (2.3). Notice that each root is isolated and that there are only ever a finite number of roots to the right of any vertical line.

This result will be most relevant to the discussion of the following section, but already hints that the characteristic roots define (most) of the possible behaviour of a system.

One may wonder to what extent we can let $\gamma \rightarrow -\infty$ to obtain a complete characterization of solutions? This is fairly well developed and is discussed at length in Hale and Verduyn Lunel (1993) and the other volumes suggested in this chapter's introduction, but would lead us too far. One aspect worth mentioning are so-called *small solutions*, solutions that go to zero faster than any exponential.

We can observe a few typical features of the spectra of RDDEs from the following example. Consider the scalar RDDE

$$\dot{x}(t) = -x(t-1), \quad (2.3)$$

which has as characteristic equation

$$s + e^{-s} = 0.$$

The characteristic roots are thus given by

$$s = -\ln |s| + i(\arg(s) + (2k+1)\pi), \quad \forall k \in \mathbb{Z}.$$

We see a plot of these on fig. 2.2. Note that the roots are isolated and that for every $\gamma \in \mathbb{R}$ there are only finitely many characteristic roots whose real part is larger than γ . One can show that both properties are

true in general for RDDES (Diekmann et al. 1995: Theorem 4.4). The arrangement of the roots is also typical of RDDES, an infinite tail with growing imaginary component as the real part goes to negative infinity.

In the next section we will start talking about concepts related to the control of these systems. In this setting, we will generally also include an input $\mathbf{u}(t) \in \mathbb{C}^p$ and an output $\mathbf{y}(t) \in \mathbb{C}^q$, giving

$$\begin{aligned}\dot{\mathbf{x}}(t) &= L(\xi_t) + B\mathbf{u}(t), \\ \mathbf{y}(t) &= C\mathbf{x}(t).\end{aligned}\tag{2.4}$$

A useful perspective to study these input-output systems is in the frequency domain, i.e. using the Laplace transform. Assume $\xi_0(\theta) = \mathbf{0}$ and let $\hat{f}(s)$ be the transform of $f(t)$. We want to find the so-called transfer function $G(s)$, such that $\hat{\mathbf{y}}(s) = G(s)\hat{\mathbf{u}}(s)$. Note that, due to the time shift property of the Laplace transform

$$\hat{\xi}_s(\theta) = \hat{\mathbf{x}}(s)e^{\theta s}.$$

Thus, from linearity, we have

$$L(\hat{\xi}_s) = L(\theta \mapsto I_n e^{\theta s})\hat{\mathbf{x}}(s).$$

Taking the Laplace transform of (2.4), we get

$$\begin{aligned}s\hat{\mathbf{x}}(s) - \mathbf{0} &= L(\theta \mapsto I_n e^{\theta s})\hat{\mathbf{x}}(s) + B\hat{\mathbf{u}}(s), \\ \hat{\mathbf{y}}(s) &= C\hat{\mathbf{x}}(s).\end{aligned}$$

After eliminating $\hat{\mathbf{x}}(s)$, we finally find

$$G(s) = C(sI_n - L(\theta \mapsto I_n e^{\theta s}))^{-1}B.$$

Note that our formulation here admits only state delays. Sometimes one also wants delays of the input signal $\mathbf{u}(t)$. Like for neutral type systems, we shall see in section 2.4 how these fit within the broader class of delay differential algebraic equations.

2.2 EXPONENTIAL STABILITY AND CONTROL

An important first step in most control problems is to check or ensure stability of the system (2.4). That is, some guarantee that solutions will not grow unboundedly. In the input-output setting, one generally requires that a bounded input always corresponds to a bounded output. The usual criterion is to require that the autonomous dynamics, i.e. those of (2.2) with the same L , satisfy the following condition.

Definition 2.2.1. The system (2.2) is said to be *exponentially stable* if, and only if, there exists some $\lambda < 0$ and $M \geq 1$ such that, for all bounded initial conditions ξ_0 , we have

$$\|\mathbf{x}(t)\|_2 \leq M e^{\lambda t} \|\xi_0\|_X, \quad \forall t \geq 0.$$

In a slight conflation of terms, we shall also call the input-output system (2.4) exponentially stable if its autonomous dynamics satisfy this definition.

If a system is exponentially stable, there is some α , called the *spectral abscissa*, such that all $\lambda < \alpha$ satisfy the bound of Definition 2.2.1. The word ‘spectral’ in its name refers to the fact that it is easily seen from Theorem 2.1.1 that this value is the real part of the right-most characteristic root of the system, i.e. the right-most element of the spectrum. Indeed, as solutions of (2.2) are dominated by the characteristic solution corresponding to this root, the worst possible decay (or growth) rate thus stems from the right-most characteristic root. If this root is simple, $\lambda = \alpha$ satisfies the bound. Conversely, if the real part of this root is non-negative, the system is not exponentially stable, as there are now solutions that grow exponentially or, for $\alpha = 0$, that have constant amplitude or polynomial growth. We will in this case still refer to this value as the spectral abscissa. Hence, the sign of α determines whether a system is exponentially stable or not.

As the roots of $\det \Delta(s)$ are precisely the values for which $sI_n - L(\theta \mapsto I_n e^{\theta s})$ is singular, it is natural to wonder if the characteristic roots and the poles of $G(s)$ are equivalent. This is only partially true, as it is possible for poles to be cancelled by zeroes in the transfer function, hiding some of the characteristic roots. That is to say, the poles are only a subset of these roots.

On the one hand, this does mean that, in the mathematical utopia of the model, the output will indeed be bounded if the input is too. However the internal state might grow unboundedly. In applications, this is also generally unwanted behaviour. In addition, an infinitesimal perturbation of the system might reintroduce the pole. Hence, if we speak of exponential stability of input-output systems, we will always look at the characteristic roots.

Furthermore, if infinitesimal perturbations can introduce poles on the right half-plane, so too can numerical round-off. If we do truly not care for those characteristic roots that do not correspond to poles of the transfer function, we should thus eliminate them from the state space realization (2.4). A realization that does not have any such additional characteristic roots is said to be a *minimal realization*.

As we saw, the worst possible behaviour in terms of decay (or growth) is given by the spectral abscissa. Hence, if we have control over some

part of the system (typically a subsystem modelling an actual controller) we can optimize its parameters to make the spectral abscissa α as small as possible. Of course reality is rarely this simple and often when α is pushed too far to the left other issues appear. Whilst it tells us something about the decay rates of the tail, the transient behaviour can still grow very large. Similarly, the input the controller receives likely contains measurement noise, the obtained controller might be highly sensitive to this noise and end up being unstable in practice.

In robust control one tries to take these issues into account as we shall see in section 2.5. Such control will thus often not bring α to its absolute minimum. However, we will almost always want to move at least all characteristic roots to the left of the imaginary axis.

2.3 ELEMENTS OF C_0 -SEMIGROUP THEORY

Another fruitful perspective on (2.2) is the so-called C_0 -semigroup formed by its bounded solution operators $\{\mathcal{S}_T : X \rightarrow X\}_{T \geq 0}$, which map a history function ξ_t to $\mathcal{S}_T \xi_t = \xi_{t+T}$. These form such a semigroup as they satisfy (Hale 1977: Lemma 7.1.1):

$$\begin{aligned} \mathcal{S}_0 &= I, \\ \mathcal{S}_{T+S} &= \mathcal{S}_T \mathcal{S}_S, \quad \forall T, S \geq 0, \quad \text{and} \\ \|\mathcal{S}_T \xi - \xi\|_X &\rightarrow 0 \quad \text{as } T \rightarrow 0^+, \quad \forall \xi \in X, \end{aligned}$$

where I is the identity operator. These properties allow us to define the *infinitesimal generator* $\mathcal{A}_X : D(\mathcal{A}_X) \rightarrow X$ of this C_0 -semigroup, with action

$$\mathcal{A}_X \xi = \lim_{T \rightarrow 0^+} \frac{1}{T} (\mathcal{S}_T - I) \xi = \xi'$$

and domain

$$D(\mathcal{A}_X) = \{\xi \in X : \xi \in C^1([-\tau, 0]; \mathbb{C}^n), \xi'(0) = L(\xi)\},$$

where $f' = \frac{d}{d\theta} f$ is the derivative with respect to the history parameter. Results relating to this formulation are more extensively developed in the aforementioned books (see e.g. Hale 1977: chapter 7), in the remainder of this section we give a rough flavour of this approach as the infinitesimal generator is the natural framing for the discretization strategy we will use in the later chapters.

This infinitesimal generator forms a particularly useful object to study the behaviour of the system. Like the characteristic roots we have seen so far, the spectrum of $\mathcal{A}_X$ can also be used to evaluate the stability of the system. In fact, one can show that its spectrum is a pure point spectrum

and that the eigenvalues correspond one to one with the characteristic roots (Hale 1977: Lemma 7.2.1). This spectrum then allows various decompositions of the state space X (Hale 1977: chapter 7), which allow for the more comprehensive discussion of possible behaviours we alluded to in section 2.1.

Note, however, that in the above formulation the actual system dynamics are hidden in the domain of $\mathcal{A}_X$. Especially in the context of control, i.e. for the infinitesimal generator associated with (2.4), this is inconvenient, as this obfuscates the system's input and output. Furthermore, $X = C([- \tau, 0]; \mathbb{C}^n)$ is famously not a Hilbert space, which means the adjoint, which is convenient for formulating some common target functions in optimal control, lives in a different, and more complicated, space.

To address both issues, we can explicitly decouple the current state from the history (resulting in the so-called 'head-tail' representation) as probably first used by Krasovskii (1964) and developed most extensively by Bernier and Manitius (1978). We now consider as state space for (2.4)

$$Z := \mathbb{C}^n \times L^2([- \tau, 0]; \mathbb{C}^n), \quad (2.5)$$

giving as infinitesimal generator the operator $\mathcal{A}_Z : D(\mathcal{A}_Z) \rightarrow Z$ with action

$$\mathcal{A}_Z(\mathbf{z}, \zeta) = (L(\zeta), \zeta')$$

and domain

$$D(\mathcal{A}_Z) = \{(\mathbf{z}, \zeta) \in Z : \zeta \in W^{1,2}([- \tau, 0]; \mathbb{C}^n), \mathbf{z} = \zeta(0)\},$$

where $W^{1,2}([- \tau, 0]; \mathbb{C}^n)$ is the Sobolev space of functions with one weak derivative in $L^2([- \tau, 0]; \mathbb{C}^n)$.⁵

As the equation itself now appears in the action of $\mathcal{A}_Z$, we can naturally reformulate (2.4) as an abstract Cauchy problem. Let operators $\mathcal{B} : \mathbb{C}^p \rightarrow Z$ and $\mathcal{C} : Z \rightarrow \mathbb{C}^q$ be defined by

$$\mathcal{B}\mathbf{u} = (B\mathbf{u}, \mathbf{0}) \quad \text{and} \quad \mathcal{C}z = Cz,$$

where $\mathbf{u} \in \mathbb{C}^p$ and $z = (\mathbf{z}, \zeta) \in Z$. We can then write

$$\begin{aligned} \dot{z}(t) &= \mathcal{A}_Z z(t) + \mathcal{B}\mathbf{u}(t), \\ \mathbf{y}(t) &= \mathcal{C}z(t), \end{aligned} \quad (2.6)$$

⁵The Sobolev spaces $W^{k,2}$ are often denoted by H^k , as they form Hilbert spaces. However, this might lead to confusion with our notation for the H^2 - and H^∞ -norms, so we will prefer the former.




Figure 2.3: The characteristic roots of (2.7). Compare to fig. 2.2 and note that, whilst they are still isolated, there is now a vertical chain with infinitely many characteristic roots.

with $z(t) = (\mathbf{z}(t), \zeta_t) \in D(\mathcal{A}_Z)$. The relation between corresponding solutions of (2.4) and (2.6) is then given by

$$\mathbf{z}(t) = \mathbf{x}(t) \quad \text{and} \quad \zeta_t(\theta) = \mathbf{x}(t + \theta), \quad \forall \theta \in [-\tau, 0].$$

For a more detailed description of the mapping between representations, and further detail on the inclusion of input and output, we refer to Curtain and Zwart (1995). We shall see in chapter 3 that abstract Cauchy problem formulations lend themselves well to discretization.

2.4 NEUTRAL SYSTEMS AND ALGEBRAIC CONSTRAINTS

As we already noted in section 2.1, the systems we can describe in the form of (2.2) or (2.4) are somewhat limited in scope. In this section we shall see that we can greatly extend the number of possibilities by allowing algebraic constraints.

A first limitation—actually the defining feature—of RDDEs is that the highest order derivative never appears delayed. However, it is not uncommon to have a delayed reading of the highest order derivative, for instance an accelerometer in a mechanical system. If we do allow such a dependence, we get a so called neutral equation (NDDE).

Such equations show some trickier behaviour. Consider for example

$$\dot{x}(t) = -\frac{1}{2}\dot{x}(t-1). \quad (2.7)$$

Through a similar process to section 2.1, we get as characteristic equation

$$s + \frac{1}{2}se^{-s} = 0.$$

The characteristic roots are thus given by $s = 0$ and

$$s = -\ln(2) + i(2k + 1)\pi, \quad \forall k \in \mathbb{Z},$$

as seen on fig. 2.3. Note that, unlike the RDDE on fig. 2.2, we now have an infinite vertical chain of characteristic roots. Hence, a vertical strip need not have a finite number of roots.

In fact it is not obvious what properties of RDDEs remain true if we allow such dependence on the derivative. Let us first introduce the general form of a linear autonomous NDDE:

$$\frac{d}{dt}D(\xi_t) = L(\xi_t), \quad (2.8)$$

with X as state space, $\xi_0 \in X$ as initial condition, D and L bounded linear functionals on X , and D atomic at 0 (see Hale and Verduyn Lunel 1993: Definition 2.5.3). The precise definition of atomic would lead us too far, but this condition essentially requires that $\mathbf{x}(t) = \xi_t(0)$ can be recovered from $D(\xi_t)$.

As for RDDEs, this system has a unique solution for this choice of state space (Hale and Verduyn Lunel 1993: Theorem 9.1.1). The solution operators also still form a C_0 -semigroup and we get, *mutatis mutandis*, essentially the same infinitesimal generator as $\mathcal{A}_X$ in section 2.3 which also has a purely point spectrum, with the eigenvalues given by the solutions of $\det \Delta(s) = 0$ with

$$\Delta(s) = sD(\theta \mapsto I_n e^{\theta s}) - L(\theta \mapsto I_n e^{\theta s})$$

the characteristic matrix of the neutral system (Hale and Verduyn Lunel 1993: Lemma 9.2.1). There is still a leftmost vertical line no eigenvalue crosses, however this need not be an attained maximum anymore. The spectral abscissa is thus now given by the supremum over the real parts of the eigenvalues. This spectral abscissa still bounds the maximal exponential growth (or slowest decay) of solutions of the system (Hale and Verduyn Lunel 1993: Corollary 9.3.1).

Thus, at first glance, they are not too different from the RDDE case, however they can also be sensitive to infinitesimal delay perturbations. Take the following example, adapted from Michiels and Niculescu (2014: chapter 1),

$$\dot{x}(t) = -\frac{1}{2}x(t) + \frac{3}{4}\dot{x}(t-1) - \frac{1}{2}\dot{x}(t-2+\epsilon). \quad (2.9)$$




Figure 2.4: The characteristic roots of (2.9), left for $\epsilon = 0$ and right for $\epsilon = 0.065$. The left example is stable, the rightmost root is to the left of the imaginary axis (the dashed line); the right is not. Note that we took a fairly large value for ϵ here. For any nonzero value the result is unstable, the smaller ϵ the larger the imaginary component of the crossing roots.

For $\epsilon = 0$ this system is stable, however under the slightest perturbation the system suddenly becomes unstable, as illustrated on fig. 2.4. When we start allowing neutral terms in chapter 5, we will thus have to take such sensitivities into account.

Whilst we could explicitly work with an additional functional as in (2.8), note that we can equally write an NDDE as an RDDE with an algebraic constraint by introducing an additional variable. Take for example (2.7), we can rewrite this into the equivalent

$$\begin{cases} \dot{x}_1(t) = -\frac{1}{2}x_2(t-1), \\ \dot{x}_1(t) = x_2(t), \end{cases}$$

or, in matrix notation,

$$\begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} \dot{\mathbf{x}}(t) = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} 0 & -\frac{1}{2} \\ 0 & 0 \end{pmatrix} \mathbf{x}(t-1),$$

giving a delay differential *algebraic* equation (DDAE).

As we see from this example, adding algebraic constraints corresponds to multiplying by a singular matrix on the left hand side. We can thus extend (2.2) to a DDAE as

$$E\dot{\mathbf{x}}(t) = L(\xi_t),$$

where E is, typically, a singular matrix. Using identical reasoning to section 2.1, we find as characteristic matrix

$$\Delta(s) = sE - L(\theta \mapsto I_n e^{\theta s}).$$

Similarly to before, we can again introduce inputs and outputs, giving

$$\begin{aligned} E\dot{\mathbf{x}}(t) &= L(\xi_t) + B\mathbf{u}(t), \\ \mathbf{y}(t) &= C\mathbf{x}(t), \end{aligned} \tag{2.10}$$

with transfer function

$$G(s) = C(sE - L(\theta \mapsto I_n e^{\theta s}))^{-1} B.$$

Such a system is sometimes also referred to as a descriptor system.

For input-output systems, DDAEs allow for additional options—or pitfalls, depending on your perspective. By setting an additional variable equal to an input, we can easily obtain input delays and even feedthrough (a direct connection between an input and an output). Particularly the latter can be an issue, as the H^2 -norm, which we will discuss in the next section, blows up for a system with feedthrough. Worse yet, such a feedthrough can again appear under infinitesimal perturbations. We shall return to both the benefits and these issues in chapter 5.

Note that these extensions might appear rather mild from our exposition, however that is in part due to our focus on the linear theory. Nonlinear NDDs form a significant further complication (see e.g. Hale and Verduyn Lunel 1993: chapter 9) and even linear DDAEs are more treacherous than we have shown here. Like for DAEs, there exists various indices that measure the ‘normalness’ of the equation. One of the main assumptions we shall make at the start of chapter 5 is a restriction to DDAEs of Kronecker index one (in the sense of Definition 5.0.1). A more complete description of linear and nonlinear DDAEs and their indices can be found in the thesis of Unger (2020).

2.5 THE H^∞ - AND H^2 -NORM

At the end of section 2.2, we noted that only optimizing for the spectral abscissa runs the risk of introducing other issues. In the final section of this chapter, we look at two other system metrics that aim to provide insight into some of these problems. Unlike the spectral abscissa, which looks only at the input-free dynamics, these system norms relate directly to the input-output behaviour by looking at the transfer function. The ones we discuss are the H^∞ -norm, given by

$$\|G\|_{H^\infty} = \sup_{s \in \mathbb{C}_{>0}} \|G(s)\|_2,$$

where $\mathbb{C}_{>0} = \{z \in \mathbb{C} : \operatorname{Re}(z) > 0\}$ is the open right half-plane and $\|A\|_2$ is the spectral norm of A (i.e. the largest singular value of A), and the H^2 -norm, defined as

$$\|G\|_{H^2} = \sup_{\beta > 0} \left(\frac{1}{2\pi} \int_{-\infty}^{\infty} \|G(\beta + i\omega)\|_F^2 d\omega \right)^{\frac{1}{2}},$$

where $\|A\|_F = \sqrt{\operatorname{tr}(AA^*)}$ is the Frobenius norm.⁶

Note that if G has poles on the closed right half-plane, i.e. a minimal realization of G is not exponentially stable, both norms are infinite. Like section 2.2, there is again some nuance for systems with pole-zero cancellations. In exact arithmetic these generally have finite norms (barring the H^2 -norm in the case of feedthrough, something we return to in chapter 5), however, for many real world applications this is arguably unwarranted, and numerical round-off can still result in an infinite norm. Going forward, we shall thus take exponential stability as a prerequisite before computing these norms.

For such an exponentially stable system, G is guaranteed to be analytic on the closed right half-plane. In this case, these norms simplify somewhat as, due to the maximum modulus principle, we now only need to evaluate them along the imaginary axis, giving

$$\|G\|_{H^\infty} = \operatorname{ess\,sup}_{\omega \in \mathbb{R}} \|G(i\omega)\|_2$$

and

$$\|G\|_{H^2} = \left(\frac{1}{2\pi} \int_{-\infty}^{\infty} \|G(i\omega)\|_F^2 d\omega \right)^{\frac{1}{2}},$$

respectively.

Like for the spectral abscissa, one can design controllers by optimizing these system norms, or a mixture of them, giving a common form of robust control. For the case of the H^∞ -norm, the interpretation (and the link to robustness) is rather straightforward; it is equal to the maximal amplification from input to output. The computation and optimization of this norm, whilst nontrivial, is well-established and we defer an interested reader to Foias et al. (1996), with a number of further algorithmic developments in Michiels and Niculescu (2014) and the thesis of Appeltans (2024). An overview of both the H^∞ - and H^2 -norm in the delay-free setting is given in K. Zhou et al. (1995).

⁶In complex analysis these are known as Hardy norms. However, note that the convention here is a bit atypical. We work with Hardy spaces on the right half-plane, instead of the upper half-plane or the unit disc.

Starting chapter 4, we will focus our attention on the H^2 -norm. Its interpretation will require a bit more exposition, but it is useful as a measure of disturbance rejection and, in the context of model reduction, to quantify the approximation error at the transfer function level. As it is a global characteristic, in the sense that it depends on the transfer function's behaviour along the entire imaginary axis, its computation is also challenging. For delay-free systems an efficient method involving the Lyapunov equation is well known (which we detail in section 4.2), how to do this for time-delay systems will be the topic of those later chapters.

A first interpretation and use case stems from the fact that the H^2 -norm is closely related to the spectral power of the output signal. Consider an input signal $\mathbf{u}(t)$ that is wide-sense stationary, i.e. with static mean and autocovariance, and finite variance. Due to the Wiener–Khinchin theorem, its spectral density $S_{\mathbf{u}\mathbf{u}}$ is then given by the Fourier transform of the autocorrelation matrix $R_{\mathbf{u}\mathbf{u}}$, that is

$$S_{\mathbf{u}\mathbf{u}}(\omega) = \int_{-\infty}^{\infty} R_{\mathbf{u}\mathbf{u}}(\tau) e^{-i\omega\tau} d\tau,$$

where

$$R_{\mathbf{u}\mathbf{u}}(\tau) := \mathbb{E}[\mathbf{u}(t + \tau)\mathbf{u}(t)^*] = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^T \mathbf{u}(t + \tau)\mathbf{u}(t)^* dt.$$

The spectral power of the signal is then defined as the square root of its variance, which, using the inverse Fourier transform, is given by

$$\|\mathbf{u}\|_P := \sqrt{\mathbb{E}[\mathbf{u}(t)^*\mathbf{u}(t)]} = \sqrt{\text{tr}(R_{\mathbf{u}\mathbf{u}}(0))} = \left(\frac{1}{2\pi} \int_{-\infty}^{\infty} \text{tr}(S_{\mathbf{u}\mathbf{u}}(\omega)) d\omega \right)^{\frac{1}{2}}.$$

As one can show (K. Zhou et al. 1995: p. 103) that

$$S_{\mathbf{y}\mathbf{y}}(\omega) = G(i\omega)S_{\mathbf{u}\mathbf{u}}(\omega)G(i\omega)^*,$$

the spectral power of the output signal is then given by

$$\|\mathbf{y}\|_P = \left(\frac{1}{2\pi} \int_{-\infty}^{\infty} \text{tr}(G(i\omega)S_{\mathbf{u}\mathbf{u}}(\omega)G(i\omega)^*) d\omega \right)^{\frac{1}{2}}.$$

The H^2 -norm of the transfer function hence gives us the spectral power of the output signal, if the input signal has spectral density $S_{\mathbf{u}\mathbf{u}}(\omega) = I_n$, i.e. is white noise. If we want to measure, and subsequently minimize, the sensitivity of a subsystem to noise on the input, the H^2 -norm is thus the natural metric. Note that its use is not limited to white noise, as




Figure 2.5: The combined system used in model approximation. Here, G is the transfer function of the original system and G_r of the approximating system. A good approximation can be found by minimizing the H^∞ - and/or H^2 -norm of the entire system.

we can measure the effect of coloured noise by inserting a filter into the measured transfer function.

Another application is the approximation of an exponentially stable system with transfer function G by another (typically smaller) system with transfer function G_r . In this context we typically want G_r to track the frequency response of G . Note that the H^2 -norm can be thought of as the L^2 -norm of the frequency response. Tracking the frequency response well in the L^2 sense is hence equivalent to the combined system on fig. 2.5, where the input is shared between G and G_r and the outputs subtracted, having a small H^2 -norm. Furthermore, as an unstable G_r would render the combined H^2 -norm infinite, this approach automatically finds a stable approximation. We work through this process for a concrete system in example 7.2.3. Alternatively, one can also use the H^∞ -norm, which would minimize the largest error between the frequency responses, i.e. finds a stable, optimal approximation in the L^∞ sense.

2.6 SUMMARY

In this chapter we gave a short overview of the basic theory of time-invariant linear time-delay systems. We started with those of retarded type, i.e. those where the highest order derivative never appears delayed, in section 2.1. We established that their true state is a function, the so-called history function, and that these systems are hence inherently infinite dimensional. Nonetheless, a lot of properties we would expect from the finite dimensional, delay-free case still hold. Solutions exist and are unique for continuous initial conditions (Hale 1977: Theorem 6.1.1) and, whilst not giving a complete description of the system, exponential solutions still largely define what is possible (in the sense of Theorem 2.1.1).

As we saw in section 2.2, the characteristic roots associated with these solutions can be used to define a quantifiable notion of stability (Definition 2.2.1). This spectral abscissa α , which corresponds to the fastest growth rate (or slowest decay rate), gives a first target for optimal control.

In section 2.3, we saw an alternative perspective on these systems by studying the C_0 -semigroup formed by their solution operators and in particular the associated infinitesimal generator. We saw that this operator has a pure point spectrum and that its eigenvalues are identical to the characteristic roots we saw before. Within this functional analytic framework we also saw an alternative formulation, popular in control applications, using the so-called ‘head-tail’ representation, which exposes the system dynamics more directly and uses a Hilbert space as state space instead of the continuous functions.

Next, we saw what happens in the neutral case, i.e. when we do allow delayed occurrences of the highest order derivative, and how they relate to the inclusion of algebraic constraints in section 2.4. A notable feature of these neutral type systems is that they introduce vertical chains of infinitely many characteristic roots into the spectrum, which also means that the spectral abscissa, whilst still bounding the growth rate, is now a supremum. These chains can also be sensitive to infinitesimal delay perturbations, as seen in the behaviour of system (2.9) which becomes unstable for any perturbation. In the remainder of the section we saw that adding algebraic constraints to an RDDE, yielding a so-called DDAE, allows one to add neutral terms. Furthermore, they also allow for the inclusion of input delays, feedthrough, and the straightforward coupling of existing models. However, this flexibility comes at a cost, as will be discussed further at the start of chapter 5.

We concluded the chapter with a discussion of the H^∞ - and H^2 -norm and their use in robust control and model order reduction in section 2.5. Both norms simplify to measurements along the imaginary axis for exponentially stable systems. We continued by describing how the H^2 -norm can be used to measure sensitivity to noise as it is equal to the spectral power of the output if the input is white noise. Finally, we saw that both the H^∞ -norm and H^2 -norm can be used to construct stable reduced order models that capture the frequency response of the original system well.

In the following chapter we address the infinite dimensional nature of these systems from the computational side using the Lanczos tau discretization, essentially approximating an infinite dimensional system by a finite dimensional one, which builds on the abstract Cauchy problem we saw in section 2.3. The remaining chapters will be dedicated to using this finite dimensional system as a proxy for the original to approximate the H^2 -norm, with the ultimate aim to compute this norm for DDAEs.

3 The Lanczos tau method

In the previous chapter we established that time-delay systems give rise to functional differential equations and are hence inherently infinite dimensional. This infinite dimensional nature implies that existing techniques for the analysis and design of delay-free systems cannot be readily applied. New methods thus have to be developed, most of which start by discretizing the infinite dimensional system into a finite-dimensional approximation, so too will our approach to approximating of the H^2 -norm we discuss in the later chapters.

The focus of this chapter is on one of the possible choices to arrive at such a delay-free approximation: the Lanczos tau method of Ito and Teglas (1986). We will introduce a novel operator based formulation of this discretization, which greatly simplifies theoretical derivations. It will allow us to extend the understanding of this method in the frequency domain, show connections between the Lanczos tau method and other popular discretization schemes, and construct highly sparse and numerically beneficial matrix realizations. In the later chapters we shall see, as a core theme of this thesis, that this choice of discretization is particularly fruitful for the computation of the H^2 -norm.

For simplicity we shall limit our discussion, in this and the following chapter, to a system with a single discrete delay. Many of the presented results straightforwardly extend to more general systems and we shall comment on this as we go along. For now, we thus consider time-delay systems described by

$$\begin{aligned}\dot{\mathbf{x}}(t) &= A_0\mathbf{x}(t) + A_1\mathbf{x}(t - \tau) + B\mathbf{u}(t), \\ \mathbf{y}(t) &= C\mathbf{x}(t),\end{aligned}\tag{3.1}$$

The primary contents of this chapter previously appeared as the first half of Provoost, E., and Michiels, W. (2024), “The Lanczos tau framework for time-delay systems: Padé approximation and collocation revisited”, *SIAM Journal on Numerical Analysis*, 62/6: 2529–48. A number of results in sections 3.4 and 3.8 have been adapted from Provoost and Michiels (2026). Proposition 3.7.4, the related discussion, and all figures have been prepared specifically for this thesis. The sparse realization of a Lanczos tau discretization truncating in a Legendre basis using generalized Gegenbauer polynomials in section 3.6 has also not previously appeared.

where $\tau \in \mathbb{R}_+$ is the constant delay, $\mathbf{x}(t) \in \mathbb{C}^n$ the state variable, $\mathbf{u}(t) \in \mathbb{C}^p$ the input, and $\mathbf{y}(t) \in \mathbb{C}^q$ the output at time t . That is, system (2.4) with $L(\xi) = A_0\xi(0) + A_1\xi(-\tau)$. The transfer function of this system is given by

$$G(s) = C(sI_n - A_0 - A_1e^{-s\tau})^{-1}B. \quad (3.2)$$

Overview. After giving a historical review of the various discretization strategies available for time-delay systems and reviewing basic properties of orthogonal polynomials in sections 3.1 and 3.2, respectively, we present an operator pencil framework for Lanczos tau discretizations (section 3.3). We continue by discussing some properties of these methods. In particular, we show how these can be interpreted as rational approximation in the frequency domain (section 3.4) and how they are deeply connected to other approximations. We prove equivalence to pseudospectral collocation, when the non-zero collocation points are chosen as the zeroes of an orthogonal polynomial (section 3.5). Collocation methods are known to be a good fit for the approximation of the characteristic roots and in section 3.6 we see how Lanczos tau methods naturally lead to sparse, nested discretizations, which is computationally beneficial in this context. In the same section we also introduce a highly sparse and numerically beneficial matrix realization of the Lanczos tau method using a Legendre basis, which will be our method of choice in the following chapter. We conclude with a surprising link to Padé approximation, when using shifted Legendre polynomials (section 3.7), which additionally allows us to provide a first set of convergence and stability proofs (section 3.8) and is part of the reason we will prefer the Legendre basis later on.

3.1 A SHORT HISTORY OF DISCRETIZATION TECHNIQUES

By far the most common approach—as this is the default in many commonly available software packages—is to replace the exponential function by a rational approximation in the frequency domain, such as the Padé approximant (see e.g. Glover et al. 1991). However, if expressed as a fraction of polynomials in monomial form, the coefficients of this approximation grow rapidly, leading to poor numerical performance. A surprisingly common assumption when working with time-delay systems then is that a Padé approximant of moderate degree is the best one can do. Clearly this leaves a lot to be desired, and for our later purposes much higher degrees would be needed. Alternative representations, such as the barycentric form (see e.g. Silviu-Ioan et al. 2018: subsection 2.2), might alleviate these issues, but are to our knowledge unexplored for time-delay systems.

Another way to not face this issue, is to discretize at the level of the state space. There are two main variants taking this approach. As the system is linear and time-invariant, one can approximate the solution operator $\mathcal{S}_T$, which, as seen in section 2.3, maps the history function ξ_t to ξ_{t+T} . Such an approach is for instance adopted for linearized stability analysis in the bifurcation analysis package of Engelborghs et al. (2002), where $\mathcal{S}_T$ is approximated using a linear multistep method. Similarly, in the context of stability analysis of periodic delay-differential equations, Butcher et al. (2004) propose to discretize $\mathcal{S}_T$ as its action on a Chebyshev series approximation of ξ_t , where T is taken to be the period.

The other option to discretize at the state space level, is to look at the infinitesimal generator of the C_0 -semigroup $\{\mathcal{S}_t\}_{t \geq 0}$ (see section 2.3). When the input and output in (3.1) are taken into account, this results in a standard state space description of a delay-free system, which captures part of the system behaviour of the original system. The main advantage of these methods is thus that one can often readily apply existing techniques for the analysis and design of finite-dimensional systems to this approximation.

The earliest method in the latter category, to the best of our knowledge, is the Lanczos tau method of Ito and Teglas (1986), which relies on a truncated Legendre basis. In a series of followup papers, they provide a spline based variation (Ito 1985) and even prove preservation of stability (Ito 1985; Ito and Kappel 1992), which for most other discretizations remains an important open problem. Their approach was extended to other bases, and shown to perform well when computing the characteristic roots, in Vyasarayani et al. (2014), be it without providing results on the stability of the discretization. Apart from the work of Scholl et al. (2024a,b), this method appears somewhat forgotten in recent years. However, as we alluded to in the introduction, we will see that it shows a lot of merit for the approximation of the H^2 -norm.

Another, more popular, approach in this class is pseudospectral collocation, introduced by Breda, Maset, et al. (2005). It was initially presented for the eigenvalue problem, but later successfully extended to construct delay-free approximations, which can for instance be used for bifurcation analysis (Breda, Diekmann, et al. 2016). This method collocates the action of the infinitesimal generator on a grid of nodes. In practice these are usually the Chebyshev extremal nodes, a set of nodes which are distributed more densely near the end points. This distribution evades Runge's phenomenon, which one can face in more naive collocation strategies (Boyd 2001). Pseudospectral collocation is tightly linked to rational approximation of the exponential, which was for instance used in the initial paper to prove super-geometric convergence of the characteristic roots. This link can also be exploited in algorithmic

design. An example is a heuristic by Wu and Michiels (2012) to select a discretization degree such that all the characteristic roots to the right of the imaginary axis are sufficiently well approximated.

3.2 A REVIEW OF ORTHOGONAL POLYNOMIALS

The Lanczos tau method presented in the next section will rely on the notion of a degree-graded series of polynomials orthogonal with respect to an inner product $\langle \cdot, \cdot \rangle$, with induced norm $\|\phi\| = \sqrt{\langle \phi, \phi \rangle}$.¹ That is, a set of polynomials $\{\phi_k\}_{k=0}^\infty$ which has as defining property that ϕ_k is of degree k and $\langle \phi_j, \phi_k \rangle = 0$ if and only if $j \neq k$. Usually, the inner product chosen is of the form

$$\langle \phi_j, \phi_k \rangle = \int_{-\tau}^0 \phi_j(\theta) \overline{\phi_k(\theta)} w(\theta) d\theta,$$

with $w(\theta) \geq 0$, $\forall \theta \in [-\tau, 0]$, the weight function. The choice of w , together with a normalization condition, then uniquely defines the orthogonal sequence.

Throughout this work we will rely on a number of such classical orthogonal polynomials shifted to the delay interval $[-\tau, 0]$. To lighten notation we shall denote these shifted polynomials by their usual names in the literature for the interval $[-1, 1]$, unless otherwise noted. In particular, we shall use the Jacobi polynomials $P_k^{(\alpha, \beta)}$ shifted to the interval $[-\tau, 0]$, which are given by the weight function

$$w(\theta) = \left(-\frac{2}{\tau}\theta\right)^\alpha \left(\frac{2}{\tau}\theta + 2\right)^\beta,$$

for $\alpha > -1$ and $\beta > -1$, and normalization condition $P_k^{(\alpha, \beta)}(0) = \binom{k+\alpha}{k}$. As special cases we have

$$T_k = \binom{k-\frac{1}{2}}{k}^{-1} P_k^{(-\frac{1}{2}, -\frac{1}{2})} \quad \text{and} \quad U_k = (k+1) \binom{k+\frac{1}{2}}{k}^{-1} P_k^{(\frac{1}{2}, \frac{1}{2})},$$

the shifted Chebyshev polynomials of, respectively, the first and second kind, and the shifted Legendre polynomials

$$P_k = P_k^{(0,0)}.$$

The latter two are more readily obtained as special cases of the Gegenbauer polynomials

$$C_k^\lambda = \frac{\Gamma(\lambda + \frac{1}{2})}{\Gamma(2\lambda)} \frac{\Gamma(k+2\lambda)}{\Gamma(k+\lambda+\frac{1}{2})} P_k^{(\lambda-\frac{1}{2}, \lambda-\frac{1}{2})},$$

¹We assume the convention that $\langle \cdot, \cdot \rangle$ is linear in the first argument and antilinear in the second.

where $\lambda > -\frac{1}{2}$, through

$$P_k = C_k^{\frac{1}{2}} \quad \text{and} \quad U_k = C_k^1.$$

The Chebyshev polynomials of both kinds are also closely related to the trigonometric functions through the identities (Szegő 1975: eq. 4.1.7)

$$T_k(\theta) = \cos(k\omega) \quad \text{and} \quad U_k(\theta) = \frac{\sin((k+1)\omega)}{\sin(\omega)},$$

where $\theta = \frac{\tau}{2}(\cos(\omega) - 1)$.

A core element of the discretization in the following section is the approximation of a function $f : [-\tau, 0] \rightarrow \mathbb{C}$ by a truncated series

$$f(\theta) \approx \sum_{k=0}^N \frac{\langle f, \phi_k \rangle}{\|\phi_k\|^2} \phi_k(\theta).$$

If f is analytically continuable to an ellipse in the complex plane (with foci at $-\tau$ and 0), it can be shown that a converging Legendre series, i.e. $\phi_k = P_k$, exists; a result that is generally attributed to Neumann (1862). Analogous to the famous Cauchy–Hadamard formula, this series decays geometrically, i.e. like $O(\rho^N)$, with $0 < \rho < 1$ inversely proportional to the size of the ellipse. In the limiting case where the function is entire, this of course becomes super-geometric decrease. These decay rates were, to the best of our knowledge, first proven for Chebyshev polynomials of the first kind, i.e. $\phi_k = T_k$, by Bernstein (1912: p. 94). A general version for Jacobi polynomials (and thus also Legendre polynomials) can be found as Theorem 9.1.1 in Szegő (1975).

Note that these T_k are generally the preferred choice for the approximation of a (given) function, as fast convergence in N is also guaranteed for functions that are not analytically continuable, but are still sufficiently smooth. In particular, Mastroianni and Szabados (1995: Corollary 2) showed that functions with $m - 1$ absolutely continuous derivatives, and the m th derivative of bounded variation, give m th order algebraic decay of the coefficients of this series. Furthermore, Ahmed and Fisher (1970) presented a way to use the highly efficient fast Fourier transform algorithm to approximate this expansion from function evaluations.

A practical introduction to function approximation and related algorithms, with a focus on Chebyshev series and nodes, can be found in the book of Trefethen (2013). A more extensive treatise on approximation theory is found in Davis (1963). Szegő (1975) gives a thorough overview of the properties of the discussed—and many other—orthogonal polynomials. Finally, application of these polynomials to the approximation of solutions of (partial) differential equations using spectral methods is discussed at length in Boyd (2001).

3.3 THE LANCZOS TAU METHOD

Analogous to what we saw for the space $Z = \mathbb{C}^n \times L^2([-\tau, 0]; \mathbb{C}^n)$ in section 2.3, we can also write the system (3.1) as an abstract Cauchy problem on the space $X = C([-\tau, 0]; \mathbb{C}^n)$, namely

$$\begin{cases} \dot{\xi}_t(0) = A_0 \xi_t(0) + A_1 \xi_t(-\tau) + B \mathbf{u}(t) \\ \dot{\xi}_t(\theta) = \frac{d}{d\theta} \xi_t(\theta), \quad \forall \theta \in [-\tau, 0], \\ \mathbf{y}(t) = C \xi_t(0), \end{cases}$$

with $\xi_t, \dot{\xi}_t \in X$. Motivated by the famous Weierstrass approximation theorem, the idea is, like any spectral method, to approximate ξ_t and $\dot{\xi}_t$ by polynomials of a finite degree.

We start by selecting an inner product $\langle \cdot, \cdot \rangle$ on the space $\mathbb{P}$ of polynomials $[-\tau, 0] \rightarrow \mathbb{C}$. Let $\{\phi_k\}_{k=0}^\infty$ be a degree-graded sequence of orthogonal polynomials with respect to this inner product, as in the previous section. Then, for any N , $\Phi_N = \{\phi_k\}_{k=0}^N$ is an orthogonal basis for $\mathbb{P}_N \subset \mathbb{P}$, the space of polynomials of degree at most N .

In order to arrive at an approximation of the abstract Cauchy problem, and thus (3.1), we will replace the infinite dimensional space of history functions X by the space $\mathbb{P}_N^n$, the space of polynomials of degree at most N that map to $\mathbb{C}^n$. The derivative with respect to θ on the right hand side of the Cauchy problem, reduces the degree of a polynomial by one. On the left hand side, we will thus also have to map an element of $\mathbb{P}_N^n$ to an element of $\mathbb{P}_{N-1}^n$. The idea of Lanczos (1938) was to do so by truncating the highest degree term from the series expansion in Φ_N ,² which was later applied to functional differential equations by Ito and Teglas (1986).

We continue by providing an operator pencil formulation of this approach. The proposed truncation corresponds to the component-wise

²He named the truncated coefficient τ , hence the name ‘tau method’.

The article itself is an interesting window into the early days of numerical analysis. A time when mechanical computers were still in their early stages and a quickly converging series was hence almost as much appreciated for the faster answer it provides as for the reduced strain it puts on the writing hand. Indeed, approximating the solutions of differential equations with the first spectral method—in fact, only a few pages later he introduces a second, ‘the method of selected points’, which we now call collocation—forms only a small part of the article.

Most of his writing is on the ‘economization’ of a power series, illustrating how the number of terms required to obtain a given accuracy from a Chebyshev series is much lower compared to a Taylor series; this is also his primary conclusion: “The practical importance of the [Chebyshev] series is the circumstance that already a small number of terms is able to yield a powerful approximation of the given function.”

orthogonal projection $\mathcal{T}_{N-1}$, with action

$$(\mathcal{T}_{N-1}\xi)_j = (\xi)_j - \frac{\langle (\xi)_j, \phi_N \rangle}{\|\phi_N\|^2} \phi_N, \quad j = 1, \dots, n.$$

Letting ε_θ denote the evaluation functional in θ , $\mathcal{D}$ the component-wise differentiation operator, and $\mathbf{0}$ the zero polynomial, we arrive at the following approximation of (3.1):

$$\begin{aligned} \left(\begin{array}{c} \varepsilon_0 \\ \mathcal{T}_{N-1} \end{array} \right) \frac{d}{dt} \xi_{tN} &= \left(\begin{array}{c} A_0 \varepsilon_0 + A_1 \varepsilon_{-\tau} \\ \mathcal{D} \end{array} \right) \xi_{tN} + \left(\begin{array}{c} B \\ \mathbf{0} \end{array} \right) \mathbf{u}(t), \\ \mathbf{y}_N(t) &= C \varepsilon_0 \xi_{tN}, \end{aligned} \quad (3.3)$$

where $\xi_{tN} \in \mathbb{P}_N^n$.

Alternatively, we can also think of (3.3) as discretizing the abstract Cauchy problem (2.6), i.e. as replacing Z by $\mathbb{P}_N^n$. The relation between solutions of (3.3) and solutions $\mathbf{x}(t)$ of (3.1) and $(\mathbf{z}(t), \zeta)$ of (2.6) can then be described by

$$\mathbf{x}(t) = \mathbf{z}(t) \approx \xi_{tN}(0) \quad \text{and} \quad \zeta_t \approx \xi_{tN}.$$

This is the more natural setting for pseudospectral collocation methods, as we shall see in section 3.5.

Note that the evolution equation (3.3) is in an implicit form. To show that solutions of the corresponding initial value problem exist and are uniquely defined, we must show that this implicit formulation can be recast as a standard delay-free system. To this end, note that $\ker(\mathcal{T}_{N-1}) = \text{span}(I_n \phi_N)$. Furthermore, a basic result on orthogonal polynomials states that their zeroes lie in the interior of their domain (Szegő 1975: Theorem 3.3.1), hence $\phi_N(0) \neq 0$. As a consequence, $\ker(\varepsilon_0) \cap \ker(\mathcal{T}_{N-1}) = \{\mathbf{0}\}$. The inverse

$$\left(\begin{array}{c} \varepsilon_0 \\ \mathcal{T}_{N-1} \end{array} \right)^{-1}$$

is thus a well defined operator from $\mathbb{C}^n \times \mathbb{P}_{N-1}^n$ to $\mathbb{P}_N^n$, and forward solutions of (3.3) are hence uniquely defined.

We shall see in section 3.6 that it is beneficial to leave the discretization in this implicit form. However, note that, depending on which side we multiply by this inverse, we get natural discretizations of both infinitesimal generators from section 2.3. Indeed, $\mathbb{P}_N^n$ is naturally thought of as approximating the state space X and $\mathbb{C}^n \times \mathbb{P}_{N-1}^n$ as approximating Z . If we multiply $\left(\begin{array}{c} A_0 \varepsilon_0 + A_1 \varepsilon_{-\tau} \\ \mathcal{D} \end{array} \right)$ by $\left(\begin{array}{c} \varepsilon_0 \\ \mathcal{T}_{N-1} \end{array} \right)^{-1}$ on the left, we thus get an object resembling $\mathcal{A}_X$. Conversely, by multiplying on the right we get

an object resembling $\mathcal{A}_Z$. The former of the two most closely resembles the original formulation of Ito and Teglas (1986).

For use in implementation, we derive a matrix-vector representation, induced by expressing elements of $\mathbb{P}_N^n$ in the basis Φ_N . That is, $\xi_{tN} = \sum_{k=0}^N \mathbf{x}_{N,k}(t) \phi_k$, with $\mathbf{x}_{N,k}(t) \in \mathbb{C}^n$, $k = 0, \dots, N$. Doing so, we get as block matrix realization of the operators

$$\begin{aligned} [\varepsilon_\theta]_k &= I_n \phi_k(\theta), \quad [\mathcal{D}]_{jk} = I_n \frac{\langle \phi'_k, \phi_j \rangle}{\|\phi_j\|^2}, \quad \text{and} \\ [\mathcal{T}_{N-1}]_{jk} &= I_n \frac{\langle \phi_k, \phi_j \rangle}{\|\phi_j\|^2} = I_n \delta_{jk}, \end{aligned}$$

where $[\cdot]$ signifies the expression in coordinates, $j = 0, \dots, N-1$, and $k = 0, \dots, N$. This then gives the explicit state space realization

$$\begin{aligned} \mathcal{E}_N \dot{\mathbf{x}}_N(t) &= \mathcal{A}_N \mathbf{x}_N(t) + \mathcal{B}_N \mathbf{u}(t), \\ \mathbf{y}_N(t) &= \mathcal{C}_N \mathbf{x}_N(t), \end{aligned} \tag{3.4}$$

with

$$\begin{aligned} \mathcal{E}_N &= \begin{pmatrix} [\varepsilon_0] \\ [\mathcal{T}_{N-1}] \end{pmatrix}, \quad \mathcal{A}_N = \begin{pmatrix} A_0[\varepsilon_0] + A_1[\varepsilon_{-\tau}] \\ [\mathcal{D}] \end{pmatrix}, \\ \mathcal{B}_N &= \begin{pmatrix} B \\ \mathbf{0} \end{pmatrix}, \quad \text{and} \quad \mathcal{C}_N = C[\varepsilon_0]. \end{aligned}$$

Note that the matrix realizations of the operators should generally not be computed explicitly element by element; see the book by Boyd (2001: chapters 3 and 21) and the work of Olver, Slevinsky, et al. (2020) for better approaches. The sparse matrix realizations we shall use in the remainder of this thesis are detailed in section 3.6.

3.4 IMPLICIT RATIONAL APPROXIMATION

Taking the Laplace transform of (3.3) gives

$$\begin{aligned} s \begin{pmatrix} \varepsilon_0 \\ \mathcal{T}_{N-1} \end{pmatrix} \hat{\xi}_{sN} &= \begin{pmatrix} A_0 \varepsilon_0 + A_1 \varepsilon_{-\tau} \\ \mathcal{D} \end{pmatrix} \hat{\xi}_{sN} + \begin{pmatrix} B \\ \mathbf{0} \end{pmatrix} \hat{\mathbf{u}}(s), \\ \hat{\mathbf{y}}_N(s) &= C \varepsilon_0 \hat{\xi}_{sN}, \end{aligned} \tag{3.5}$$

where $\hat{f}(s)$ is the transform of $f(t)$. In this way, we arrive at an expression of the transfer function of (3.3)

$$G_N(s) = C \varepsilon_0 \left(\begin{matrix} s \varepsilon_0 - A_0 \varepsilon_0 - A_1 \varepsilon_{-\tau} \\ s \mathcal{T}_{N-1} - \mathcal{D} \end{matrix} \right)^{-1} \begin{pmatrix} B \\ \mathbf{0} \end{pmatrix}, \tag{3.6}$$

which can of course be thought of as an approximation of $G(s)$. Analogous to an existing result for pseudospectral collocation (see Proposition 3.5.1), we can provide a result on the structure of this approximation.

Proposition 3.4.1. *The transfer function (3.6) satisfies*

$$G_N(s) = C(sI_n - A_0 - A_1 r_N(s, -\tau))^{-1} B,$$

where the function

$$[-\tau, 0] \ni \theta \mapsto r_N(s, \theta)$$

is the unique polynomial of degree N satisfying

$$\begin{cases} r_N(s, 0) = 1, \\ \mathcal{D}r_N(s, \cdot) = s\mathcal{T}_{N-1}r_N(s, \cdot). \end{cases}$$

Furthermore, $r_N(s, \theta)$ is a rational function of s for all θ .

Proof. By expanding $r_N(s, \cdot)$ in a basis, it can easily be seen that it is uniquely defined. The second row of the top expression in (3.5) corresponds to a set of homogeneous equations, each of which has $r_N(s, \cdot)$ as a solution, by the latter's definition. As a consequence, the solutions of this set are of the form $\hat{\xi}_{sN} = r_N(s, \cdot)\mathbf{z}(s)$, with $\mathbf{z}(s) \in \mathbb{C}^n$. Substituting this form in the first row of the top equation leads us to

$$s\mathbf{z}(s) = A_0\mathbf{z}(s) + A_1r_N(s, -\tau)\mathbf{z}(s) + B\hat{\mathbf{u}}(s),$$

while the output equation becomes $\hat{\mathbf{y}}_N(s) = C\mathbf{z}(s)$. The assertions follow from solving for $\mathbf{z}(s)$. $\square$

The effect of approximating (3.1) by (3.3) can thus be interpreted, in the frequency domain, as the effect of approximating the exponential function $s \mapsto e^{-s\tau}$ in (3.2) by the rational function $s \mapsto r_N(s, -\tau)$. Note that this result trivially extends to any linear time-delay system where now e^{-s} is replaced by $r_N(s, \cdot)$. In Proposition 5.3.1, for instance, we shall see the analogous result for DDAEs with discrete delays.

As seen by comparing fig. 3.1 and fig. 1.1, this approximation performs well in an ellipse around zero. As the poles cannot change for different delay-values, we see the zeroes shift around and compensate the presence of poles, losing the symmetry seen for the maximal delay. We shall see in chapter 6 how the use of a spline instead of a single polynomial can partially alleviate this constraint, leading to better performance.

We can also derive a closed form formula for r_N .

Proposition 3.4.2. *The rational function r_N , as defined in Proposition 3.4.1, is given by the explicit expression*

$$r_N(s, \theta) = \frac{\sum_{k=0}^N \phi_N^{(N-k)}(\theta) s^k}{\sum_{k=0}^N \phi_N^{(N-k)}(0) s^k},$$

where $\phi^{(k)}$ is the k th derivative of ϕ .

Proof. Note that, from the definitions of $\mathcal{D}$ and $\mathcal{T}_{N-1}$, we have the equivalent

$$\begin{cases} r_N(s, \theta) = \alpha(s) \left(s - \frac{d}{d\theta}\right)^{-1} \phi_N(\theta), \\ r_N(s, 0) = 1, \end{cases}$$

where $\alpha(s) = \frac{s}{\|\phi_N\|^2} \langle r_N(s, \cdot), \phi_N \rangle$. Using a Neumann series we then find

$$r_N(s, \theta) = \frac{\alpha(s)}{s} \sum_{j=0}^{\infty} \frac{1}{s^j} \frac{d^j}{d\theta^j} \phi_N(\theta) = \alpha(s) \sum_{k=0}^N \frac{1}{s^{N+1-k}} \phi_N^{(N-k)}(\theta).$$

From the boundary condition we have

$$\alpha(s) = \left(\sum_{k=0}^N \frac{1}{s^{N+1-k}} \phi_N^{(N-k)}(0) \right)^{-1}$$

and thus, after multiplying by $\frac{s^{N+1}}{s^{N+1}}$, we find the claimed formula. $\square$

Generally, the coefficients of this expression grow rapidly with N . Alternatively, the definition of r_N also admits the compact representation

$$r_N(s, \theta) = \varepsilon_{\theta} \begin{pmatrix} \varepsilon_0 \\ s\mathcal{T}_{N-1} - \mathcal{D} \end{pmatrix}^{-1} \begin{pmatrix} 1 \\ \mathbf{0} \end{pmatrix}. \quad (3.7)$$

For numerical reasons, solving this system or using the state space realization is usually preferred in implementation.

Later on, we want to use the H^2 -norm of G_N as an approximation of the H^2 -norm of G . An important element is whether G_N converges to G as $N \rightarrow \infty$. As mentioned before, Proposition 3.4.1 allows us to think of G_N as G with the function $s \mapsto e^{-\tau s}$ replaced by $s \mapsto r_N(s, -\tau)$. The quality of the approximation of G by G_N thus depends on the quality of the approximation of the exponential. We conclude this section by proving convergence of G_N , for which we shall need r_N to converge.

Assumption 3.4.3. The basis $\{\phi_k\}_{k=0}^N$ is chosen such that, for any $s \in \mathbb{C}$, the polynomial $[-\tau, 0] \ni \theta \mapsto r_N(s, \theta)$ converges uniformly to $\theta \mapsto e^{\theta s}$ as $N \rightarrow \infty$.

Note that, in the definition of r_N , the condition

$$\mathcal{D}r_N(s, \cdot) = s\mathcal{T}_{N-1}r_N(s, \cdot)$$

is equivalent to requiring $\frac{d}{d\theta} r_N(s, \theta_j) = sr_N(s, \theta_j)$ for $\theta_0, \theta_1, \dots$, and θ_N the zeroes of ϕ_N (Lanczos 1938). The tau method defining $r_N(s, \cdot)$ thus has an equivalent definition using collocation. In fact, this equivalence also exists at the level of the state space discretization, relating the Lanczos tau method to pseudospectral collocation, as we shall see in section 3.5. Because of this relation, Assumption 3.4.3 can often be guaranteed through the following result which is a direct corollary of Proposition 5.1 of Breda, Maset, et al. (2015).

Corollary 3.4.4. *If the Lebesgue constant associated with the zeroes of ϕ_N grows sublinearly, i.e. is of order $o(N)$ as $N \rightarrow \infty$, then Assumption 3.4.3 holds.*

We shall see later that the Legendre polynomials are generally preferred for our purposes. As the Lebesgue constants of Legendre zeroes grow as $O(\sqrt{N})$ (Szegő 1975: bottom of p. 338), we thus have the following from Corollary 3.4.4.

Proposition 3.4.5. *A Legendre basis satisfies Assumption 3.4.3.*

Using the following lemma we will be able to extend the convergence of Assumption 3.4.3 to G_N .

Lemma 3.4.6. *If the pointwise matrix inverse of a matrix valued function $F_N(s)$ is bounded in the spectral norm on some compact subset Ω of the complex plane, i.e. there is some M such that $\|F_N(s)^{-1}\|_2 \leq M$ for all $s \in \Omega$, and $F_N(s)$ converges as $O(f(N))$ to $F(s)$ for $N \rightarrow \infty$ in the Frobenius norm on Ω , then $F_N(s)^{-1}$ converges at the same order to $F(s)^{-1}$ on Ω .*

Proof. Let $E_N(s) = F_N(s) - F(s)$, then this result simply follows from

$$\begin{aligned} F_N(s)^{-1} &= (F(s) + E_N(s))^{-1} \\ &= F(s)^{-1} - F(s)^{-1}E_N(s)F(s)^{-1} + O(\|E_N(s)\|_F^2) \\ &= F(s)^{-1} + O(\|F(s)^{-1}\|_2^2 \|E_N(s)\|_F) \\ &= F(s)^{-1} + O(f(N)). \end{aligned}$$

A derivation for the approximation of the perturbed matrix inverse can be found in Stewart and Sun (1990: section 2.4). The penultimate step is found by vectorizing the middle term. $\square$

We can then prove the following result.

Theorem 3.4.7. *If Assumption 3.4.3 holds, then G_N converges uniformly to G as $N \rightarrow \infty$ on compact sets in $\mathbb{C}$ that do not contain a singularity for N sufficiently large.*

Proof. As a straightforward corollary of Assumption 3.4.3, we have, by combining the largest error and slowest error decay on a compact set, that the function $s \mapsto r_N(s, \theta)$ converges uniformly to $s \mapsto e^{\theta s}$ on that set. Since multiplication with a constant matrix and addition of functions does not change the convergence properties we finally have the required result from Proposition 3.4.1 and Lemma 3.4.6. $\square$

Whilst $r_N(s, \theta)$ is only ever evaluated for $\theta = -\tau$ in the single delay case, nothing in the proof assumed this, and this result, like Proposition 3.4.1, thus readily applies to more general linear time-delay systems. Notably, from Proposition 3.4.5, we have that this theorem holds for a Lanczos tau method based on Legendre polynomials.

3.5 CONNECTION TO PSEUDOSPECTRAL COLLOCATION

As we hinted to in the previous section, there is a direct correspondence between pseudospectral collocation and the Lanczos tau method. Before proving this relation, we outline how the system (2.6), and thus also (3.1), can be discretized using the former method. A more comprehensive treatment is given in Breda, Maset, et al. (2015).

A review of pseudospectral collocation

Given a positive integer N , we consider a mesh Ω of $N + 1$ distinct points in the interval $[-\tau, 0]$, namely

$$\Omega = \{\theta_k : k = 0, \dots, N\}, \quad (3.8)$$

where

$$-\tau \leq \theta_0 < \dots < \theta_{N-1} < \theta_N = 0.$$

This allows us to replace the continuous space Z , defined in (2.5), with the space X_N of discrete functions, defined on the mesh Ω , in the abstract Cauchy problem (2.6), i.e. any tuple $(\mathbf{z}, \zeta) \in Z$ is approximated by a block vector $\mathbf{x}_N \in X_N$, with

$$\mathbf{x}_{N,k} = \zeta(\theta_k), \quad k = 0, \dots, N-1, \quad \text{and} \quad \mathbf{x}_{N,N} = \mathbf{z}.$$

Let $\mathcal{P}_{\mathbf{x}_N}$ be the unique $\mathbb{C}^n$ -valued interpolating polynomial of degree at most N , satisfying

$$(\mathcal{P}_{\mathbf{x}_N})(\theta_k) = \mathbf{x}_{N,k}, \quad k = 0, \dots, N.$$

This way we can approximate the infinitesimal generator $\mathcal{A}_Z$ by the finite-dimensional operator $\mathcal{A}_N : X_N \rightarrow X_N$, defined by

$$\begin{cases} (\mathcal{A}_N \mathbf{x}_N)_k = (\mathcal{D} \mathcal{P}_{\mathbf{x}_N})(\theta_k), & k = 0, \dots, N-1, \\ (\mathcal{A}_N \mathbf{x}_N)_N = (A_0 \varepsilon_0 + A_1 \varepsilon_{-\tau})(\mathcal{P}_{\mathbf{x}_N}). \end{cases}$$

Note that in doing so, we implicitly enforce the boundary condition of the ‘head-tail’ representation, namely $(\mathcal{P}_{\mathbf{x}_N})(0) = \mathbf{x}_{N,N} = \mathbf{z}$, where $\mathcal{P}_{\mathbf{x}_N}$ can be seen as the approximation of ζ .

Using the Lagrange representation of $\mathcal{P}\mathbf{x}_N$,

$$(\mathcal{P}\mathbf{x}_N)(\theta) = \sum_{k=0}^N \mathbf{x}_{N,k} \ell_k(\theta),$$

where the Lagrange polynomials ℓ_k are those real-valued polynomials of degree N satisfying $\ell_k(\theta_j) = \delta_{jk}$, with δ_{jk} the usual Kronecker delta, one can get an explicit matrix expression

$$\mathcal{A}_N = \begin{pmatrix} [\underline{\mathcal{D}}] \\ \mathbf{a} \end{pmatrix},$$

where $[\underline{\mathcal{D}}]$ consists of the first N block rows of the $(N+1) \times (N+1)$ differentiation matrix

$$[\underline{\mathcal{D}}]_{jk} = I_n \ell'_k(\theta_j),$$

and $\mathbf{a}$ is a block row vector with

$$\mathbf{a}_k = A_0 \ell_k(0) + A_1 \ell_k(-\tau),$$

where $j = 0, \dots, N$ and $k = 0, \dots, N$.

In the same way we can approximate the operators $\mathcal{B}$ and $\mathcal{C}$ from the abstract Cauchy problem (2.6) by

$$\mathcal{B}_N = (\mathbf{0} \quad \cdots \quad \mathbf{0} \quad B^T)^T \quad \text{and} \quad \mathcal{C}_N = (\mathbf{0} \quad \cdots \quad \mathbf{0} \quad C).$$

As such, we arrive at a finite-dimensional approximation of (3.1)

$$\begin{aligned} \dot{\mathbf{x}}_N(t) &= \mathcal{A}_N \mathbf{x}_N(t) + \mathcal{B}_N \mathbf{u}(t), \\ \mathbf{y}_N(t) &= \mathcal{C}_N \mathbf{x}_N(t), \end{aligned} \tag{3.9}$$

and we can thus also approximate the transfer function (3.2) by

$$J_N(s) = \mathcal{C}_N (sI_{n(N+1)} - \mathcal{A}_N)^{-1} \mathcal{B}_N. \tag{3.10}$$

The following result on the structure of this approximation was proved by Gumussoy and Michiels (2010), which is the counterpart to Proposition 3.4.1.

Proposition 3.5.1. *The transfer function (3.10) satisfies*

$$J_N(s) = C(sI_n - A_0 - A_1 p_N(s, -\tau))^{-1} B,$$

where the function

$$[-\tau, 0] \ni \theta \mapsto p_N(s, \theta)$$

is the unique polynomial of degree N satisfying

$$\begin{cases} p_N(s, 0) = 1, \\ p'_N(s, \theta_k) = sp_N(s, \theta_k), \quad k = 0, \dots, N-1. \end{cases}$$

Furthermore, $p_N(s, \theta)$ is a rational function of s for all θ .

A common choice of mesh points in the literature consists of scaled and shifted Chebyshev extremal points, that is

$$\theta_k = -\frac{\tau}{2} \left(\cos\left(\frac{\pi k}{N}\right) + 1 \right), \quad k = 0, \dots, N.$$

The choice of this mesh is motivated by the resulting fast convergence of the eigenvalues of $\mathcal{A}_N$ to the corresponding characteristic roots of (3.1). More specifically, in Theorem 5.1 of Breda, Maset, et al. (2015), it is proved that super-geometric accuracy, i.e. approximation error $O(N^{-N})$, is obtained using these nodes.

Note that, as for the Lanczos tau method, we can rewrite $\mathbf{a}$ in terms of a block vector expression of the evaluation functionals, namely

$$\mathbf{a} = A_0[\varepsilon_0] + A_1[\varepsilon_{-\tau}].$$

Similarly, we have for $\mathcal{C}_N$

$$C \begin{pmatrix} \mathbf{0} & \cdots & \mathbf{0} & I_n \end{pmatrix} = C[\varepsilon_0].$$

An equivalence theorem

As we saw in section 3.4, Lanczos tau methods for differential equations are well known to correspond to collocation in the zeroes of the truncated polynomial (Lanczos 1938). A similar intimate connection between pseudospectral collocation and the approximation scheme introduced in this chapter is thus expected; the following result holds.

Theorem 3.5.2. *Assume that the non-zero mesh points of Ω , as defined in (3.8), are chosen as the zeroes of ϕ_N , that is $\theta_N = 0$ and*

$$\phi_N(\theta_k) = 0, \quad k = 0, \dots, N-1,$$

then, for all N , the finite-dimensional approximations (3.9) and (3.3) are equivalent, and thus $J_N(s) = G_N(s)$.

Proof. Note that (3.9) can also be derived from the relations

$$\begin{cases} \varepsilon_0 \frac{d}{dt} \eta_{tN} = (A_0 \varepsilon_0 + A_1 \varepsilon_{-\tau}) \eta_{tN} + B \mathbf{u}(t), \\ \varepsilon_{\theta_k} \frac{d}{dt} \eta_{tN} = \varepsilon_{\theta_k} \mathcal{D} \eta_{tN}, \quad k = 0, \dots, N-1, \\ \mathbf{y}_N(t) = C \varepsilon_0 \eta_{tN}, \end{cases}$$

with $\eta_{tN} \in \mathbb{P}_N^n$, by expressing η_{tN} in the Lagrange basis with respect to (3.8). Additionally, using the definition of $\mathcal{T}_{N-1}$, the bottom row of the first equation in (3.3) reads

$$\left(\frac{d}{dt}\xi_{tN}\right)_j - \left\langle \left(\frac{d}{dt}\xi_{tN}\right)_j, \phi_N \right\rangle \frac{\phi_N}{\|\phi_N\|^2} = (\mathcal{D}\xi_{tN})_j, \quad j = 1, \dots, n,$$

which, under the above conditions on Ω , implies

$$\varepsilon_{\theta_k} \frac{d}{dt}\xi_{tN} = \varepsilon_{\theta_k} \mathcal{D}\xi_{tN}, \quad k = 0, \dots, N-1.$$

Hence, the conditions imposed on the evolution of the polynomial ξ_{tN} , as a function of t , imply the conditions imposed on the evolution of the polynomial η_{tN} . But since each set of conditions uniquely defines the flow, they must be equivalent. $\square$

This connection allows one to reuse the large number of results that were developed for collocation-based methods, as we have already, implicitly, done for Corollary 3.4.4. In particular, we recover the aforementioned super-geometric convergence of the characteristic roots when using a Legendre basis or a Chebyshev basis of the first or second kind, as this corresponds to using the zeroes of P_N , T_N , or U_N , respectively, as the non-zero collocation points (Breda, Maset, et al. 2015: Theorem 5.1). Such a result is not unexpected as the conditions in Propositions 3.4.1 and 3.5.1, in the limit, define the function $\theta \mapsto e^{s\theta}$. These methods are both grounded in the approximation of the exponential by a polynomial, obtained through collocation or the truncation of a series, respectively. From the discussion in section 3.2, and analogous results for interpolation through Chebyshev and Legendre nodes, super-geometric convergence is observed for this function as it is entire. Indeed on figs. 3.2 and 3.3 we see that this fast convergence indeed also holds for appropriate Lanczos tau methods.

3.6 SPARSE, NESTED DISCRETIZATIONS

A major advantage of the Lanczos tau method, particularly in the implicit form we use in eq. (3.3), is that it readily admits highly sparse and self-nesting matrices. Indeed, from a degree-graded orthogonal sequence $\{\phi_k\}_{k=0}^\infty$, we can trivially form a sequence of nested bases $\Phi_0 \subset \Phi_1 \subset \dots$ for $\mathbb{P}_0 \subset \mathbb{P}_1 \subset \dots$ respectively. As a consequence, it is easy to construct a matrix realization of (3.3) which is also nested, in the sense that the resulting matrices for degree N_1 are submatrices of those at degree N_2 , for $N_1 < N_2$. In fact, the matrices in (3.4) have this property.




Figure 3.1: Domain-colourings of $r_N(s, -\tau)$ (left) and $r_N(s, -0.675)$ (right) for $\phi_N = P_N$, $\tau = 1$, and $N = 6$. Note that the location of the poles is identical on both figures.




Figure 3.2: The characteristic roots of (3.1), with $A_0 = \begin{pmatrix} -2 & 1 \\ 3 & -8 \end{pmatrix}$, $A_1 = \begin{pmatrix} -1 & -1 \\ -1 & -1 \end{pmatrix}$, $\tau = 1$, and $B = C = I_2$, and those of two low degree Lanczos tau discretizations using a Legendre basis. The remaining roots of the discretizations lie to the left of -4 and close to the real line, poorly approximating any of the true roots.




Figure 3.3: Convergence rate of the rightmost root of fig. 3.2. Other roots show similar curves, but requiring increasingly higher degrees for the same accuracy the farther away from zero.

This nesting of state space representations can be exploited in Krylov algorithms for characteristic roots computation and model reduction of time delay systems of high dimension, as in the infinite Arnoldi method introduced by Jarlebring, Meerbergen, et al. (2010). This leads to far cheaper and far more flexible methods, as this nesting allows the reuse of previous computations whilst adaptively changing the discretization degree.

As an example, the self-nesting discretization initially derived for this purpose in the above article is based on collocation in zero and the zeroes of U_N . By Theorem 3.5.2 we can thus recast this in the framework of section 3.3, as this corresponds, up to a basis transform, to the choice $\Phi_N = \{U_k\}_{k=0}^N$.

Furthermore, note that when expressing (3.3) in a basis, it is not necessary to use Φ_N as basis for the input side of the operators (on the contrary, using Φ_N to represent the output yields neater representations of $\mathcal{T}_{N-1}$ and is thus generally preferred). For Chebyshev polynomials we have, from elementary trigonometric identities, the relations

$$\frac{d}{d\theta} T_k(\theta) = \frac{2}{\tau} k U_{k-1}(\theta)$$

and

$$T_k(\theta) = \begin{cases} \frac{1}{2}(U_k(\theta) - U_{k-2}(\theta)), & k \geq 0, \\ U_0(\theta), & k = 0, \end{cases}$$




Figure 3.4: The non-zeroes of the Lanczos tau discretization of a single delay system using an ultraspherical matrix realization based on Chebyshev polynomials for $N = 10$. On the left are the non-zeroes of $\mathcal{E}_N$, on the right those of $\mathcal{A}_N$. Each dot represents an $n \times n$ matrix.

where $U_k(\theta) = 0$ for $k < 0$. We can thus use the freedom in the choice of bases to represent the derivative as a single non-zero diagonal, whilst making $\mathcal{T}_{N-1}$ only tridiagonal, by taking $\Phi_N = \{U_k\}_{k=0}^N$ and representing the input side of the operators in the basis $\{T_k\}_{k=0}^N$. As seen on fig. 3.4, this indeed gives the highly sparse representation we expected, with $O(N)$ non-zeroes instead of $O(N^2)$, yielding even more computational gains. In fact, the matrix representation of Jarlebring, Meerbergen, et al. (2010) corresponds to this choice, up to a scaling of the rows.

This approach is also taken for the solution of differential equations in the ultraspherical methods of Olver and Townsend (2013), using similar properties of Gegenbauer polynomials to approximate higher order derivatives. In particular, these polynomials satisfy the relations (see Szegő 1975: eq. 4.7.14)

$$\frac{d}{d\theta} C_k^\lambda(\theta) = \frac{4}{\tau} \lambda C_{k-1}^{\lambda+1}(\theta) \quad (3.11)$$

and (see Szegő 1975: eq. 4.7.29)

$$C_k^\lambda(\theta) = \frac{\lambda}{\lambda+k} (C_k^{\lambda+1}(\theta) - C_{k-2}^{\lambda+1}(\theta)), \quad (3.12)$$

where $C_k^\lambda(\theta) = 0$ for $k < 0$. Hence, similar sparsity patterns to fig. 3.4 can be obtained for $\Phi_N = \{C_k^\lambda\}_{k=0}^N$, by representing the input side of the operators in the basis $\{C_k^{\lambda-1}\}_{k=0}^N$.

Apart from sparsity, another important reason for such a matrix representation are its improved numerical properties (Olver and Townsend 2013: section 4). We observe similar benefits when discretizing DDEs.




Figure 3.5: The non-zeroes of the Lanczos tau discretization of a single delay system using an ultraspherical matrix realization based on Legendre and generalized Gegenbauer polynomials with $\lambda = -\frac{1}{2}$ for $N = 10$. On the left are the non-zeroes of $\mathcal{E}_N$, on the right those of $\mathcal{A}_N$. Each dot represents an $n \times n$ matrix.

All numerical experiments involving a Lanczos tau method in this thesis were conducted using such a matrix pencil, if not the error in the tail of e.g. figs. 3.3 and 4.1 would rise slightly with N .

We shall see later that a Lanczos tau method based on Legendre polynomials will be the most beneficial when approximating the H^2 -norm. Hence, we would like to construct a similar sparse discretization for this choice. However, this would require us to use the basis $\{C_k^{-\frac{1}{2}}\}_{k=0}^N$ on the input side. As we saw in section 3.2, the weighted inner product of the Gegenbauer polynomials is only integrable for $\lambda > -\frac{1}{2}$. Luckily, the Gegenbauer polynomials can still be generalized to $\lambda = \frac{1}{2} - n$, for $n > 0$, by using their three term recurrence relation

$$\begin{aligned} kC_k^\lambda(\theta) &= 2(k + \lambda - 1)\left(\frac{2}{\tau}\theta + 1\right)C_{k-1}^\lambda(\theta) \quad \text{for } k \geq 2, \\ C_0^\lambda(\theta) &= 1, \quad \text{and} \quad C_1^\lambda(\theta) = 2\lambda\left(\frac{2}{\tau}\theta + 1\right), \end{aligned}$$

instead, preserving most of the relevant properties (Álvarez de Morales et al. 1998). In particular they satisfy (3.11) (Álvarez de Morales et al. 1998: Proposition 1 (iii)) and as they have the same three term recurrence relation we can use the same derivation as Szegő (1975) to obtain (3.12).

Whilst they are not orthogonal with respect to a standard weighted integral, they are orthogonal with respect to a family of Sobolev type inner products, one of which is

$$\langle \phi, \psi \rangle_W = \phi(-\tau)\psi(-\tau) + \phi(0)\psi(0) + \int_{-\tau}^0 \phi'(\theta)\psi'(\theta) d\theta,$$

hence, importantly, they still form bases for the spaces of polynomials of fixed degree. At first glance this inner product is a downside, however, to construct a matrix realization we never need to use the inner product directly. Better yet, the resulting polynomials need not satisfy Theorem 3.3.1 of Szegő (1975), and can thus have zeroes in the end points of the interval. Indeed, we have (see Álvarez de Morales et al. 1998: Proposition 2 (iv))

$$C_k^{-\frac{1}{2}}(-\tau) = C_k^{-\frac{1}{2}}(0) = 0 \quad \text{for } k \geq 2.$$

Thus, not only are the matrix realizations of $\mathcal{D}$ and $\mathcal{T}_{N-1}$ highly sparse, so too are the realizations of $\varepsilon_{-\tau}$ and ε_0 . Notably, as seen on fig. 3.5, the matrix realization of $\mathcal{E}_N$ is thus always block tridiagonal and, in the single delay case, so too the matrix realization of $\mathcal{A}_N$. For more general functionals than $A_0\varepsilon_0 + A_1\varepsilon_{-\tau}$, the matrix realization of $\mathcal{A}_N$ will of course generically have a dense top row. Still, this is an improvement, even compared to the infinite Arnoldi matrices, as solving a system involving a tridiagonal $\mathcal{E}_N$ is cheaper than for one that is only nearly tridiagonal.

3.7 RELATION TO PADÉ APPROXIMATION

The Padé approximant of e^s near zero and the Legendre polynomials are well known to be linked, as shown by Ahmad (1998). In the context of the approximation of time-delay systems using the Lanczos tau method, such a connection has also been demonstrated by Ito (1985). We start by providing an alternative argument, working directly at the level of the transfer function.

Theorem 3.7.1. *For the choice $\Phi_N = \{P_k\}_{k=0}^N$, the rational function*

$$s \mapsto r_N(s, -\tau),$$

as in Proposition 3.4.1, is an (N, N) Padé approximant of $e^{-\tau s}$ near zero.

Proof. From Proposition 3.4.1 we know that $r_N(s, -\tau)$ is a rational function of (at most) type (N, N) . The defining property of a Padé approximant of this type, and thus what we must show, is that the first $2N + 1$ moments match those of the exponential at zero, i.e.

$$\left[\frac{d^n}{ds^n} r_N(s, -\tau) \right]_{s=0} = \left[\frac{d^n}{ds^n} e^{-\tau s} \right]_{s=0} = (-\tau)^n \quad \forall n \leq 2N.$$

As the operators involved in (3.7) linearly map between finite-dimensional spaces, we can apply the analogue of the derivative of an inverse

matrix, yielding

$$\begin{aligned}
 & \left[\frac{d^n}{ds^n} r_N(s, -\tau) \right]_{s=0} \\
 &= \left[\frac{d^n}{ds^n} \varepsilon_{-\tau} \left(s \mathcal{T}_{N-1}^{\varepsilon_0} - \mathcal{D} \right)^{-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right]_{s=0} \\
 &= \left[-\frac{d^{n-1}}{ds^{n-1}} \varepsilon_{-\tau} \left(s \mathcal{T}_{N-1}^{\varepsilon_0} - \mathcal{D} \right)^{-1} \begin{pmatrix} \mathbf{0}^* \\ \mathcal{T}_{N-1} \end{pmatrix} \left(s \mathcal{T}_{N-1}^{\varepsilon_0} - \mathcal{D} \right)^{-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right]_{s=0} \\
 &= \left[n! (-1)^n \varepsilon_{-\tau} \left[\left(s \mathcal{T}_{N-1}^{\varepsilon_0} - \mathcal{D} \right)^{-1} \begin{pmatrix} \mathbf{0}^* \\ \mathcal{T}_{N-1} \end{pmatrix} \right]^n \left(s \mathcal{T}_{N-1}^{\varepsilon_0} - \mathcal{D} \right)^{-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix} \right]_{s=0} \\
 &= n! \varepsilon_{-\tau} \left[\begin{pmatrix} \varepsilon_0 \\ \mathcal{D} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{0}^* \\ \mathcal{T}_{N-1} \end{pmatrix} \right]^n \begin{pmatrix} \varepsilon_0 \\ -\mathcal{D} \end{pmatrix}^{-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix},
 \end{aligned}$$

with $\mathbf{0}^* f = 0$. Let $\mathcal{M}_N = \begin{pmatrix} \varepsilon_0 \\ \mathcal{D} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{0}^* \\ \mathcal{T}_{N-1} \end{pmatrix}$ and note that

$$(\mathcal{M}_N f)(\theta) = \int_0^\theta \left[f(\xi) - \frac{\langle f, P_N \rangle}{\|P_N\|^2} P_N(\xi) \right] d\xi.$$

Additionally, for $f_0 = \begin{pmatrix} \varepsilon_0 \\ -\mathcal{D} \end{pmatrix}^{-1} \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ we have $f_0(\theta) = 1$, hence

$$\left[\frac{d^n}{ds^n} r_N(s, -\tau) \right]_{s=0} = n! \varepsilon_{-\tau} \mathcal{M}_N^n f_0.$$

It thus remains to show that, for $n \leq 2N$,

$$n! \varepsilon_{-\tau} \mathcal{M}_N^n f_0 = (-\tau)^n.$$

The operator $\mathcal{M}_N$ maps from $\mathbb{P}_N$ to $\mathbb{P}_N$. If we embed this space in the space $\mathbb{P}$ of polynomials of arbitrary degree, we can split $\mathcal{M}_N = \mathcal{I} + \mathcal{K}_N$, where $(\mathcal{I}f)(\theta) = \int_0^\theta f(\xi) d\xi$ and $(\mathcal{K}_N f)(\theta) = -\int_0^\theta \frac{\langle f, P_N \rangle}{\|P_N\|^2} P_N(\xi) d\xi$. We know that $\frac{2}{\tau}(2k+1)P_k(\theta) = \frac{d}{d\theta}(P_{k+1}(\theta) - P_{k-1}(\theta))$ for $k \geq 1$ (consequence of Szegő 1975: eq. 4.7.29), hence, by the fundamental theorem of calculus and the fact that $P_{k+1}(0) = P_{k-1}(0) = 1$,

$$\mathcal{K}_N f = -\frac{\langle f, P_N \rangle}{\frac{2}{\tau}(2N+1)\|P_N\|^2} (P_{N+1} - P_{N-1}).$$

As P_N is orthogonal to all polynomials of degree less than N , we have

$$\mathcal{K}_N g = \mathbf{0}, \quad \forall g \in \mathbb{P}_{N-1}. \quad (3.13)$$

Note that $(\mathcal{I} + \mathcal{K}_N)^n$ is the sum of all possible sequences of $\mathcal{I}$ and $\mathcal{K}_N$ of length n . As a consequence of (3.13), we know that any subsequence of operators at most N long and containing $\mathcal{K}_N$, gives the zero function when applied to f_0 . Thus $(\mathcal{I} + \mathcal{K}_N)^n f_0$ consists of $(\mathcal{I}^n f_0)(\theta) = \frac{\theta^n}{n!}$ and terms of the form $\mathcal{I}^m \mathcal{K}_N f$, with $m < n - N$. Applying the same integration rule for P_k from before, we get the following forms for $\mathcal{I}^m \mathcal{K}_N f$.

m	expression
0	$\alpha(P_{N+1} - P_{N-1})$
1	$\beta(P_{N+2} - P_N) + \gamma(P_N - P_{N-2})$
$\vdots$	$\vdots$
$N-1$	$\chi(P_{2N} - P_{2N-2}) + \cdots + \omega(P_2 - P_0)$

As $P_k(-\tau) = (-1)^k$ (Szegő 1975: eq. 4.1.4), we get

$$\varepsilon_{-\tau} \mathcal{I}^m \mathcal{K}_N f = 0, \quad \forall m < N.$$

Since $n \leq 2N$ we have $m < 2N - N = N$, thus the last result holds, hence

$$n! \varepsilon_{-\tau} \mathcal{M}_N^n f_0 = n! \varepsilon_{-\tau} \mathcal{I}^n f_0 = (-\tau)^n, \quad \forall n \leq 2N,$$

which is precisely what had to be shown. $\square$

This result and Proposition 3.4.1 then readily give the following.

Corollary 3.7.2. *For the choice $\Phi_N = \{P_k\}_{k=0}^N$, the Lanczos tau discretization (3.3) is a state space representation of the transfer function of the original, single delay, system (3.2) where the exponential $s \mapsto e^{-\tau s}$ is replaced by a Padé approximant of type (N, N) around zero.*

Of course, an analogous result holds for pseudospectral collocation, via Theorem 3.5.2.

Finally, we can use this link and Proposition 3.4.2 to give an explicit expression for the resulting Padé approximant.

Corollary 3.7.3. *The (N, N) Padé approximant of $s \mapsto e^{-\tau s}$ around zero has as rational expression*

$$\frac{\sum_{k=0}^N P_N^{(N-k)}(-\tau) s^k}{\sum_{k=0}^N P_N^{(N-k)}(0) s^k}.$$

Note that this expression of the Padé approximant is identical to expression (3.4) in Ahmad (1998). This correspondence can thus be used to give an alternative proof of Theorem 3.7.1, be it more indirect than the approach presented here.

Yet another alternative proof tells us a little more about r_N for other choices of Gegenbauer polynomials, which, as they include both Legendre and Chebyshev polynomials, are the most relevant in practice. For this family of bases we have the following connection to confluent hypergeometric functions.

Proposition 3.7.4. *For the choice $\Phi_N = \{C_k^\lambda\}_{k=0}^N$, the rational function*

$$s \mapsto r_N(s, -\tau),$$

as in Proposition 3.4.1, is given by

$$\frac{{}_1F_1^{(N)}(-(N + \lambda - \frac{1}{2}); -2(N + \lambda - \frac{1}{2}); -\tau s)}{{}_1F_1^{(N)}(-(N + \lambda - \frac{1}{2}); -2(N + \lambda - \frac{1}{2}); \tau s)},$$

with ${}_1F_1^{(N)}(a; b; z) = \sum_{k=0}^N \frac{(a)_k z^k}{(b)_k k!}$ the truncated confluent hypergeometric function, i.e. ${}_1F_1 = {}_1F_1^{(\infty)}$, and $(a)_k = a(a+1) \cdots (a+k-1)$ the rising factorial.

Proof. We start by proving the expression for the denominator of r_N , given, up to scaling, by

$$q(s) = \sum_{k=0}^N C_N^{\lambda(k)}(0) s^{N-k}.$$

From Szegő (1975: eq. 4.7.6) we have

$$\begin{aligned} C_N^\lambda(s) &= \frac{(2\lambda)_N}{N!} {}_2F_1(-N, N + 2\lambda; \lambda + \frac{1}{2}; -\tau s) \\ &= \frac{(2\lambda)_N}{N!} \sum_{k=0}^N \frac{(-N)_k (N + 2\lambda)_k}{(\lambda + \frac{1}{2})_k} \frac{(-\tau s)^k}{k!}, \end{aligned}$$

hence

$$q(s) = \frac{(2\lambda)_N}{N!} \sum_{k=0}^N \frac{(-\tau)^k (-N)_k (N + 2\lambda)_k}{(\lambda + \frac{1}{2})_k} s^{N-k}.$$

After normalization, we find

$$\begin{aligned} \frac{q(s)}{q(0)} &= \sum_{k=0}^N \frac{(-1)^k (-N)_k (N + 2\lambda)_k (\lambda + \frac{1}{2})_N}{(-1)^N (-N)_N (N + 2\lambda)_N (\lambda + \frac{1}{2})_k} (\tau s)^{N-k} \\ &= \sum_{k=0}^N \frac{(\lambda + \frac{1}{2} + k)_{N-k}}{(N + 2\lambda + k)_{N-k}} \frac{(\tau s)^{N-k}}{(N - k)!}, \end{aligned}$$

where we used the easily verified identities $\frac{(x)_N}{(x)_k} = (x + k)_{N-k}$, for $k \leq N$, and $(-N)_k = (-1)^k \frac{N!}{(N-k)!}$. Substituting $N - k$ for k and using the identity $(x - k)_k = (-1)^k (-x + 1)_k$ gives

$$\frac{q(s)}{q(0)} = {}_1F_1^{(N)}(-(N + \lambda - \frac{1}{2}); -2(N + \lambda - \frac{1}{2}); \tau s).$$

Using an analogous argument we find the required result for the denominator, concluding the proof. $\square$

The choice of Legendre polynomials corresponds to picking $\lambda = \frac{1}{2}$ (see section 3.2). Furthermore, note that $(-N)_k = 0$ for all $k > N$. We thus have, from the above proposition, that

$$r_N(s, -\tau) = \frac{{}_1F_1(-N; -2N; -\tau s)}{{}_1F_1(-N; -2N; \tau s)},$$

which is a well known expression for the (N, N) Padé approximant of $s \mapsto e^{-\tau s}$ around zero (Perron 1913: p. 431). Another way to interpret Proposition 3.7.4 is thus that all other Gegenbauer polynomials give some sort of perturbed Padé approximation.

Finally, note that this implies that for any choice of Gegenbauer polynomial, we have

$$|r_N(i\omega, -\tau)| = 1, \quad \forall \omega \in \mathbb{R}.$$

As we shall see in the next chapter (see in particular Proposition 4.1.2), this appears to be beneficial when computing the H^2 -norm.

Whilst the results preceding this section straightforwardly extend to more general settings such as multiple or distributed delays, those of this section are unique to the single delay case. Indeed, if we replace $A_0\varepsilon_0 + A_1\varepsilon_{-\tau}$ by some other bounded linear functional in (3.3), we see by the reasoning of Proposition 3.4.1 that $\theta \mapsto r_N(s, \theta)$ will be evaluated in multiple points, not just $-\tau$. From Proposition 3.4.2, however, we know that the poles of r_N are independent of θ , hence an analogue of Corollary 3.7.2 cannot hold, as the (N, N) Padé approximants of e.g. $s \mapsto e^{-s}$ and $s \mapsto e^{-2s}$ do not have common poles.

3.8 PRESERVATION OF STABILITY

To conclude this chapter we show that the link to Padé approximation we established in the previous section allows us to prove preservation of stability when using Legendre polynomials in this, be it somewhat trivial, single delay setting. The proof is in part inspired by the argument of Theorem 5.3 of Ito (1985), of which it is a special case (his proof includes distributed delay with a sufficiently well-behaved integration kernel). In later chapters we shall extend this argument to DDAES, where it is novel.

Theorem 3.8.1. *For a single delay system (3.1), the discretization (3.3) preserves stability for N sufficiently large, when using a Legendre basis.*

Proof. Let $\Delta_N(s) = sI_n - A_0 - A_1 r_N(s, -\tau)$. From Theorem 3.7.1 we know that, for this choice of basis, $r_N(s, -\tau)$ is the N th diagonal Padé approximant of $s \mapsto e^{-\tau s}$. Birkhoff and Varga (1965: Lemma 8) showed of these approximations that they are bounded in magnitude by one on the closed right half-plane, i.e. $|r_N(s, -\tau)| \leq 1$ for all $s \in \mathbb{C}_{\geq 0} = \{z \in \mathbb{C} : \operatorname{Re}(z) \geq 0\}$.

All characteristic roots λ of the discretization in this half-plane satisfy $\det(\Delta_N(\lambda)) = 0$ and thus

$$|\lambda| \leq \|A_0 + A_1 r_N(s, -\tau)\|_2.$$

The bound on r_N then implies that the zeroes of $\det(\Delta_N(\lambda))$ that lie in $\mathbb{C}_{\geq 0}$ are bounded in magnitude by some finite constant M and are thus constrained to a hemicircle Ω of radius M .

As the original system is stable, it has no characteristic roots in Ω . From the argument for Theorem 3.4.7, we know that $\Delta_N(s)$ converges on compact sets. Hence, for N sufficiently large, the discretized system cannot have any characteristic roots in Ω and thus has none in the entire closed right half-plane. $\square$

It would be interesting to extend this result to other bases, if possible, which to our knowledge is an open problem. It is sufficient to show that $r_N(s, -\tau)$ is bounded on the right half-plane by some constant uniform in N . For other Gegenbauer polynomials, a good starting point could be Proposition 3.7.4 and the resulting interpretation of $r_N(s, -\tau)$ as a perturbed Padé approximant. However,³ such a bound clearly does not exist for all Gegenbauer polynomials, as can be seen on fig. 3.6. There is thus also the possibility that Legendre polynomials are in some sense ‘special’. In part for this reason, together with the sparse matrix realization we saw in section 3.6 and the better convergence characteristics we shall see in the next chapter, we recommend to stick to a Legendre based discretization.

3.9 SUMMARY

In this chapter, we developed a framework of operator pencil formulations of the Lanczos tau method (3.3) for the discretization of linear systems with a single state delay (section 3.3). Throughout, we saw that this interpretation in terms of actions on polynomials aids theoretical derivations.

³This should, at least in part, be taken as a warning. I spent a great many months on this problem and the present text unfortunately represents the limits of my understanding.

We showed equivalence to rational approximation in the frequency domain and provided an explicit expression of the resulting rational function in Propositions 3.4.1 and 3.4.2, respectively. If the underlying approximation of the exponential converges (for which we can often use Corollary 3.4.4), the transfer function converges as well (Theorem 3.4.7).

The Lanczos tau method is closely related to pseudospectral collocation as seen in Theorem 3.5.2. This means we can reuse many results known for that discretization, such as the fast convergence of the characteristic roots (using Theorem 5.1 of Breda, Maset, et al. (2015)).

In section 3.6 we saw how Lanczos tau methods naturally lead to nested and sparse matrix realizations, which can be exploited in Krylov methods, allowing for improved performance for the eigenvalue problem (Jarlebring, Meerbergen, et al. 2010). In the same section we introduced, using a similar technique, a highly sparse matrix realization of a Lanczos tau method using a Legendre basis, which will be our primary method in the following chapters.

Particularly interesting is the equivalence to Padé approximation for the choice of a shifted Legendre basis (Theorem 3.7.1); our proof of which strongly relied on the interpretation in terms of operations on polynomials. This connection finally allowed us to prove preservation of stability when discretizing a single delay system using a tau method based on a Legendre basis (Theorem 3.8.1).

Although we pursued neither multiple nor distributed delays in this chapter, the development of section 3.3, in particular Propositions 3.4.1 and 3.4.2, the benefits discussed in section 3.6, and the equivalence Theorem 3.5.2, can easily be extended to these cases. In fact, any bounded linear functional (in the sense of L in (2.4)) can be substituted for $A_0\varepsilon_0 + A_1\varepsilon_{-\tau}$ in (3.3). As already noted at the end of section 3.7, a link to Padé approximation cannot be obtained from such an extension. Similarly, the proof of preservation of stability (Theorem 3.8.1) would need to be adapted (see Theorem 5.3 of Ito (1985) for a version supporting distributed delay).




Figure 3.6: Domain-colouring of $r_N(s, -\tau)$ for $\phi_N = C_N^{10}$, $\tau = 1$, and $N = 15$. Note the four poles to the right of the imaginary axis.

4 *Computing the H^2 -norm of a single delay system*

In the previous chapter, we introduced the Lanczos tau method to obtain a delay-free approximation of a time-delay system. In this and the remaining chapters, we will aim to approximate the H^2 -norm of a time-delay system by using this discretization. Remember from section 2.5 that this norm, for an exponentially stable system, is given by

$$\|G\|_{H^2} = \left(\frac{1}{2\pi} \int_{-\infty}^{\infty} \|G(i\omega)\|_F^2 d\omega \right)^{\frac{1}{2}},$$

where i is the imaginary unit and $\|A\|_F = \sqrt{\text{tr}(AA^*)}$ the Frobenius norm.

It is well known that we can compute this norm for a delay-free system by solving a Lyapunov equation (see section 4.2). To approximate the H^2 -norm of a time-delay system, we shall thus use a discretization of the system, with transfer function G_N , as a delay-free proxy for the original, with transfer function G , i.e. we will approximate

$$\|G\|_{H^2} \approx \|G_N\|_{H^2},$$

as was proposed by Vanbiervliet, Michiels, et al. (2011).

The natural extension in the delay setting of the algebraic Lyapunov equation is the so-called delay-Lyapunov equation, a boundary value problem defining a matrix-valued function. In Jarlebring, Vanbiervliet, et al. (2011), a spectral discretization of this equation is proposed to compute the H^2 -norm, yielding geometric convergence of this norm

The primary contents of this chapter previously appeared as the second half of Provoost, E., and Michiels, W. (2024), “The Lanczos tau framework for time-delay systems: Padé approximation and collocation revisited”, *SIAM Journal on Numerical Analysis*, 62/6: 2529–48. The results in section 4.1 have been adapted from Provoost and Michiels (2026). Sections 4.2 and 4.6 have been prepared specifically for this thesis. In light of new insights since original publication, some figures were updated to use a Lanczos tau method based on Legendre polynomials instead of a Chebyshev basis and a comparison of the two was added.

in the single delay setting, at the cost of $O(n^6 N^3)$ operations, where N is the degree of the approximation. The method we base ourselves on improves the time complexity to $O(n^3 N^3)$ operations, but appears to reduce the convergence rate to third order algebraic convergence (Vanbiervliet, Michiels, et al. 2011).

In this chapter, we illustrate how using a Lanczos tau method for the approximation of the system allows us to recover geometric convergence for the same systems, and sometimes even results in super convergence (i.e. it gives the exact result for all N larger than some finite N_0), giving us the best of both worlds. This unexpected improvement served as the initial motivation to revisit the Lanczos tau method and thus this thesis.

As in the previous chapter, we mostly limit ourselves to the single delay case of system (3.1). In this setting, we also prove the proposed method to converge. Whilst we cannot prove the improved convergence rates, we do provide some hints as to why this rate might be obtained. A beautiful argument from symmetry will allow us to prove the super convergence. Finally, in section 4.6, we take a first look at what happens for multiple and distributed delays.

Overview. We start by showing that $\|G_N\|_{H^2}$ indeed converges to $\|G\|_{H^2}$ in section 4.1. We subsequently discuss how to compute $\|G_N\|_{H^2}$ given the state-space representation of the discretization (section 4.2). As mentioned before, using a Lanczos tau method satisfying a symmetry condition as proxy results in greatly improved convergence. This property and the required symmetry condition are the topic of section 4.3, where we also observe that the Legendre basis gives an additional improvement to the convergence rate compared to other symmetric bases. Part of the explanation of the geometric convergence likely stems from an intricate link between the algebraic Lyapunov equation of the discretization and the delay-Lyapunov equation of the original system (section 4.4). When using a symmetric Lanczos tau method to compute the H^2 -norm of a class of scalar time-delay systems we see super convergence. In section 4.5, we show how the aforementioned link to the delay Lyapunov equation can be used to prove this result. Finally, as a prelude to the following chapters, we look at a system with multiple delays and one with distributed delay in section 4.6.

4.1 PROOF OF CONVERGENCE

Before turning to the computational details and the convergence rate, we first establish that the H^2 -norm of G_N actually converges to that of G , a result that, to our knowledge, is novel. To this end, we shall need the following assumption on the basis used in the discretization (3.3).

Assumption 4.1.1. The basis function ϕ_k is symmetric when k is even and antisymmetric when k is odd, i.e. $\phi_k(-\tau - \theta) = (-1)^k \phi_k(\theta)$, $\forall \theta \in [-\tau, 0]$.

Note that when shifted from the interval $[-\tau, 0]$ to the interval $[-\frac{\tau}{2}, \frac{\tau}{2}]$, symmetric and antisymmetric functions correspond to the classical notion of even and odd functions, respectively.

For many of the bases used in this thesis, this property is a direct consequence of the defining weight function being symmetric (Szegő 1975: eq. 2.3.3). The usual pseudospectral collocation method (see section 3.5) using Chebyshev extremal points does not satisfy any analogous property, as the non-zero collocation points are not symmetric. The less common Chebyshev and Legendre zeroes, both extended with the node 0, however, are, per Theorem 3.5.2, equivalent to a Lanczos tau method that does satisfy this symmetry condition.

The reason this symmetry matters, is because, under this assumption, the underlying rational approximation qualitatively matches the exponential better on the imaginary axis. In the sense that its magnitude equals one, as shown in the following result.

Proposition 4.1.2. *If the basis $\{\phi_k\}_{k=0}^N$ underlying the Lanczos tau discretization (3.3) is chosen such that Assumption 4.1.1 is satisfied, we have*

$$|r_N(i\omega, -\tau)| = 1, \quad \forall \omega \in \mathbb{R},$$

where r_N is as in Proposition 3.4.1.

Proof. From Proposition 3.4.2 we have

$$r_N(i\omega, -\tau) = \frac{\sum_{k=0}^N i^k \phi_N^{(N-k)}(-\tau) \omega^k}{\sum_{k=0}^N i^k \phi_N^{(N-k)}(0) \omega^k}.$$

Assumption 4.1.1 implies $\phi_N^{(N-k)}(-\tau) = (-1)^k \phi_N^{(N-k)}(0)$, hence the denominator is the complex conjugate of the numerator. $\square$

From this property we readily have the following convergence theorem.

Theorem 4.1.3. *For an exponentially stable single delay system (3.1), we have $\|G_N\|_{H^2} \rightarrow \|G\|_{H^2}$ as $N \rightarrow \infty$ when the discretization (3.3) preserves stability for N sufficiently large and Assumptions 3.4.3 and 4.1.1 hold.*

Proof. Note that, by Proposition 3.4.1, we can write

$$G_N(i\omega) = \frac{1}{i\omega} C \left(I_n - \frac{1}{i\omega} (A_0 + A_1 r_N(i\omega, -\tau)) \right)^{-1} B,$$

and an analogous structure holds for G . As we know from Assumption 4.1.1 and Proposition 4.1.2 that $|r_N(i\omega, -\tau)| = |e^{-\tau i\omega}| = 1$ for all N and ω , $G(i\omega)$ and $G_N(i\omega)$ decay linearly for $\omega \rightarrow \infty$, independently from N . We thus have that, for all $\gamma > 0$, there is some $\omega^* > 0$ such that for all $|\omega| \geq \omega^*$ we can bound

$$\max \{\|G(i\omega)\|_F, \|G_N(i\omega)\|_F\} \leq \gamma|\omega|^{-1}$$

uniformly in N .

Let $e_N(\omega) = \|G(i\omega)\|_F^2 - \|G_N(i\omega)\|_F^2$ and take some arbitrarily small $\delta > 0$. The uniform, linearly decaying bound on $G(i\omega)$ and $G_N(i\omega)$ then implies that there must be some $\omega^{**} > 0$ such that, for all N ,

$$\left| \frac{1}{2\pi} \int_{-\infty}^{-\omega^{**}} e_N(\omega) d\omega + \frac{1}{2\pi} \int_{\omega^{**}}^{\infty} e_N(\omega) d\omega \right| \leq \frac{\delta}{2}.$$

Furthermore, preservation of stability implies that there are no singularities on the imaginary axis for N sufficiently large. Together with Assumption 3.4.3, we can thus apply Theorem 3.4.7 and have $e_N \rightarrow 0$ as $N \rightarrow \infty$, uniformly on $[-\omega^{**}, \omega^{**}]$. Hence, there is some N^* such that, for all $N > N^*$,

$$\left| \frac{1}{2\pi} \int_{-\omega^{**}}^{\omega^{**}} e_N(i\omega) d\omega \right| \leq \frac{\delta}{2}.$$

Taking these two bounds together we have, for all $N \geq N^*$,

$$\left| \frac{1}{2\pi} \int_{-\infty}^{\infty} e_N(i\omega) d\omega \right| \leq \delta.$$

Again invoking preservation of stability, there is some $N^{**} \geq N^*$ such that $\|G_N\|_{H^2} < \infty$ for all $N \geq N^{**}$. As δ was arbitrary and using the positivity of norms we can, for all $\epsilon > 0$, choose a δ such that, for all $N \geq N^{**}$,

$$|\|G\|_{H^2} - \|G_N\|_{H^2}| \leq \epsilon.$$

□

Note that, if the discretized system were not to be stable, the proof still implies that the L^2 -norm along the imaginary axis converges, as long as the false poles are not on the imaginary axis. From experience, however, the discretization does seem to preserve stability, for N sufficiently large, for most classical choices of basis. This is not unreasonable to assume as we saw preservation of stability when using a Legendre basis in Theorem 3.8.1.

In fact, as previously established, Assumption 4.1.1 holds for this basis and we have Assumption 3.4.3 from Proposition 3.4.5, the following thus holds.

Corollary 4.1.4. *For an exponentially stable, single delay system (3.1), we have $\|G_N\|_{H^2} \rightarrow \|G\|_{H^2}$ as $N \rightarrow \infty$, when using Legendre polynomials as basis.*

Like the convergence of G_N , this result will also be generalized to multiple delays in later chapters, see in particular Corollary 5.3.7.

4.2 THE H^2 -NORM OF A DELAY-FREE SYSTEM

To compute $\|G_N\|_{H^2}$, we need to know how to compute the H^2 -norm of a delay-free system, i.e. system (3.1) with $A_1 = \mathbf{0}$, that is exponentially stable. To avoid confusion with the delay case, let us call the transfer function of this delay-free system H .

Note that, using the Parseval–Plancherel identity, we can move to the time domain to find

$$\|H\|_{H^2} = \left(\int_0^\infty \|h(t)\|_F^2 dt \right)^{\frac{1}{2}},$$

where

$$h(t) = Ce^{A_0 t} B$$

is the so-called impulse response. From the definition of the Frobenius norm and the linearity of the trace, we subsequently have

$$\|H\|_{H^2}^2 = \text{tr} \left(C \left(\int_0^\infty e^{A_0 t} B B^T e^{A_0^T t} dt \right) C^T \right).$$

Call the inner integral W_c .¹ As an ansatz, we look at the following expression:

$$\begin{aligned} A_0 W_c + W_c A_0^T &= \int_0^\infty A_0 e^{A_0 t} B B^T e^{A_0^T t} dt + \int_0^\infty e^{A_0 t} B B^T e^{A_0^T t} A_0^T dt \\ &= \int_0^\infty \frac{d}{dt} \left(e^{A_0 t} B B^T e^{A_0^T t} \right) dt. \end{aligned}$$

Noting that exponential stability implies that $\lim_{t \rightarrow \infty} e^{A_0 t} = \mathbf{0}$, we thus find that W_c satisfies the Lyapunov equation

$$A_0 W_c + W_c A_0^T = -B B^T.$$

¹This quantity is known as the controllability Gramian in the control literature. If W_c is positive definite, the system is controllable. The W_o matrix we shall introduce shortly is, wholly analogously, known as the observability Gramian. See Brockett (1970) for a lengthier discussion of these matrices.

Finally we have, for an exponentially stable, delay-free system, that

$$\|H\|_{H^2} = \sqrt{\text{tr}(CW_c C^T)}.$$

Through analogous reasoning we also find

$$\|H\|_{H^2} = \sqrt{\text{tr}(B^T W_o B)},$$

where W_o is the solution of

$$W_o A_0 + A_0^T W_o = -C^T C.$$

Returning to the problem at hand, we thus have

$$\|G_N\|_{H^2} = \sqrt{\text{tr}(C_N V C_N^T)},$$

where the matrix $V \in \mathbb{C}^{nN \times nN}$, assuming the notation of (3.4), is the solution of the Lyapunov equation

$$\mathcal{E}_N^{-1} \mathcal{A}_N V + V \mathcal{A}_N^T \mathcal{E}_N^{-T} = -\mathcal{E}_N^{-1} \mathcal{B}_N \mathcal{B}_N^T \mathcal{E}_N^{-T},$$

and thus of the generalized symmetric Lyapunov equation

$$\mathcal{A}_N V \mathcal{E}_N^T + \mathcal{E}_N V \mathcal{A}_N^T = -\mathcal{B}_N \mathcal{B}_N^T. \quad (4.1)$$

Note that we will always opt to use the generalized Lyapunov equation. This is because, for some choices of discretization, $\mathcal{E}_N$ can be ill-conditioned for inversion. This matters less here, but will become crucial in chapter 5.

4.3 AN UNEXPECTED IMPROVEMENT TO GEOMETRIC RATE

We don't make mistakes; we just have happy accidents.

Robert 'Bob' Ross

If we approximate the H^2 -norm using different state space discretizations, namely pseudospectral collocation in Chebyshev extremal nodes and a Lanczos tau method satisfying Assumption 4.1.1, we consistently see results similar to fig. 4.1.² We get geometric convergence for the Lanczos tau method instead of the third order convergence of the other discretization that was observed in Vanbiervliet, Michiels, et al. (2011).




Figure 4.1: Convergence of the H^2 -norm of (3.1) with $A_0 = \begin{pmatrix} -2 & 1 \\ 3 & -8 \end{pmatrix}$, $A_1 = \begin{pmatrix} -1 & -1 \\ -1 & -1 \end{pmatrix}$, $\tau = 1$, and $B = C = I_2$, for different discretizations. The solid line corresponds to the Lanczos tau method using Legendre polynomials, the dashed line to pseudospectral collocation in Chebyshev extremal nodes.




Figure 4.2: Convergence of the H^2 -norm of the system of fig. 4.1, for different discretizations. The dashed line corresponds to the Lanczos tau method with Jacobi $\left(-\frac{1}{2}, -\frac{3}{4}\right)$ polynomials, which do not satisfy Assumption 4.1.1, the solid line to Legendre polynomials, which do. Note how the resulting curves are almost identical to fig. 4.1.

Note that the H^2 -norm depends on the behaviour of the transfer function along the entire imaginary line; however, the rational approximation of Proposition 3.4.1 is only accurate around zero (see e.g. fig. 3.1). From this perspective, the algebraic convergence reported in the literature is already somewhat unexpected. The observed improvement in convergence is thus highly surprising, especially as we did not change the computational cost.

Part of the explanation likely stems from a link to the delay-Lyapunov equation that we shall see in the next section. We could get geometric convergence for classical bases if we implicitly solve this equation using a spectral method (Tadmor 1986). However, this can only be part of the explanation as similar convergence would then be expected of other reasonably behaved bases, such as the Jacobi polynomials $P_k^{(-\frac{1}{2}, -\frac{3}{4})}$, which nonetheless degrade back to third order convergence (see fig. 4.2).

As we only note this deterioration when Assumption 4.1.1 is not satisfied, we presume that such a symmetry condition also plays a major part in explaining the observed geometric convergence. A partial answer is given by Proposition 4.1.2, where we saw that the rational approximation qualitatively matches the exponential better on the imaginary axis, in the sense that its magnitude equals one, precisely when this symmetry property is satisfied.

Finally, even between symmetric bases there is some difference in convergence rate. On fig. 4.3, we see that a Legendre basis results in faster convergence than Chebyshev polynomials of the second kind. Of all classical orthogonal bases, Legendre appears to perform best for the approximation of the H^2 -norm. Part of the explanation of this difference likely stems from the distribution of the poles. Comparing those of the Chebyshev basis of the second kind (fig. 4.4) to those of the Legendre basis (fig. 3.1), we see that the poles of the former sit significantly closer to the imaginary axis. That those obtained from a Legendre based discretization are, in some sense, optimal is not entirely unexpected given that they correspond to a diagonal Padé approximant (see section 3.7).

In the remainder of this and the following chapters we shall hence prefer Legendre polynomials. They have the fastest convergence rate, the most extensive supporting theory (see section 3.8 and section 4.1) and still admit a highly sparse discretization that displays improved numerical behaviour when approximating the H^2 -norm (see section 3.6).

²In a way, this section forms the seed from which this thesis grew. The original aim was to extend discretization-first approximation of the H^2 -norm to DDAES (the topic of chapter 5). In implementing this method, I used an ultraspherical discretization out of habit (see section 3.6), leading to the surprise on fig. 4.1. This is the ‘happy accident’ of the epigraph.




Figure 4.3: Convergence of the H^2 -norm of the system of fig. 4.1, for different discretizations. The dashed line corresponds to the Lanczos tau method using Chebyshev polynomials of the second kind, the solid line to Legendre polynomials.




Figure 4.4: Domain-colouring of $r_N(s, -\tau)$ for $\phi_N = U_N$, $\tau = 1$, and $N = 6$. Note how much closer the poles are to the imaginary axis compared to fig. 3.1.

4.4 CONNECTION TO THE DELAY-LYAPUNOV EQUATIONS

We can reinterpret the symmetric Lyapunov equation (4.1) in terms of operations on polynomials, which, as announced, allows us to partially explain the convergence seen in the previous section. To this end, note that we can identify the solution V with a bivariate polynomial

$$U(\theta, \theta') = \sum_{j=0}^N \sum_{k=0}^N V_{jk} \phi_j(\theta) \phi_k(\theta') \in \mathbb{C}^{n \times n},$$

with V_{jk} the relevant $n \times n$ subblock of V . Formally, as our choice of coordinates with respect to the basis $\Phi_N = \{\phi_k\}_{k=0}^N$ imposes a bijection between $\mathbb{C}^{nN}$ and $\mathbb{P}_N^n$, we similarly have an induced bijection between the tensor spaces $\mathbb{C}^{nN \times nN} \cong \mathbb{C}^{nN} \otimes \mathbb{C}^{nN}$ and $\mathbb{P}_N^n \otimes \mathbb{P}_N^n$ via the basis $\Phi_N \otimes \Phi_N$.

Inspired by this bijection, we extend our operator notation to bivariate polynomials, to have left multiplication mean application to the first variable and transposed right multiplication to be application to the second variable, i.e. we have $\varepsilon_\theta U \varepsilon_{\theta'}^T := U(\theta, \theta')$. This notation is justified by the tensor nature of U . Indeed, denoting by $[\cdot]$ the matrix realization as in section 3.3, we get

$$\varepsilon_\theta U \varepsilon_{\theta'}^T = [\varepsilon_\theta] V [\varepsilon_{\theta'}]^T \in \mathbb{C}^{n \times n}.$$

We then have the following result.

Proposition 4.4.1. *The H^2 -norm of the transfer function G_N (3.6) of the Lanczos tau approximation (3.3), if exponentially stable, is given by*

$$\|G_N\|_{H^2} = \sqrt{\text{tr}(C \varepsilon_0 U \varepsilon_0^T C^T)},$$

where the bivariate polynomial $U : [-\tau, 0] \times [-\tau, 0] \rightarrow \mathbb{C}^{n \times n}$, of degree (N, N) , is the unique solution of the system

$$\begin{cases} \mathcal{D}U \varepsilon_0^T + \mathcal{T}_{N-1}U(\varepsilon_0^T A_0^T + \varepsilon_{-\tau}^T A_1^T) = \mathbf{0}, \\ \mathcal{D}U \mathcal{T}_{N-1}^T + \mathcal{T}_{N-1}U \mathcal{D}^T = \mathbf{0}, \\ \varepsilon_\theta U \varepsilon_{\theta'}^T = (\varepsilon_{\theta'} U \varepsilon_\theta^T)^T, \\ (A_0 \varepsilon_0 + A_1 \varepsilon_{-\tau}) U \varepsilon_0^T + \varepsilon_0 U (\varepsilon_0^T A_0^T + \varepsilon_{-\tau}^T A_1^T) = -B B^T. \end{cases} \quad (4.2)$$

Proof. Note that we can write (4.1), without loss of generality, as

$$\begin{aligned} \begin{pmatrix} A_0[\varepsilon_0] + A_1[\varepsilon_{-\tau}] \\ [\mathcal{D}] \end{pmatrix} V \begin{pmatrix} [\varepsilon_0] \\ [\mathcal{T}_{N-1}] \end{pmatrix}^T + \begin{pmatrix} [\varepsilon_0] \\ [\mathcal{T}_{N-1}] \end{pmatrix} V \begin{pmatrix} A_0[\varepsilon_0] + A_1[\varepsilon_{-\tau}] \\ [\mathcal{D}] \end{pmatrix}^T \\ = -\begin{pmatrix} B \\ \mathbf{0} \end{pmatrix} \begin{pmatrix} B \\ \mathbf{0} \end{pmatrix}^T. \end{aligned} \quad (4.3)$$

Similarly, we have $\|G_N\|_{H^2}^2 = \text{tr}(C[\varepsilon_0]V[\varepsilon_0]^T C^T)$. The assertions can now be recovered using the earlier identification between matrices and bivariate polynomials, by splitting these expressions into their subblocks and noting that $V = V^T$ implies $\varepsilon_\theta U \varepsilon_{\theta'}^T = (\varepsilon_{\theta'} U \varepsilon_\theta^T)^T$. Finally, uniqueness and existence are direct consequences of the corresponding properties of (4.1). $\square$

We can use this result to connect the approach adopted here with the method of Jarlebring, Vanbiervliet, et al. (2011). The latter is based on the following characterization of the H^2 -norm, again assuming the system (3.1) is exponentially stable:

$$\|G\|_{H^2} = \sqrt{\text{tr}(C\lambda(0)C^T)},$$

where $\lambda \in C^0([-\tau, \tau]; \mathbb{C}^{n \times n})$ is the solution of the so-called delay-Lyapunov equation

$$\begin{cases} \lambda'(t) + \lambda(t)A_0^T + \lambda(t+\tau)A_1^T = \mathbf{0}, & t \in [-\tau, 0), \\ \lambda(-t) = \lambda(t)^T, \\ A_0\lambda(0) + A_1\lambda(-\tau) + \lambda(0)A_0^T + \lambda(\tau)A_1^T = -BB^T. \end{cases} \quad (4.4)$$

At least intuitively, we now have a link between both methods via $U(\theta, \theta') \approx \lambda(\theta - \theta')$. Indeed, assuming $\mathcal{T}_{N-1}$ to go to the identity operator as $N \rightarrow \infty$, the top and bottom equations of (4.2) and (4.4) match. We recover $\lambda(-t) = \lambda(t)^T$ from $\varepsilon_\theta U \varepsilon_{\theta'}^T = (\varepsilon_{\theta'} U \varepsilon_\theta^T)^T$ by setting $\theta = t$ and $\theta' = 0$ for $t \leq 0$ and $\theta = 0$ and $\theta' = -t$ for $t > 0$. Finally, the remaining equation guarantees that, in the limit, $U(\theta, \theta')$ is a function of $\theta - \theta'$. We can thus interpret the bivariate polynomial associated with the solution of (4.1) as an approximation of $(\theta, \theta') \mapsto \lambda(\theta - \theta')$ with respect to $\Phi_N \otimes \Phi_N$.

In the single delay case, λ is analytic on $[-\tau, 0)$ and $(0, -\tau]$ with a discontinuity in the first derivative at 0. The bivariate approximation of λ ensures that this discontinuity is always in an end point of the approximation interval, which would allow for geometric convergence in a spectral method. When multiple delays are introduced, and especially neutral terms, the function loses smoothness also on the interior of these segments (Jarlebring, Vanbiervliet, et al. 2011). We shall see in section 4.6 and the later chapters that the observed convergence rates are (mostly) consistent with what we would expect of a spectral method.

Finally, as an additional property, these equations show an interesting symmetry in the scalar case.

Lemma 4.4.2. *For a Lanczos tau discretization satisfying Assumption 4.1.1 of a scalar system (3.1), i.e. a system with $n = 1$, it holds that the solution of (4.2) satisfies*

$$U = \mathcal{R}U\mathcal{R}^T,$$

with $\mathcal{R}$ the reversal operator, such that $\varepsilon_\theta \mathcal{R} = \varepsilon_{-\theta-\tau}$.

Proof. Decomposing U in the derivative basis of ϕ_N , that is, writing

$$U = \sum_{j=0}^N \sum_{k=0}^N W_{jk} \Psi_{jk},$$

with $\Psi_{jk}(\theta, \theta') = \phi_N^{(N-j)}(\theta) \phi_N^{(N-k)}(\theta')$, we note that

$$\mathcal{D}U\mathcal{T}_{N-1}^T = -\mathcal{T}_{N-1}U\mathcal{D}^T$$

implies

$$\sum_{j=1}^N \sum_{k=0}^{N-1} W_{jk} \Psi_{j-1,k} = -\sum_{j=0}^{N-1} \sum_{k=1}^N W_{jk} \Psi_{j,k-1},$$

and thus $W_{j+1,k} = -W_{j,k+1}$. Hence, the coefficients on an anti-diagonal of W have constant magnitude and alternating sign. As we additionally have $W_{jk} = W_{kj}$ from $\varepsilon_\theta U \varepsilon_{\theta'}^T = (\varepsilon_{\theta'} U \varepsilon_\theta^T)^T$, the anti-diagonals of even length must be zero, thus $W_{jk} = 0$ if $j+k$ is odd. Furthermore, Assumption 4.1.1 implies that $\phi_N^{(N-k)}$ is symmetric when k is even and antisymmetric when k is odd. Hence, for $i+j$ even, the bivariate polynomial Ψ_{jk} is the product of either two symmetric or two antisymmetric polynomials and thus $\Psi_{jk} = \mathcal{R}\Psi_{jk}\mathcal{R}^T$. The assertion then follows from the linearity of $\mathcal{R}$ and the structure of W . $\square$

This result is useful in proving a surprising result seen in this scalar setting, the topic of the penultimate section of this chapter.

4.5 A CASE WITH SUPER CONVERGENCE

We consider the single delay system (3.1) in the scalar case, with $A_0 = A_1 = a < 0$. The zero solution of this system is then known to be exponentially stable for any $\tau \geq 0$, as proved by Hayes (1950: Theorem 1). Impressively, we see that the discretizations based on the Lanczos tau method, again under Assumption 4.1.1, already give the exact result at $N = 1$ (see fig. 4.5); the slight increase in error as N increases is due to rounding errors in the underlying matrix operations. This is once again rather surprising from the perspective of an integral over an unbounded domain, as the transfer functions barely match, even near zero (see fig. 4.6). We continue by proving this property.




Figure 4.5: Convergence of the H^2 -norm of (3.1) with $A_0 = A_1 = -1$, $\tau = 1$, and $B = C = 1$, for different discretizations. The solid line corresponds to the Lanczos tau method with Legendre polynomials, the dashed line to pseudospectral collocation in Chebyshev extremal nodes. For clarity of the figure, the relative error was lower bounded by 10^{-16} .




Figure 4.6: Transfer function G of the system of fig. 4.5, and the approximation G_1 from a Lanczos tau method using Legendre polynomials with $N = 1$.

For the class of systems showing this super convergence, we get as transfer function

$$G(s) = (s - a(1 + e^{-\tau s}))^{-1}. \quad (4.5)$$

We can compute its H^2 -norm analytically by solving the delay-Lyapunov equation. As before, we have $\|G\|_{H^2} = \sqrt{\lambda(0)}$, where

$$\begin{cases} \lambda'(t) + a(\lambda(t) + \lambda(t + \tau)) = 0, & t \in [-\tau, 0), \\ \lambda(-t) = \lambda(t), \\ 2a(\lambda(0) + \lambda(\tau)) = -1. \end{cases}$$

It is easy to verify that $\lambda(t) = \frac{a\tau - 2a|t| - 1}{4a}$ is the solution, yielding

$$\|G\|_{H^2} = \sqrt{\frac{a\tau - 1}{4a}}.$$

We can now show the following.

Proposition 4.5.1. *For the system with transfer function (4.5), we have*

$$\|G_N\|_{H^2} = \|G\|_{H^2}$$

for $N \geq 1$, when discretized using a Lanczos tau method satisfying Assumption 4.1.1.

Proof. From Proposition 4.4.1 we have $\|G_N\|_{H^2} = \sqrt{\varepsilon_0 U \varepsilon_0^T}$ with U the solution of

$$\begin{cases} \mathcal{D}U\varepsilon_0^T + a\mathcal{T}_{N-1}U(\varepsilon_0^T + \varepsilon_{-\tau}^T) = \mathbf{0}, \\ \mathcal{D}U\mathcal{T}_{N-1}^T + \mathcal{T}_{N-1}U\mathcal{D}^T = \mathbf{0}, \\ \varepsilon_\theta U \varepsilon_{\theta'}^T = (\varepsilon_{\theta'} U \varepsilon_\theta^T)^T, \\ a(\varepsilon_0 + \varepsilon_{-\tau})U\varepsilon_0^T + a\varepsilon_0 U(\varepsilon_0^T + \varepsilon_{-\tau}^T) = -1. \end{cases}$$

Let μ be the univariate polynomial $U\varepsilon_0^T$. From Lemma 4.4.2, we then have $U\varepsilon_{-\tau}^T = \mathcal{R}U\varepsilon_0^T = \mathcal{R}\mu$. For the top equation this gives

$$\mathcal{D}\mu = -a\mathcal{T}_{N-1}(\mu + \mathcal{R}\mu). \quad (4.6)$$

The bottom equation can be rewritten as $(\varepsilon_0 + \varepsilon_0\mathcal{R})\mu + \varepsilon_0(\mu + \mathcal{R}\mu) = -\frac{1}{a}$ or

$$\varepsilon_0(\mu + \mathcal{R}\mu) = -\frac{1}{2a}. \quad (4.7)$$

By expressing these equations in a basis, it is easily seen that (4.6) and (4.7) form a system of full rank, hence this system uniquely determines a polynomial. It is straightforwardly verified that $\mu(\theta) = \lambda(-\theta) = \frac{a\tau + 2a\theta - 1}{4a}$ solves this system for a degree-graded, orthogonal basis, if $N \geq 1$, as $\mu + \mathcal{R}\mu$ is a constant function. We thus get the exact result under the given conditions. $\square$

4.6 DISTRIBUTED AND MULTIPLE DELAYS

So far we have limited ourselves to RDDES with a single, discrete delay, i.e. (2.4) with $L(\xi) = A_0\xi(0) + A_1\xi(-\tau)$. In the following chapters we shall say more on the effect of multiple discrete delays, i.e. $L(\xi) = \sum_{k=0}^m A_k\xi(-\tau_k)$, but it is worthwhile to already take a first look at the effect of different functionals L on the convergence rate of the H^2 -norm.

Consider the RDDE systems

$$\begin{aligned}\dot{x}(t) &= -5x(t) - 2x(t-1) - x(t-1.9) + u(t), \\ y(t) &= x(t),\end{aligned}$$

and

$$\begin{aligned}\dot{x}(t) &= - \int_{-1}^0 e^{2\theta+1} x(t+\theta) d\theta + u(t), \\ y(t) &= x(t).\end{aligned}$$

We obtain the respective Lanczos tau discretizations of these systems by taking $[-\tau_m, 0]$, with τ_m the maximal delay, as support for the polynomial basis and substituting the row vector

$$[L] = (L(\phi_0) \quad \cdots \quad L(\phi_N))$$

for $[A_0\varepsilon_0 + A_1\varepsilon_{-\tau}]$ in (3.4). For the choice $\phi_k = P_k$, with P_k the k th Legendre polynomial shifted to the interval $[-1, 0]$, we obtain this row vector for the latter system from the identity

$$\int_{-1}^0 e^{2\theta+1} P_k(\theta) d\theta = 2i_k(1),$$

with $i_k(x)$ the k th modified spherical Bessel function of the first kind (Whittaker 1902).

On fig. 4.7 we see the convergence rates of the H^2 -norm for these two systems. As expected, the multiple delay example reverts to third order algebraic convergence. The distributed delay case appears unbothered and maintains geometric convergence.

In the multiple delay case, we now evaluate the approximation $r_N(s, \cdot)$ of the exponential function e^{-s} on the interior of the delay interval and Proposition 4.1.2 does thus not hold. Visually, we see why on fig. 3.1, the symmetry between poles and zeroes is broken. From this perspective, the convergence rate of the distributed delay is somewhat unexpected, as we now evaluate $r_N(s, \cdot)$ throughout the entire delay interval. However, the approximation of the exponential is only poor at high frequencies, where the integral would be close to zero anyway.




Figure 4.7: Convergence of the H^2 -norm for a system with two discrete delays (dashed line) and a distributed delay system (solid line) when using a Lanczos tau discretization using Legendre polynomials.

Both rates are consistent with a spectral discretization of the delay-Lyapunov equation. For a sufficiently well-behaved integration kernel, its solution has the same properties as for a single discrete delay, thus there is no apparent reason for a loss of geometric convergence.

For multiple discrete delays, Jarlebring, Vanbiervliet, et al. (2011) showed that we get bounded jumps in the derivatives at integer combinations of those delays, the order of the derivative depending on the type and number of delays. All such combinations of delays result in jumps in the $(k + 1)$ th derivative with k the absolute sum of the coefficients, i.e. a discontinuity in the first derivative at 0, in the second at the delays, in the third at two times a delay and sums and differences of two delays, etc. Delays that correspond to neutral terms behave worse, giving a jump in the first derivative at all their combinations. For a stable system, the size of these jumps does, however, decay rapidly in the number of involved delays. For both types of system, when the delays are commensurate, i.e. when all delays are some integer multiple of some $h > 0$, there are jumps at all integer multiples of the largest such h . When they are non-commensurate, the locations of jumps are dense in $[-\tau_m, \tau_m]$.

Hence, in the case of and RDDE with multiple discrete delays, the second derivative of λ is of bounded variation. From the discussion in section 3.2, we would thus expect second order convergence to λ . As the H^2 -norm is obtained from $\lambda(0)$, which is the value at an end point, it is not atypical to gain an additional order, which explains the observed third

order convergence. For NDDEs with multiple delays the first derivative is of bounded variation and we would thus expect first order convergence of λ . Unfortunately, we shall see in chapter 5 that, unlike RDDEs, we do not gain an additional order and we will thus only observe first order convergence.

These properties suggest that using a spline with knots at the delays, or even combinations thereof, could further increase the convergence order for RDDEs. In chapter 6 we will see that this is indeed the case and that this additionally simplifies parts of the theory of DDAEs we will develop in the next chapter.

4.7 SUMMARY

In this chapter we discussed the discretize-first approach of Vanbiervliet, Michiels, et al. (2011) for the approximation of the H^2 -norm of a single delay system. We proved—to our knowledge for the first time—that this method converges (Theorem 4.1.3) under the assumption that the underlying rational approximation converges (Assumption 3.4.3) and that a symmetric basis is used (Assumption 4.1.1). Using the convergence and stability guarantees of the Legendre based discretization established in the previous chapter, we could guarantee convergence for this specific discretization (Corollary 4.1.4).

We saw how a well-chosen Lanczos tau method can benefit the approximation of the H^2 -norm of a single delay system. In section 4.3, we observed significantly improved convergence rates compared to earlier work for bases satisfying the symmetry condition (Assumption 4.1.1); going from third algebraic order to geometric convergence. For the specific choice of Legendre polynomials the convergence seems the fastest, making them both the theoretical and practical choice for the approximation of the H^2 -norm.

From links to the delay-Lyapunov equation (section 4.4) and from qualitative properties of the rational approximation (Proposition 4.1.2), we could provide partial arguments for the observed geometric convergence. Proving this property, however, remains an open problem, though we were able to prove a case of super convergence in section 4.5. Note that this geometric convergence and, in particular, super convergence are at first glance somewhat unexpected, as the H^2 -norm is inherently a global characteristic; it depends on the behaviour of the transfer function along the entire imaginary axis.

In section 4.6 we took a first look at what happens for multiple and distributed delays. Whilst distributed delays also show geometric convergence, we do not see this improvement for multiple discrete delays, which reverted back to third order algebraic convergence.

5 *Computing the H^2 -norm of a delay differential algebraic system*

In the previous chapters we developed a discretization and an approximation method for the H^2 -norm for systems with a single discrete delay, adopting the discretize-first approach of Vanbiervliet, Michiels, et al. (2011). In this chapter we turn to the far more general case of approximating the H^2 -norm of systems described by delay differential algebraic equations (DDAEs) of index one (in the sense of Definition 5.0.1), with an arbitrary number of discrete state delays. To our knowledge, no method existed in the literature for this model class. The alternative approach of Jarlebring, Vanbiervliet, et al. (2011), using the delay-Lyapunov equation, already admits neutral type systems (which was part of our motivation to introduce DDAEs in chapter 2), however DDAEs are more general still and, as we saw in chapter 4, the discretize-first approach is computationally more tractable.

Remember from chapter 2 that DDAEs are particularly suitable for modelling interconnected systems and networks, where delays are often inherently present and algebraic equations simplify modelling. Moreover, the introduction of algebraic equations greatly eases the design of fixed order controllers, and the synthesis of stable approximate models, as these can now be readily included without needing to eliminate variables, leaving the parameters exposed for direct optimization (see chapter 7). Finally, DDAEs also allow for the inclusion of input delays and, detrimental for the H^2 -norm, feedthrough. Thus, DDAEs are—for better or worse—a much richer model class than their apparent simplicity might suggest.

The contents of this chapter previously appeared as the first part of Provoost, E., and Michiels, W. (2026), “Computing and optimizing the H^2 -norm of delay differential algebraic systems”, arXiv: 2603.10734 [math.NA]. The connection between the convergence rates and the properties of the solution of the delay-Lyapunov equation in section 5.3 is novel to this thesis.

In what follows, we consider the system

$$\begin{aligned} E\dot{\mathbf{x}}(t) &= \sum_{k=0}^m A_k \mathbf{x}(t - \tau_k) + B\mathbf{u}(t), \\ \mathbf{y}(t) &= C\mathbf{x}(t), \end{aligned} \quad (5.1)$$

where the leading matrix E is allowed to be singular, $\tau_0 = 0 < \tau_1 < \dots < \tau_m < \infty$ are the constant delays, $\mathbf{x}(t) \in \mathbb{C}^n$ is the state variable, $\mathbf{u}(t) \in \mathbb{C}^p$ is the input, and $\mathbf{y}(t) \in \mathbb{C}^q$ is the output at time t . That is, system (2.10) with $L(\xi) = \sum_{k=0}^m A_k \xi(-\tau_k)$. This system's transfer function is given by

$$G(s) = C(sE - \sum_{k=0}^m A_k e^{-\tau_k s})^{-1} B. \quad (5.2)$$

Finally, to prevent overly exotic cases, we will assume the system to have at most Kronecker index one. Whilst there is no full consensus on the definition of a DDAE's indices, we will, in following with e.g. Fridman (2002) and Unger (2020), use the following.

Definition 5.0.1. Let $\tilde{V}$ and $\tilde{U}$ be matrices with orthogonal columns such that $\text{span } \tilde{V} = \ker E$ and $\text{span } \tilde{U}^* = \ker E^*$. If $\tilde{U}A_0\tilde{V}$ is non-singular, the DDAE (5.1) is said to be of index one. If E is non-singular, and $\tilde{U}A_0\tilde{V}$ thus a 0×0 matrix, the system is said to be of index zero.

In other words, we will assume the system with the delayed terms eliminated, i.e. the DAE obtained by setting $A_k = \mathbf{0}$ for $k > 0$, to have at most Kronecker index one.

Remember from section 2.4 that, even though they are not explicitly present, neutral terms can easily be added using slack variables. For example,

$$\begin{aligned} \dot{x}(t) &= \dot{x}(t-1) + u(t), \\ y(t) &= x(t), \end{aligned}$$

and

$$\begin{aligned} \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} \dot{\mathbf{x}}(t) &= \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \mathbf{x}(t-1) + \begin{pmatrix} 1 \\ 0 \end{pmatrix} u(t), \\ y(t) &= \begin{pmatrix} 1 & 0 \end{pmatrix} \mathbf{x}(t), \end{aligned}$$

are equivalent from an input-output and stability perspective. In a similar fashion, input and output delays can also be added using such slack variables. A number of additional examples of such reformulations can be found in section 7.2.

More treacherously, there can be a frequency independent connection between input and output, i.e. $\limsup_{\omega \rightarrow \infty} G(i\omega) \neq \mathbf{0}$, hidden in

model (5.1), as this can, once again, be achieved using slack variables. The absence of this so-called feedthrough, is necessary for a finite H^2 -norm. In section 2.5, we saw that, for an exponentially stable system, this norm is given by

$$\|G\|_{H^2} = \left(\frac{1}{2\pi} \int_{-\infty}^{\infty} \|G(i\omega)\|_F^2 d\omega \right)^{\frac{1}{2}}.$$

Feedthrough blows up this integral; we will thus need additional tools to check for its presence.

Worse yet, this feedthrough can now, like the NDDE instability of fig. 2.4, appear after infinitesimal delay perturbations. Take the algebraic constraints

$$\begin{aligned} \mathbf{0} &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \mathbf{x}(t) - \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \mathbf{x}(t - \tau_1) \\ &\quad - \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \mathbf{x}(t - \tau_2) - \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \mathbf{x}(t - \tau_3) - \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix} u(t), \\ y(t) &= (1 \ 1 \ 0 \ 0) \mathbf{x}(t), \end{aligned}$$

for instance, which could quite easily hide within a large model. These simplify to

$$y(t) = u(t - (\tau_1 + \tau_2)) - u(t - \tau_3).$$

If $\tau_1 + \tau_2 = \tau_3$ this is the zero output system, which has zero as H^2 -norm. The slightest perturbation, however, clearly introduces feedthrough, resulting in an infinite H^2 -norm.

Overview. In section 5.2 we extend the discretization of chapter 3 and the approximation method for the H^2 -norm of chapter 4 to the DDAE setting, dealing with the additional challenges of this more general model class as needed. In particular, we prove this method to be sound under the assumption of a finite strong H^2 -norm. We continue in section 5.3 by proving that, under mild assumptions, the method converges and illustrate typical convergence rates. A number of preliminaries for these sections are discussed in the following section.

5.1 FINITENESS OF THE STRONG H^2 -NORM OF A DDAE

As noted, DDAEs can be used to model neutral systems. For such systems, we saw in chapter 2 that exponential stability is not necessarily preserved under infinitesimal perturbations of the delays. Equally, we know from the example in the introduction of this chapter that feedthroughs can appear under such perturbation. As both instability and feedthrough

yield an infinite H^2 -norm, its finiteness is thus not guaranteed under infinitesimal perturbations. This prompted Gomez, Jungers, et al. (2022) to introduce a strong notion of the H^2 -norm that takes such perturbations into account.

Let $\boldsymbol{\tau} = (\tau_1, \dots, \tau_m) \in \mathbb{R}_+^m$ and $B_\delta(\boldsymbol{\tau}) = \{\boldsymbol{\theta} \in \mathbb{R}_+^m : \|\boldsymbol{\tau} - \boldsymbol{\theta}\| < \delta\}$ the open ball around $\boldsymbol{\tau}$ of radius δ . We can then define the following norm.

Definition 5.1.1. The *strong H^2 -norm* of the system (5.1) is given by

$$\|G\|_{H^2}^s = \limsup_{\delta \rightarrow 0^+} \{\|G(\cdot; \boldsymbol{\theta})\|_{H^2} : \boldsymbol{\theta} \in B_\delta(\boldsymbol{\tau})\},$$

where the transfer function G is parameterized by the delay values $\boldsymbol{\theta}$.

Previous work allows us to characterize the finiteness of this norm using two components: strong exponential stability and an algebraic test to detect feedthrough under infinitesimal perturbations. In the remainder of this section we review how to establish both.

Remember that a system is exponentially stable if and only if the spectral abscissa $\alpha < 0$, where α is the supremum of the real parts of the characteristic roots of the system. Clearly, α is a function of $\boldsymbol{\tau}$, we can thus, analogously to before, define the following due to Hale and Verduyn Lunel (2002).

Definition 5.1.2. The system (5.1) is *strongly exponentially stable* if, and only if, there exists a finite $\delta > 0$ such that

$$\max_{\boldsymbol{\theta} \in B_\delta(\boldsymbol{\tau})} \alpha(\boldsymbol{\theta}) < 0.$$

Numerically, one can assess strong stability using the method described by Michiels (2011: section 4.3). We can also give an analytical characterization using the following standard form.

Let E be of rank $n - \nu$. Choose matrices $\bar{V}$ and $\bar{U}$ with orthogonal columns, such that $\text{span } \bar{V} = \ker E$, $\text{span } \bar{U}^* = \ker E^*$, $\bar{U} A_0 \bar{V} = -I_\nu$, and $\bar{U}^\perp E \bar{V}^\perp = I_{n-\nu}$, where the columns of $A^\perp$ span the orthogonal complement of $\text{span } A$. Applying the transform $\mathbf{x}(t) = \bar{V}^\perp \mathbf{x}_1(t) + \bar{V} \mathbf{x}_2(t)$ and left multiplying by $\begin{pmatrix} \bar{U}^\perp \\ \bar{U} \end{pmatrix}$ gives us

$$\begin{cases} \dot{\mathbf{x}}_1(t) = \sum_{k=0}^m A_{k,11} \mathbf{x}_1(t - \tau_k) + \sum_{k=0}^m A_{k,12} \mathbf{x}_2(t - \tau_k) + B_1 \mathbf{u}(t), \\ \mathbf{x}_2(t) = \sum_{k=0}^m A_{k,21} \mathbf{x}_1(t - \tau_k) + \sum_{k=1}^m A_{k,22} \mathbf{x}_2(t - \tau_k) + B_2 \mathbf{u}(t), \\ \mathbf{y}(t) = C_1 \mathbf{x}_1(t) + C_2 \mathbf{x}_2(t). \end{cases} \quad (5.3)$$

Note that we can write system (5.1) in the above form if, and only if, it is of index one.

We now have the following as a direct consequence of Corollary 2.1 of Hale and Verduyn Lunel (2002).

Theorem 5.1.3. *The system (5.3), and thus also (5.1) when of index one, is strongly exponentially stable if, and only if, the spectral abscissa is negative for the nominal delays, i.e. $\alpha(\tau) < 0$, and*

$$\max_{\theta \in [0, 2\pi]^m} \rho\left(\sum_{k=1}^m A_{k,22} e^{i\theta_k}\right) < 1,$$

where $\rho(A)$ is the spectral radius of A .

Given strong stability, we still need to verify that no feedthrough exists, nor can appear under infinitesimal perturbations. To this end, define the following family of matrices.

Definition 5.1.4. Let $\mathcal{P}_{k_0 \dots k_m}$ be the sum of all unique matrix products consisting of k_0 times $A_{0,22} = -I_\nu$, k_1 times $A_{1,22}$, ..., and k_m times $A_{m,22}$.

For example, $\mathcal{P}_{0,\dots,0} = I_\nu$, $\mathcal{P}_{1,2,0,\dots,0} = -A_{1,22}^2$, and

$$\mathcal{P}_{0,1,2,0,\dots,0} = A_{1,22}A_{2,22}^2 + A_{2,22}A_{1,22}A_{2,22} + A_{2,22}^2A_{1,22}.$$

These $\mathcal{P}_{\mathbf{k}}$ appear naturally in light of results like Theorem 5.1.3. Indeed, from the multinomial theorem directly follows the following property.

Proposition 5.1.5. *For $r \in \mathbb{N}$ we have*

$$\left(\sum_{k=0}^m z_k A_{k,22}\right)^r = \sum_{\substack{\mathbf{k} \in \mathbb{N}^{m+1} \\ \|\mathbf{k}\|_1 = r}} \binom{r}{\mathbf{k}} \mathbf{z}^{\mathbf{k}} \left(\|\mathbf{k}\|_1\right)^{-1} \mathcal{P}_{\mathbf{k}} = \sum_{\substack{\mathbf{k} \in \mathbb{N}^{m+1} \\ \|\mathbf{k}\|_1 = r}} \mathbf{z}^{\mathbf{k}} \mathcal{P}_{\mathbf{k}},$$

where $\mathbf{z}^{\mathbf{k}} = z_0^{k_0} z_1^{k_1} \dots z_m^{k_m}$.

Using these $\mathcal{P}_{\mathbf{k}}$ we can now formulate the following theorem due to Gomez, Jungers, et al. (2020: Corollary 9) as sharpened by Mattenet et al. (2022: Theorem 13).

Theorem 5.1.6. *The system (5.3), and thus also (5.1) when of index one, has a finite strong H^2 -norm, i.e. $\|G\|_{H^2}^s < \infty$, if, and only if, it is strongly stable and $C_2 \mathcal{P}_{\mathbf{k}} B_2 = \mathbf{0}$ for all $\mathbf{k} \in \mathbb{N}^{m+1}$. Furthermore, it is sufficient to only verify the latter condition for $\|\mathbf{k}\|_1 < \nu$.*

We now move to our proposed approximation method for the H^2 -norm of system (5.1). In the first section we describe the modifications needed and prove those to be sound for a system with a finite strong H^2 -norm. We continue by proving that the resulting method converges, under mild assumptions on the approximation, and illustrate typical convergence rates in the final section of this chapter.

5.2 LANCZOS TAU METHODS FOR DDAES

Extending the method of chapter 3 to DDAEs is rather natural. We start by multiplying by E in the left hand side of the top row of the discretization (3.4), yielding

$$\begin{aligned}\mathcal{E}_N \dot{\mathbf{x}}_N(t) &= \mathcal{A}_N \mathbf{x}_N(t) + \mathcal{B}_N \mathbf{u}(t), \\ \mathbf{y}_N(t) &= \mathcal{C}_N \mathbf{x}_N(t),\end{aligned}\tag{5.4}$$

with

$$\begin{aligned}\mathcal{E}_N &= \begin{pmatrix} E[\varepsilon_0] \\ [\mathcal{T}_{N-1}] \end{pmatrix}, \quad \mathcal{A}_N = \begin{pmatrix} \sum_{k=0}^m A_k[\varepsilon_{-\tau_k}] \\ [\mathcal{D}] \end{pmatrix}, \\ \mathcal{B}_N &= \begin{pmatrix} B \\ \mathbf{0} \end{pmatrix}, \quad \text{and} \quad \mathcal{C}_N = C[\varepsilon_0],\end{aligned}$$

where $[\cdot]$ signifies the expression in coordinates.

The result is a DAE, from which we will now have to eliminate the algebraic part before we can continue as for the H^2 -norm of an RDDE system without algebraic constraints (chapter 4). We start by writing the DAE such that the differential and algebraic parts are separated, analogous to (5.3).

Let V and U be matrices consisting of orthogonal columns, such that $\text{span } V = \ker \mathcal{E}_N$ and $\text{span } U^* = \ker \mathcal{E}_N^*$. Furthermore, let the columns of $A^\perp$ span the orthogonal complement of $\text{span } A$. Note that such U , V , $U^\perp$, and $V^\perp$ are easily obtained from a rank revealing decomposition of $\mathcal{E}_N$. Applying the transform $\mathbf{x}_N(t) = V^\perp \mathbf{x}_{N,1}(t) + V \mathbf{x}_{N,2}(t)$ and left multiplying by $\begin{pmatrix} U^\perp \\ U \end{pmatrix}$ gives us

$$\begin{cases} \mathcal{E}_{11} \dot{\mathbf{x}}_{N,1}(t) = \mathcal{A}_{11} \mathbf{x}_{N,1}(t) + \mathcal{A}_{12} \mathbf{x}_{N,2}(t) + \mathcal{B}_1 \mathbf{u}(t), \\ \mathbf{0} = \mathcal{A}_{21} \mathbf{x}_{N,1}(t) + \mathcal{A}_{22} \mathbf{x}_{N,2}(t) + \mathcal{B}_2 \mathbf{u}(t), \\ \mathbf{y}_N(t) = \mathcal{C}_1 \mathbf{x}_{N,1}(t) + \mathcal{C}_2 \mathbf{x}_{N,2}(t), \end{cases}\tag{5.5}$$

where

$$\begin{aligned}\mathcal{E}_{11} &= U^\perp \mathcal{E}_N V^\perp, & \mathcal{A}_{11} &= U^\perp \mathcal{A}_N V^\perp, & \mathcal{A}_{12} &= U^\perp \mathcal{A}_N V, \\ \mathcal{A}_{21} &= U \mathcal{A}_N V^\perp, & \mathcal{A}_{22} &= U \mathcal{A}_N V, & \mathcal{B}_1 &= U^\perp \mathcal{B}_N, \\ \mathcal{B}_2 &= U \mathcal{B}_N, & \mathcal{C}_1 &= \mathcal{C}_N V^\perp, & \text{and} & \quad \mathcal{C}_2 = \mathcal{C}_N V.\end{aligned}$$

Unlike (5.3), we did not pick V and U such that $\mathcal{E}_{11}$ is an identity matrix. Whilst this is indeed possible, it can lead to numerical issues when the largest and smallest non-zero singular values of $\mathcal{E}_N$ are highly separated. Nonetheless, $\mathcal{E}_{11}$ is invertible by construction, which is sufficient for the remainder of our results.

To finally eliminate the algebraic part from (5.5), we need to show $\mathcal{A}_{22}$ to be invertible. To this end, we first need the following result, extending Theorem 5.1.3.

Lemma 5.2.1. *If the system (5.3) is strongly exponentially stable then*

$$\max_{\mathbf{z} \in \mathbb{D}^m} \rho(\sum_{k=1}^m A_{k,22} z_k) < 1,$$

where $\mathbb{D} = \{z \in \mathbb{C} : |z| \leq 1\}$ is the closed unit disc.

Proof. The result is shown through contradiction. Assume (5.3) is strongly exponentially stable and there do exist some $|z_1|, |z_2|, \dots, |z_m| \leq 1$ such that

$$\rho(\sum_{k=1}^m A_{k,22} z_k) \geq 1.$$

Then there exists some $|\lambda| \geq 1$ such that

$$\det(a_1 a_2 \cdots a_m \lambda I_\nu - \sum_{k=1}^m A_{k,22} z_k) = 0,$$

where, for now, $a_1 = a_2 = \cdots = a_m = 1$. As a consequence of Rouché's theorem, an eigenvalue cannot disappear unless it moves to infinity. Through continuation, and keeping the other variables fixed, we thus have a mapping $[1, \infty) \ni a_1 \mapsto z_1(a_1)$. As the spectral radius is bounded by any natural matrix norm, $|z_1(a_1)| \rightarrow \infty$ as $a_1 \rightarrow \infty$, there is thus some $a_1^* \geq 1$ such that $z_1^* = z_1(a_1^*)$ has modulus one. Repeating this procedure, we find $|z_1^*| = |z_2^*| = \cdots = |z_m^*| = 1$ and $a_1^*, a_2^*, \dots, a_m^* \geq 1$ such that

$$\det(a_1^* a_2^* \cdots a_m^* \lambda I_\nu - \sum_{k=1}^m A_{k,22} z_k^*) = 0.$$

However, as $|\lambda| \geq 1$, this would mean that

$$\rho(\sum_{k=1}^m A_{k,22} z_k^*) \geq |a_1^* a_2^* \cdots a_m^* \lambda| \geq 1,$$

contradicting Theorem 5.1.3. □

As we shall see shortly, for our discretization the z_k of the above lemma correspond to $\phi_N(-\tau_k)$, $k = 0, 1, \dots, m$. To use this result then, we impose the following assumption on the basis.

Assumption 5.2.2. Choose the basis $\{\phi_j\}_{j=0}^N$ such that $|\phi_N(-\tau_k)| \leq |\phi_N(0)|$ for $k = 1, \dots, m$.

This condition on ϕ_N is satisfied for common choices of orthogonal polynomials. For many Jacobi polynomials, and in particular those used in our experiments, this is guaranteed by the well-known limits on their extrema (see Szegő 1975: p. 168).

We can then prove the following.

Proposition 5.2.3. *If the original DDAE (5.1) is of index one and strongly exponentially stable, and ϕ_N satisfies Assumption 5.2.2, then $\mathcal{A}_{22}$, as defined by (5.5), is invertible.*

Proof. Note that the vector $\xi \in \mathbb{C}^{n(N+1)}$, corresponding to the polynomial $\xi \in \mathbb{P}_N^n$, is in the right nullspace of $\mathcal{E}_N$ (and thus in $\text{span } V$) if, and only if, $\varepsilon_0 \xi \in \ker E$ and $\xi \in \ker \mathcal{T}_{N-1} = \text{span } I_n \phi_N$. The tuple $(\mathbf{z}, \zeta) \in \mathbb{C}^n \times \mathbb{C}^{nN}$ is in the left nullspace of $\mathcal{E}_N$ (and its complex transpose thus in $\text{span } U^*$) if, and only if, $\mathbf{z}^* \in \ker (E[\varepsilon_0])^* = \ker E^*$ and $\zeta^* \in \ker [\mathcal{T}_{N-1}]^* = \{\mathbf{0}^*\}$. In short

$$\begin{aligned} \xi \in \text{span } V &\iff \varepsilon_0 \xi \in \ker E \wedge \xi \in \text{span } I_n \phi_N \quad \text{and} \\ (\mathbf{z}, \zeta)^* \in \text{span } U^* &\iff \mathbf{z}^* \in \ker E^* \wedge \zeta = \mathbf{0}. \end{aligned}$$

These results imply that there exists a basis transform,¹ i.e. an invertible matrix, T , such that

$$T\mathcal{A}_{22}T^{-1} = TUA_NV T^{-1} = \begin{pmatrix} \bar{U}_0^* \\ \mathbf{0} \end{pmatrix}^* \mathcal{A}_N [\phi_N \bar{V}] = \sum_{k=0}^m A_{k,22} \phi_N(-\tau_k),$$

where $\bar{U}$, $\bar{V}$, $A_{0,22} = -I_\nu$, $A_{1,22}$, ..., and $A_{m,22}$ are as defined in (5.3). As $\phi_N(0) \neq 0$ (Szegő 1975: Theorem 3.3.1), $\frac{1}{\phi_N(0)}\mathcal{A}_{22}$ is invertible by Lemma 5.2.1 and Assumption 5.2.2 and so too $\mathcal{A}_{22}$. $\square$

Having established the invertibility of $\mathcal{A}_{22}$, we can apply the Schur complement formula to eliminate the algebraic part from (5.5) and obtain

$$\begin{aligned} \mathcal{E}_{11} \dot{\mathbf{x}}_{N,1}(t) &= \tilde{A} \mathbf{x}_{N,1}(t) + \tilde{B} \mathbf{u}(t), \\ \mathbf{y}_N(t) &= \tilde{C} \mathbf{x}_{N,1}(t) + \tilde{D} \mathbf{u}(t), \end{aligned} \tag{5.6}$$

with

$$\begin{aligned} \tilde{A} &= \mathcal{A}_{11} - \mathcal{A}_{12} \mathcal{A}_{22}^{-1} \mathcal{A}_{21}, & \tilde{B} &= \mathcal{B}_1 - \mathcal{A}_{12} \mathcal{A}_{22}^{-1} \mathcal{B}_2, \\ \tilde{C} &= \mathcal{C}_1 - \mathcal{C}_2 \mathcal{A}_{22}^{-1} \mathcal{A}_{21}, \quad \text{and} & \tilde{D} &= -\mathcal{C}_2 \mathcal{A}_{22}^{-1} \mathcal{B}_2. \end{aligned}$$

The remaining challenge novel to DDAEs is then to show that if the system (5.1) does not have feedthrough, we always have $\tilde{D} = \mathbf{0}$. Luckily, this turns out to be the case as shown in the following result.

¹Loosely speaking, V and U project onto to the right and left nullspaces of E respectively, where in V the information can only be ‘carried’ by ϕ_N and in U the tail has to be the zero function.

Theorem 5.2.4. *If the DDAE (5.1) is of index one, has a finite strong H^2 -norm, and ϕ_N satisfies Assumption 5.2.2, then the system (5.6) does not have feedthrough, i.e. $\tilde{D} = \mathbf{0}$.*

Proof. As a finite strong H^2 -norm implies strong exponential stability, we can continue the proof of Proposition 5.2.3. Analogously to the argument there, we have $T\mathcal{B}_2 = TUB_N = B_2$ and $C_2T^{-1} = C_NV T^{-1} = C_2$, thus

$$\tilde{D} = -C_2T^{-1}TA_{22}^{-1}T^{-1}T\mathcal{B}_2 = -C_2(\sum_{k=0}^m A_{k,22} \phi_N(-\tau_k))^{-1}B_2.$$

Because of the invertibility of A_{22} and a standard corollary of the Cayley–Hamilton theorem, there exists some polynomial p of finite degree such that

$$\tilde{D} = -C_2 p(\sum_{k=0}^m A_{k,22} \phi_N(-\tau_k)) B_2. \quad (5.7)$$

From Proposition 5.1.5 we have, for $r \in \mathbb{N}$,

$$\begin{aligned} C_2(\sum_{k=0}^m A_{k,22} \phi_N(-\tau_k))^r B_2 \\ = \sum_{\substack{\mathbf{k} \in \mathbb{N}^{m+1} \\ \|\mathbf{k}\|_1 = r}} \phi_N(-\tau_0)^{k_0} \cdots \phi_N(-\tau_m)^{k_m} C_2 \mathcal{P}_{\mathbf{k}} B_2. \end{aligned}$$

As (5.7) entails that $\tilde{D}$ is a weighted sum of such terms, we finally have $\tilde{D} = \mathbf{0}$ from Theorem 5.1.6. $\square$

As was the case for RDDEs, the system obtained by discretizing a stable system need not itself be stable. If this were to occur, one can reflect the poles on the right half-plane across the imaginary axis, as is commonly done in reduced order modelling (see e.g. Conni et al. 2024). Typically, these poles have a high frequency. As the transfer function converges rapidly on compact sets (see Theorem 5.3.3), flipping thus induces only a minor additional error on the transfer function and hence the computed H^2 -norm. As we shall comment on more extensively in the next section, like in the RDDE setting, the used discretization seems to preserve stability for N sufficiently large, so alternatively one can attempt to increase N if the discretized system is not stable. In either case, unstable systems obtained from stable ones are exceedingly rare.

Assuming we obtain a stable discretized system and the conditions for Theorem 5.2.4 hold, we can finally, as in chapter 4, approximate the strong H^2 -norm as

$$\|G\|_{H^2}^s \approx \|G_N\|_{H^2} = \sqrt{\text{tr}(\tilde{C}P\tilde{C}^T)}, \quad (5.8)$$

where P is the solution of $\tilde{A}P\mathcal{E}_{11}^T + \mathcal{E}_{11}P\tilde{A}^T = -\tilde{B}\tilde{B}^T$.

We thus arrive at the following procedure to approximate the strong H^2 -norm of the DDAE system (5.1):

1. Test for strong exponential stability using the method described in Michiels (2011: section 4.3). When this check fails, we have $\|G\|_{H^2}^s = \infty$.
2. Test for feedthrough under infinitesimal perturbations using Theorem 5.1.6. If this test fails, we again know that $\|G\|_{H^2}^s = \infty$.
3. Construct (5.6) and, if exponentially stable, use (5.8) to approximate $\|G\|_{H^2}^s$. If the discretized system is not stable, first flip the poles on the right half-plane across the imaginary axis.

5.3 CONVERGENCE

In the previous section we proposed to approximate the H^2 -norm of a DDAE from a Lanczos tau approximation and addressed the problem of feedthrough in Theorem 5.2.4, however, we have yet to show it to converge. In this section we shall show this to be the case, assuming the discretization preserves stability and has a sufficiently good (in the sense of Assumptions 3.4.3, 5.3.2 and 5.3.4) underlying approximation of the exponential, effectively extending Theorem 4.1.3, the convergence result in the single delay RDDE case. We will conclude the chapter with a number of experiments illustrating typical convergence rates of the method. We start with the following result, which is a straightforward extension of Proposition 3.4.1.

Proposition 5.3.1. *The transfer function of (5.4) can be expressed as*

$$G_N(s) = C(sE - \sum_{k=0}^m A_k r_N(s, -\tau_k))^{-1} B,$$

where $r_N(s, \theta)$ is a rational function of type (N, N) in $s \in \mathbb{C}$ and a polynomial of degree N in $\theta \in [-\tau_m, 0]$, such that

$$\begin{cases} r_N(s, 0) = 1, \\ \mathcal{D}r_N(s, \cdot) = s\mathcal{T}_{N-1}r_N(s, \cdot). \end{cases}$$

Proof. Taking the Laplace transform of (5.4), assuming zero initial conditions, gives

$$\begin{cases} sE[\varepsilon_0]\hat{\mathbf{x}}_N(s) = \sum_{k=0}^m A_k[\varepsilon_{-\tau_k}]\hat{\mathbf{x}}_N(s) + B\hat{\mathbf{u}}(s), \\ s[\mathcal{T}_{N-1}]\hat{\mathbf{x}}_N(s) = [\mathcal{D}]\hat{\mathbf{x}}_N(s), \\ \hat{\mathbf{y}}_N(s) = C[\varepsilon_0]\hat{\mathbf{x}}_N(s), \end{cases} \quad (5.9)$$

where $\hat{f}(s)$ is the Laplace transform of $f(t)$. Let $\mathbf{r}_N(s)$ be the coefficients of a matrix valued polynomial such that

$$\begin{cases} [\varepsilon_0] \mathbf{r}_N(s) = I_n, \\ s[\mathcal{T}_{N-1}] \mathbf{r}_N(s) = [D] \mathbf{r}_N(s). \end{cases}$$

From this definition we easily see that

$$\hat{\mathbf{x}}_N(s) = \mathbf{r}_N(s)(sE[\varepsilon_0] \mathbf{r}_N(s) - \sum_{k=0}^m A_k[\varepsilon_{-\tau_k}] \mathbf{r}_N(s))^{-1} B \hat{\mathbf{u}}(s)$$

satisfies (5.9), and hence

$$\begin{aligned} G_N(s) &= C[\varepsilon_0] \mathbf{r}_N(s)(sE[\varepsilon_0] \mathbf{r}_N(s) - \sum_{k=0}^m A_k[\varepsilon_{-\tau_k}] \mathbf{r}_N(s))^{-1} B \\ &= C(sE - \sum_{k=0}^m A_k[\varepsilon_{-\tau_k}] \mathbf{r}_N(s))^{-1} B. \end{aligned}$$

As $[\varepsilon_\theta] \mathbf{r}_N(s) = I_n r_N(s, \theta)$, the claim holds. $\square$

Later we will need an expression of G_N where the differential and algebraic parts are separated. Using the standard form (5.3) for a system of index one, we can rewrite this transfer function as

$$G_N(s) = \begin{pmatrix} C_1 & C_2 \end{pmatrix} \begin{pmatrix} sI_{n-\nu} - A_{N,11}(s) & -A_{N,12}(s) \\ -A_{N,21}(s) & -A_{N,22}(s) \end{pmatrix}^{-1} \begin{pmatrix} B_1 \\ B_2 \end{pmatrix},$$

where $A_{N,ij}(s) = \sum_{k=0}^m A_{k,ij} r_N(s, -\tau_k)$, $A_{0,22} = -I_\nu$, and all other matrices as in (5.3). Using the Schur complement formula we obtain the following expression for G_N , where we elided the s dependence for readability and A^{-1} should be read as $A(s)^{-1}$,

$$\begin{pmatrix} C_1 & C_2 \end{pmatrix} \begin{pmatrix} D_N^{-1} & -D_N^{-1} A_{N,12} A_{N,22}^{-1} \\ -A_{N,22}^{-1} A_{N,21} D_N^{-1} & A_{N,22}^{-1} + A_{N,22}^{-1} A_{N,21} D_N^{-1} A_{N,12} A_{N,22}^{-1} \end{pmatrix} \begin{pmatrix} B_1 \\ B_2 \end{pmatrix}, \quad (5.10)$$

with $D_N(s) = sI_{n-\nu} - A_{N,11}(s) - A_{N,12}(s) A_{N,22}(s)^{-1} A_{N,21}(s)$.

By inspection of the definition of $r_N(s, \theta)$, we see that for every $s \in \mathbb{C}$ we have a polynomial approximation of $[-\tau_m, 0] \ni \theta \mapsto e^{\theta s}$, hence, by comparing G_N to (5.2), Proposition 5.3.1 can also in the DDAE case be interpreted as approximating each $s \mapsto e^{-\tau_k s}$ in $G(s)$ by $s \mapsto r_N(s, -\tau_k)$. Like in the single delay case, the quality of the approximation of G by G_N , and hence the convergence of the H^2 -norm, thus depends on the quality of the approximation of the exponential. In what follows we shall need the following assumption on the convergence of these r_N , in addition to the local convergence of Assumption 3.4.3.

Assumption 5.3.2. The basis $\{\phi_k\}_{k=0}^N$ is chosen such that there is a function f of order $o(|\omega|^{\frac{1}{4}})$, i.e. $\lim_{\omega \rightarrow \infty} |\omega|^{-\frac{1}{4}} f(\omega) = 0$, bounding $|r_N(i\omega, -\tau_k)|$ for all $k = 0, \dots, m$, $\omega \in \mathbb{R}$, and N .

This assumption is suggested by the common experience that, for a good choice of basis, the result of a tau method (such as, by its definition, $r_N(s, \cdot)$) tends to be quasi-optimal, i.e. with an error within some factor $C \geq 1$ from that of the best polynomial approximation of the true solution of the same degree (Boyd 2001). In the case of Chebyshev polynomials, this observation is motivated by the spectral accuracy of the derivative (Tadmor 1986).

As we saw in section 3.4, Assumption 3.4.3 can generally be asserted by Corollary 3.4.4. For Assumption 5.3.2 we know of no such readily available result. As the exponential is an entire function, the best polynomial approximation of $[-\tau_m, 0] \ni \theta \mapsto e^{s\theta}$ converges super-geometrically in the supremum norm (Bernstein 1912). On the imaginary axis, i.e. $s = i\omega$, the constant zero polynomial has unit error, trivially and uniformly bounding the error of the best polynomial approximation in ω and N . A bounded approximation error to a bounded target function implies that the best polynomial approximation itself is uniformly bounded in ω and N . However, the quasi-optimality factor is problem dependent and can thus grow with ω . Numerical experiments fortunately suggest that the growth of the maximal error is only logarithmic in N , that it is concentrated around a single frequency, and that it moves away from the origin at a superlinear rate, attaining maximal error around $\omega \sim \frac{1}{\tau_m} N \log(N)$ after which the approximation, if $\theta \neq -\tau_m$, rapidly approaches zero. All this suggests that the obtained approximations are all well within the requirements of Assumption 5.3.2.

As in the single delay case we can use the structure of Proposition 5.3.1 and Lemma 3.4.6 to extend the local convergence of Assumption 3.4.3 to convergence on compact sets of the transfer function.

Theorem 5.3.3. *If Assumption 3.4.3 holds, then G_N converges uniformly to G as $N \rightarrow \infty$ on compact sets in $\mathbb{C}$ that do not contain a singularity for N sufficiently large.*

Proving convergence of the H^2 -norm will rely on the local convergence of the transfer function established in the previous theorem and bounding the ‘tail’. For the latter we will need, in conjunction with Assumption 5.3.2, the following.

Assumption 5.3.4. There exists an M such that, for all N and ω , we have the bound $\|A_{N,22}(i\omega)^{-1}\|_F \leq M$.

As previously discussed, in experiments the magnitude of the rational approximations only exceed one near $\omega \sim \frac{1}{\tau_m} N \log(N)$. As the singular matrices form a set of measure zero, this is thus all but guaranteed by strong stability and Lemma 5.2.1, and appears to hold in experiments.

This then is sufficient to prove the final convergence result, which is the natural extension of Theorem 4.1.3 to DDAES.

Theorem 5.3.5. *For a strongly exponentially stable system (5.1) of index one with finite strong H^2 -norm, the method described in section 5.2 converges, i.e. $\|G_N\|_{H^2} \rightarrow \|G\|_{H^2}$ for $N \rightarrow \infty$, when the discretization preserves stability for N sufficiently large and Assumptions 3.4.3, 5.3.2 and 5.3.4 hold.*

Proof. We start by bounding the tails of $G_N(i\omega)$ using the form in equation (5.10). Note that, as Assumption 5.3.4 ensures that $A_{N,22}(i\omega)$ is always invertible, we have $C_2 A_{N,22}(i\omega)^{-1} B_2 = \mathbf{0}$ through an analogous argument to the proof of Theorem 5.2.4. Additionally, Assumption 5.3.2 implies that $D_N(i\omega)^{-1}$ decays at least linearly for $\omega \rightarrow \infty$, uniformly in N . Using this decay and Assumptions 5.3.2 and 5.3.4, we see that the remaining non-zero terms, and thus $G_N(i\omega)$, decay at least as $|\omega|^{-\frac{1}{2}-\alpha}$ for some $\alpha > 0$. As $G(i\omega)$ decays linearly for $\omega \rightarrow \infty$, we thus have that, for all $\gamma > 0$, there is some $\omega^* > 0$ such that for all $|\omega| \geq \omega^*$ we can bound

$$\max \{ \|G(i\omega)\|_F, \|G_N(i\omega)\|_F \} \leq \gamma |\omega|^{-\frac{1}{2}-\alpha}$$

uniformly in N .

Let $e_N(\omega) = \|G(i\omega)\|_F^2 - \|G_N(i\omega)\|_F^2$ and take some arbitrarily small $\delta > 0$. The uniform, decaying bound on $G(i\omega)$ and $G_N(i\omega)$ then implies that there must be some $\omega^{**} > 0$ such that, for all N ,

$$\left| \frac{1}{2\pi} \int_{-\infty}^{-\omega^{**}} e_N(\omega) d\omega + \frac{1}{2\pi} \int_{\omega^{**}}^{\infty} e_N(\omega) d\omega \right| \leq \frac{\delta}{2}.$$

Furthermore, preservation of stability implies that there are no singularities on the imaginary axis for N sufficiently large. Together with Assumption 3.4.3, we can thus apply Theorem 5.3.3 and have $e_N \rightarrow 0$ as $N \rightarrow \infty$, uniformly on $[-\omega^{**}, \omega^{**}]$. Hence, there is some N^* such that, for all $N > N^*$,

$$\left| \frac{1}{2\pi} \int_{-\omega^{**}}^{\omega^{**}} e_N(i\omega) d\omega \right| \leq \frac{\delta}{2}.$$

Taking these two bounds together we have, for all $N \geq N^*$,

$$\left| \frac{1}{2\pi} \int_{-\infty}^{\infty} e_N(i\omega) d\omega \right| \leq \delta.$$

Again invoking preservation of stability, there is some $N^{**} \geq N^*$ such that $\|G_N\|_{H^2} < \infty$ for all $N \geq N^{**}$. As δ was arbitrary and using the positivity of norms we can, for all $\epsilon > 0$, choose a δ such that, for all $N \geq N^{**}$,

$$|\|G\|_{H^2} - \|G_N\|_{H^2}| \leq \epsilon.$$

□

Note that for RDDEs with multiple discrete delays the argument simplifies dramatically as Assumption 5.3.4 becomes irrelevant. We then also have $D_N(s) = sI_n - A_{N,11}(s)$ and $G_N(s) = C_1 D_N(s)^{-1} B_1$, thus Assumption 5.3.2 can be relaxed to sublinear growth.

Like Theorem 4.1.3, if the discretized system were not to be stable, the proof still implies that the L^2 -norm along the imaginary axis converges, as long as the false poles are not on the imaginary axis. From experience, however, the discretization does seem to preserve stability, for N sufficiently large, for most classical choices of basis. As we saw in chapter 3, this is not unreasonable to assume as similar results on the preservation of stability have been proven for RDDEs by Ito (1985) and Ito and Kappel (1992) when using a Legendre basis. We saw a special case of these results in Theorem 3.8.1 and we could use this to guarantee convergence of the H^2 -norm of a single delay RDDE in Corollary 4.1.4. We will now extend this result to a single delay DDAE.

Remember that we have Assumption 3.4.3 from Proposition 3.4.5. Furthermore, Assumption 5.2.2 is guaranteed by equation 7.21.1 of Szegő (1975). Hence, up to preservation of stability and Assumptions 5.3.2 and 5.3.4, all conditions on the discretization required for Theorem 5.3.5 to hold are met when using Legendre polynomials. As these polynomials are symmetric on their domain, we have

$$|r_N(i\omega, -\tau_m)| = 1$$

for all $\omega \in \mathbb{R}$ (Proposition 4.1.2). In the single delay case we thus trivially satisfy Assumption 5.3.2 and have Assumption 5.3.4 from Lemma 5.2.1, if the system is strongly exponentially stable. This leaves preservation of stability, which we establish for the single delay case in the following result as an extension to the RDDE version in Theorem 3.8.1.

Theorem 5.3.6. *For a strongly exponentially stable, single delay DDAE system (5.1), i.e. $m = 1$, of index one, the discretization (5.4) preserves stability for N sufficiently large, when using a Legendre basis.*

Proof. From Theorem 3.7.1 we know that, for this choice of basis, $r_N(s, -\tau_m)$ is the N th diagonal Padé approximant of $s \mapsto e^{-\tau_m s}$. Birkhoff and Varga (1965: Lemma 8) showed of these approximations that they are bounded in magnitude by one on the closed right half-plane, i.e. $|r_N(s, -\tau_m)| \leq 1$ for all $s \in \mathbb{C}_{\geq 0} = \{z \in \mathbb{C} : \operatorname{Re}(z) \geq 0\}$.

Together with Lemma 5.2.1 this bound on r_N implies that $A_{N,22}(s)^{-1}$ is bounded for $s \in \mathbb{C}_{\geq 0}$. From equation (5.10) we then conclude that all characteristic roots λ of the discretization in this half-plane satisfy $\det(D_N(\lambda)) = 0$ and thus

$$|\lambda| \leq \|A_{N,11}(\lambda) + A_{N,12}(\lambda) A_{N,22}(\lambda)^{-1} A_{N,21}(\lambda)\|_2.$$

The bound on r_N further implies that $A_{N,11}$, $A_{N,12}$, and $A_{N,21}$ are bounded on $\mathbb{C}_{\geq 0}$. Thus, the zeroes of $\det(D_N(\lambda))$ that lie in $\mathbb{C}_{\geq 0}$ are bounded in magnitude by some finite constant M and are thus constrained to a semicircle Ω of radius M .

As the original system is stable, it has no characteristic roots in Ω . From the argument for Theorem 5.3.3, we know that $D_N(s)$ converges on compact sets. Hence, for N sufficiently large, the discretized system cannot have any characteristic roots in Ω and thus has none in the entire closed right half-plane. $\square$

Using this theorem and the preceding discussion we finally have the following generalization of Corollary 4.1.4.

Corollary 5.3.7. *For a strongly exponentially stable, single delay system (5.1) of index one with finite strong H^2 -norm, the method described in section 5.2 converges, i.e. $\|G_N\|_{H^2} \rightarrow \|G\|_{H^2}$ for $N \rightarrow \infty$, when using Legendre polynomials as basis.*

The reason we could provide a complete argument for the single delay case mainly hinges on the boundedness of $r_N(s, -\tau_m)$. In chapter 6 we shall see that, for a spline based discretization with its knots on the delays, $|r_N(i\omega, -\tau_k)| = 1$ holds for all delays τ_k and $\omega \in \mathbb{R}$, if the underlying polynomials are symmetric on their domains. This suggests that this corollary can be extended to multiple delays when using a spline based discretization, as we shall detail in that chapter. To keep the formulations similar between the chapters, we used the strong definitions of exponential stability and the H^2 -norm in Theorem 5.3.6 and Corollary 5.3.7, even though these are identical to the non-strong versions in the case of a single delay system.

Having established conditions for convergence, we conclude this section discussing the convergence rate. As mentioned before, in the case of retarded type systems, Vanbiervliet, Michiels, et al. (2011) report third order algebraic convergence of $\|G_N\|_{H^2}$ to $\|G\|_{H^2}$ and provide partial arguments for this observation. Whilst similar derivations to theirs are likely possible for our extension, we will limit ourselves to arguments based on the properties of the delay-Lyapunov equation discussed in section 4.4.

For our illustrations we shall use the following group of systems

$$\begin{aligned} \dot{\mathbf{x}}(t) = & \begin{pmatrix} -5 & 1 \\ 3 & 8 \end{pmatrix} \mathbf{x}(t) + \delta_{r_1} \begin{pmatrix} -2 & 0 \\ -2 & 1 \end{pmatrix} \mathbf{x}(t-1) - \delta_{r_2} I_2 \mathbf{x}(t-1.9) \\ & + \delta_{n_1} \begin{pmatrix} -1/5 & 1/10 \\ 0 & -1/10 \end{pmatrix} \dot{\mathbf{x}}(t-1) - \frac{1}{10} \delta_{n_2} I_2 \dot{\mathbf{x}}(t-1.9) \\ & + \begin{pmatrix} 1 \\ 1 \end{pmatrix} u(t), \\ y(t) = & \begin{pmatrix} 1 & 1 \end{pmatrix} \mathbf{x}(t), \end{aligned} \quad (5.11)$$




Figure 5.1: Convergence rate for different RDDE configurations of system (5.11). The dashed curve corresponds to the single delay case, $\vec{\delta} = (1, 0; 0, 0)$, the solid curve to two delays, $\vec{\delta} = (1, 1; 0, 0)$.

where $\vec{\delta} = (\delta_{r_1}, \delta_{r_2}; \delta_{n_1}, \delta_{n_2}) \in \{0, 1\}^4$ is used to determine the system's composition. Introducing the slack variables $x_3(t) = \dot{x}_1(t)$ and $x_4(t) = \dot{x}_2(t)$, we can write this system in DDAE form as

$$\begin{pmatrix} I_2 & 0 \\ I_2 & 0 \end{pmatrix} \dot{\mathbf{x}}(t) = \begin{pmatrix} \begin{pmatrix} -5 & 1 \\ 3 & 8 \end{pmatrix} & 0 \\ 0 & I_2 \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} \delta_{r_1} \begin{pmatrix} -2 & 0 \\ -2 & 1 \end{pmatrix} & \delta_{n_1} \begin{pmatrix} -1/5 & 1/10 \\ 0 & -1/10 \end{pmatrix} \\ 0 & 0 \end{pmatrix} \mathbf{x}(t-1) \\ - \begin{pmatrix} \delta_{r_2} I_2 & \frac{1}{10} \delta_{n_2} I_2 \\ 0 & 0 \end{pmatrix} \mathbf{x}(t-1.9) + \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} u(t), \\ y(t) = (1 \ 1 \ 0 \ 0) \mathbf{x}(t).$$

We will use the appropriately shifted and scaled N th Legendre polynomial P_N as ϕ_N for our discretization, as these have the most developed theory and appeared to perform best overall in chapter 4.

From our discussion in sections 4.3 and 4.6, we would expect geometric convergence in the single delay RDDE and NDDE cases, i.e. $m = 1$, as the Legendre polynomials are symmetric and the solution of the delay-Lyapunov equation is analytic away from zero. For an RDDE with multiple discrete delays we know this to reduce to third order algebraic convergence. Both are confirmed on figs. 5.1 and 5.2.

As systems of neutral type have vertical chains of infinitely many characteristic roots parallel to the imaginary axis and the H^2 -norm involves an integral along this axis, a further reduction in algebraic convergence rate is expected for such systems. Also from the perspective




Figure 5.2: Convergence rate for different NDDE configurations of system (5.11). The dashed curve has a single delay, $\vec{\delta} = (0, 0; 1, 0)$, the solid curve two, $\vec{\delta} = (0, 0; 1, 1)$.

of the delay-Lyapunov equation the loss of a second derivative suggests slower convergence. Indeed, on fig. 5.2 we see a reduction to first order algebraic convergence for multiple delays. As mentioned in section 4.6, this shows a discrepancy between RDDEs and NDDEs. Whereas the H^2 -norm of an RDDE converges at one order higher than the full solution of the delay-Lyapunov equation, as is typical if only the value in the end point of the solution is needed, NDDEs do not seem to gain this order.

On fig. 5.3 we see that the algebraic convergence rate is only determined by the type of the inner delays. If any of those correspond to terms that render the equation of neutral type, we reduce to first order convergence, else we maintain third order convergence. This can again be understood from the delay-Lyapunov equation. The jumps in the first derivative only appear at combinations of delays corresponding to neutral terms, if there is only one such delay and it is the maximal delay, there are only jumps in the first derivative at the end points of the approximation interval, which do not impact the convergence rate.

5.4 SUMMARY

In this chapter we extended the method of chapter 4 to approximate the H^2 -norm of an DDAE system (5.1) of index one, using a Lanczos tau discretization as proxy. We showed that, under the assumption of a finite strong H^2 -norm, this approach is sound (Theorem 5.2.4). Furthermore,




Figure 5.3: Convergence rate for different mixed configurations of system (5.11). The dashed curve has a neutral term on the maximal delay, $\vec{\delta} = (1, 0; 0, 1)$, the solid curve on the interior delay, $\vec{\delta} = (0, 1; 1, 0)$.

we were able to extend Theorem 4.1.3 and prove convergence under reasonable assumptions (Theorem 5.3.5). For the specific case of a Legendre based discretization of a single delay system, we were able to formally prove preservation of stability (Theorem 5.3.6) and the other requirements (Corollary 5.3.7). Underlying the convergence of the H^2 -norm is the convergence of the transfer function established in Theorem 5.3.3 which is thus the natural extension of Theorem 3.4.7.

Whilst an NDDE with a single delay still displays geometric convergence for a symmetric basis, we see, like in section 4.6, that multiple discrete delays slow convergence to algebraic rates. For RDDEs this is the third order algebraic convergence we saw before (fig. 5.1), for NDDEs we obtain only first order convergence (fig. 5.2). These algebraic rates are determined only by the type of the delays on the interior. Hence, an NDDE with multiple delays, but where the derivative is only evaluated in the maximal delay, still gives third order convergence (see fig. 5.3).

Although we did not treat it explicitly here—as we do not know of any analogues to Theorems 5.1.3 and 5.1.6—the presented method can be straightforwardly extended to more general functional differential algebraic equations of index one, by replacing $\sum_{k=0}^m A_k[\varepsilon - \tau_k]$ by the discretized functional in (5.4), like we did in section 4.6.

6 Accelerating H^2 -norm convergence using splines

At the end of section 5.3, we proved convergence for the specific case of a DDAE of index one (in the sense of Definition 5.0.1) with a single delay, when discretized using Legendre polynomials (Corollary 5.3.7). The argument mainly relied on the fact that, for a single delay, the rational approximation of the exponential underlying this discretization has magnitude one on the imaginary axis. This property is also the one of which we hypothesized in chapter 4 that it allowed the improvement from algebraic to geometric convergence of the H^2 -norm.

In section 4.6, we noted that the properties of the solution of the delay-Lyapunov equation suggest that using a spline might result in improved convergence rates for RDDEs with multiple delays. In this chapter we shall see that using a discretization based on splines, with the knots at the delays, indeed improves convergence rates and also gives us modulus one at all delays. We sketch how the latter property can be used to adapt the theory developed in the preceding chapters to splines, often ending up simpler and more general. Finally, we conclude by illustrating that the convergence rate can be improved further still by introducing additional knots.

Overview. In section 6.1 we adapt the Lanczos tau method of chapter 5 to use a spline with knots at the delays instead of a single polynomial. In the subsequent section, we show that this results in a similar interpretation in the frequency domain as we saw in the previous chapters. These properties are used to prove convergence of the H^2 -norm of a DDAE

The results of this chapter previously appeared in Provoost, E., and Michiels, W. (2025), “Accelerated H^2 -norm approximation for time-delay systems with discrete delays”, *IFAC-PapersOnLine*, 59/13, Presented at the *19th IFAC Workshop on Time-Delay Systems*: 93–8 and in the third part of Provoost, E., and Michiels, W. (2026), “Computing and optimizing the H^2 -norm of delay differential algebraic systems”, arXiv: 2603.10734 [math.NA]. The connection between convergence rates and the properties of the solution of delay-Lyapunov equation and the introduction of additional knots at integer combinations of delays are novel to this thesis.

of index one with multiple discrete delays in section 6.3 when using a Legendre basis. In this section, we also observe accelerated convergence rates compared to what we saw in the previous chapter. Finally, in section 6.4 we see that placing additional knots at integer combinations of delays can improve convergence further, up to geometric in the case of commensurate delays.

6.1 A SPLINE BASED LANCZOS TAU METHOD

In Ito and Teglas (1987) and Ito and Kappel (1992) the authors already propose to use splines in a tau method for a system with discrete delays, where the spline is allowed to be discontinuous to prevent jump discontinuities in the history function from resulting in a poor approximation. In our application, however, the latter does not occur; similarly to the spline pseudospectral collocation method of Breda, Maset, et al. (2005), we will rather require the spline to be continuous. Let Ξ_{tN} be a set of polynomials $\{\xi_{tN}^{(k)} : [-\tau_k, -\tau_{k-1}] \rightarrow \mathbb{C}\}_{k=1}^m$, each of degree N , i.e. a spline with its knots at the delays. Analogous to section 5.2, we can discretize the DDAE system (5.1) by expanding each $\xi_{tN}^{(k)}$ in an orthogonal, degree-graded basis $\{\phi_j^{(k)} : [-\tau_k, -\tau_{k-1}] \rightarrow \mathbb{C}\}_{j=0}^N$. As mentioned, we complete the discretization by requiring continuity at the knots, i.e. $\xi_{tN}^{(k)}(-\tau_k) = \xi_{tN}^{(k+1)}(-\tau_k)$ for $k = 1, \dots, m-1$, giving us

$$\begin{pmatrix} E\varepsilon_0^{(1)} \\ \left(\mathcal{T}_{N-1}^{(k)}\right)_{k=1}^m \\ \left(\mathbf{0}\right)_{k=1}^{m-1} \end{pmatrix} \dot{\Xi}_{tN} = \begin{pmatrix} A_0\varepsilon_0^{(1)} + \sum_{k=1}^m A_k\varepsilon_{-\tau_k}^{(k)} \\ \left(\mathcal{D}^{(k)}\right)_{k=1}^m \\ \left(\varepsilon_{-\tau_k}^{(k)} - \varepsilon_{-\tau_k}^{(k+1)}\right)_{k=1}^{m-1} \end{pmatrix} \Xi_{tN} + \begin{pmatrix} B \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix} \mathbf{u}(t), \quad (6.1)$$

$$\mathbf{y}_N(t) = C\varepsilon_0^{(1)}\Xi_{tN},$$

where $\mathcal{P}^{(k)}$ indicates that $\mathcal{P}$ is applied to segment $\xi_{tN}^{(k)}$ and $(\mathcal{P}_k)_{k=1}^m$ is the vertical concatenation of $\{\mathcal{P}_1, \dots, \mathcal{P}_m\}$.

Like the polynomial Lanczos tau method, the spline version is mathematically equivalent to the spline extension of pseudospectral collocation as introduced in Breda, Maset, et al. (2005). In their work, an ordinary differential equation is obtained directly from an RDDE, i.e. they do not obtain a DAE, by having both end points of a segment as collocation nodes, which allows the elimination of algebraic constraints. Whilst a Lanczos tau method corresponds to collocation of the relation $\frac{d}{dt}\xi_t(\theta) = \frac{d}{d\theta}\xi_t(\theta)$ in the zeroes of the $\phi_N^{(k)}$ (analogous to Theorem 3.5.2), we saw before that orthogonal polynomials cannot have zeroes in the end points and hence

we cannot collocate the derivative in the delays. This simplification can thus not be used in a Lanczos tau method based on orthogonal polynomials. However, we shall see in section 6.3 that this is generally compensated for by the increase in convergence rate.

6.2 PROPERTIES IN THE FREQUENCY DOMAIN

Before computing the H^2 -norm using this spline discretization, we first take a closer look at what happens in the frequency domain. Taking the Laplace transform of discretization (6.1) gives:

$$\begin{cases} sE\varepsilon_0^{(1)}\hat{\Xi}_{sN} = \left(A_0\varepsilon_0^{(1)} + \sum_{k=1}^m A_k\varepsilon_{-\tau_k}^{(k)}\right)\hat{\Xi}_{sN} + B\hat{\mathbf{u}}(s), \\ s\mathcal{T}_{N-1}^{(k)}\hat{\Xi}_{sN} = \mathcal{D}^{(k)}\hat{\Xi}_{sN}, \quad k = 1, \dots, m, \\ \mathbf{0} = \left(\varepsilon_{-\tau_k}^{(k)} - \varepsilon_{-\tau_k}^{(k+1)}\right)\hat{\Xi}_{sN}, \quad k = 1, \dots, m-1, \\ \hat{\mathbf{y}}_N(s) = C\varepsilon_0^{(1)}\hat{\Xi}_{sN}, \end{cases} \quad (6.2)$$

where $\hat{f}(s)$ is the transform of $f(t)$.

Analogous to Proposition 5.3.1, we can prove that the transfer function of (6.1) is given by

$$G_N^{\text{spl}}(s) = C \left(sE - A_0 - \sum_{k=1}^m A_k r_N^{(k)}(s, -\tau_k) \right)^{-1} B,$$

where $r_N^{(k)}(s, \theta)$ is polynomial in $\theta \in [-\tau_k, -\tau_{k-1}]$ and a rational function of s , such that

$$\begin{cases} r_N^{(1)}(s, 0) = 1, \\ \mathcal{D}r_N^{(k)}(s, \cdot) = s\mathcal{T}_{N-1}r_N^{(k)}(s, \cdot), \quad k = 1, \dots, m, \\ r_N^{(k)}(s, -\tau_{k-1}) = r_N^{(k-1)}(s, -\tau_{k-1}), \quad k = 2, \dots, m. \end{cases}$$

Like Proposition 5.3.1, the argument relies on noting that we can write $\hat{\Xi}_{sN}$ in terms of these $r_N^{(k)}$. The conditions on the latter then end up corresponding to the middle two equations of (6.2).

As a corollary, this discretization has a distinct rational approximation $r_N^{(k)}(s, -\tau_k)$ of $s \mapsto e^{-\tau_k s}$ for each delay interval $[-\tau_k, -\tau_{k-1}]$. Hence, the poles—and even the type—of the rational approximations need not be the same for all delay intervals, which is unlike the case with a single polynomial where only the zeroes can change. This is clear from their respective domain-colourings, compare, for instance, fig. 6.1 to fig. 3.1.

We can rewrite this rational function into a more explicit form. Let $\rho_N^{(k)}(s, \theta)$ be a rational function of s and polynomial in $\theta \in [-\tau_k, -\tau_{k-1}]$,




Figure 6.1: Domain-colourings of $r_N^{(2)}(s, -\tau_2)$ (left) and $r_N^{(1)}(s, -\tau_1)$ (right) for $\phi_N = P_N$, $\tau_1 = 0.675$, $\tau_2 = 1$, and $N = 6$. Note that the poles and zeroes of the right figure are also present on the left and that these are symmetric sets for both delays. Compare this to the analogous results for a polynomial discretization on fig. 3.1.




Figure 6.2: Domain-colouring of $r_N^{(2)}(s, -\tau_2)$ for $\phi_N = P_N$, $\tau_1 = 1$, $\tau_2 = 2$, i.e. equidistant delays, and $N = 6$. Note that the poles and zeroes are all of order two.

such that

$$\begin{cases} \rho_N^{(k)}(s, -\tau_{k-1}) = 1, \\ \mathcal{D}\rho_N^{(k)}(s, \cdot) = s\mathcal{T}_{N-1}\rho_N^{(k)}(s, \cdot). \end{cases}$$

It is then easy to see that the rational approximants of the exponential in G_N^{spl} are given by $r_N^{(1)} = \rho_N^{(1)}$ and

$$r_N^{(k)}(s, \cdot) = r_N^{(k-1)}(s, -\tau_{k-1})\rho_N^{(k)}(s, \cdot), \quad k = 2, \dots, m. \quad (6.3)$$

Hence, each $r_N^{(k)}(s, -\tau_k)$ has the preceding rational approximations, i.e. those with lower k , as factors. As $\rho_N^{(k)}(s, \theta)$ is a type (N, N) rational function of s , $r_N^{(k)}(s, \theta)$ is of type (kN, kN) . When the delays are equidistant, i.e. $\tau_k = k\tau$ for some $\tau > 0$, we find $\rho_N^{(k)}(s, \theta) = \rho_N^{(k-1)}(s, \theta + \tau)$, and thus $r_N^{(k)}(s, -k\tau) = (r_N^{(1)}(s, -\tau))^k$, as seen on fig. 6.2.

From Proposition 3.4.2, we have the closed form expression

$$\rho_N^{(k)}(s, \theta) = \frac{\sum_{j=0}^N \frac{d^{N-j}}{d\theta^{N-j}} \phi_N^{(k)}(\theta) s^j}{\sum_{j=0}^N \frac{d^{N-j}}{d\theta^{N-j}} \phi_N^{(k)}(0) s^j},$$

where the superscript of $\phi_N^{(k)}$, contrary to Proposition 3.4.2, indicates the relevant delay interval, not the derivative. Hence, we also have a closed form expression for all $r_N^{(k)}(s, \theta)$. From this and the discussed structure of $r_N^{(k)}$, we straightforwardly see, for symmetric basis functions $\phi_N^{(k)}$ (in the sense of Assumption 4.1.1) that

$$|r_N^{(k)}(i\omega, -\tau_k)| = 1, \quad k = 1, \dots, m, \quad (6.4)$$

for all $\omega \in \mathbb{R}$.

Note that, with the exception of $r_N^{(1)}(s, 0) = 1$, all $r_N^{(k)}$ in G_N^{spl} are only evaluated with θ in their left end point, i.e. $\theta = -\tau_k$. This, together with the above result, then means that, for symmetric basis functions, the rational approximants in G_N^{spl} have a better qualitative match with the exponential—in the sense that their magnitudes match—than if using a non-symmetric set of $\phi_N^{(k)}$ or a single polynomial for multiple delay intervals. This is one the properties that we presume allowed for the geometric convergence observed for the single delay case in sections 4.3 and 5.3. We shall see in the next section that the convergence is also accelerated for multiple discrete delays by the use of splines based on symmetric polynomials.

6.3 CONVERGENCE OF THE H^2 -NORM

The properties we saw in the last section simplify the arguments needed to prove convergence of the H^2 -norm for the polynomial case in section 5.3 somewhat. First, note that the absence of feedthrough under the condition of strong stability becomes rather straightforward using the decomposed form of G_N (5.10), which, *mutatis mutandis*, applies to G_N^{spl} as well. Using (6.4) and Theorem 5.1.3, it is easily seen that $A_{N,22}(i\omega)$ is invertible for all $\omega \in \mathbb{R}$. As $D_N(i\omega)^{-1} \rightarrow 0$ as $\omega \rightarrow \infty$, the feedthrough term is given by

$$\lim_{\omega \rightarrow \infty} G_N^{\text{spl}}(i\omega) = - \lim_{\omega \rightarrow \infty} C_2 A_{N,22}(i\omega)^{-1} B_2.$$

Using the argument of Theorem 5.2.4, we readily have

$$C_2 A_{N,22}(i\omega)^{-1} B_2 = \mathbf{0}$$

for all ω and so too at infinity.

The theorems on the convergence of our H^2 -norm approximation from section 5.3 also readily extend to splines. Better yet, most of the required assumptions now trivially hold. The bound of Assumption 5.3.2 is only required to hold in the delays, which is trivially true when the underlying polynomials are symmetric, because of (6.4). From the same property and Theorem 5.1.3, Assumption 5.3.4 holds.

All that is left to show is thus convergence of the $r_N^{(k)}$, i.e. Assumption 3.4.3, and preservation of stability. As in section 5.3, we can do so when we use splines based on Legendre polynomials. Using an analogous argument to the one used to establish the convergence of r_N in Corollary 3.4.4, we can show that $\rho_N^{(k)}(s, \theta) \rightarrow e^{(\theta - \tau_{k-1})s}$ as $N \rightarrow \infty$. Convergence of $r_N^{(k)}$ then follows from (6.3) and $e^a e^b = e^{a+b}$. For preservation of stability, note that for this choice of basis, equation (6.3) implies that $r_N^{(k)}(i\omega, -\tau_k)$ is the product of k diagonal Padé approximants of various decaying exponential functions. The argument and conclusion of Theorem 3.8.1 thus readily extend to this case, giving the following result.

Theorem 6.3.1. *For a strongly exponentially stable DDAE system (5.1) of index one, the discretization (6.1) preserves stability for N sufficiently large, when using a spline based on a Legendre basis.*

From the preceding discussion we thus have the following generalization of Corollary 5.3.7.




Figure 6.3: Convergence rate for the two delay RDDE configuration of system (5.11), i.e. $\vec{\delta} = (1, 1; 0, 0)$. The dashed line uses a single polynomial for the entire interval and the solid line uses a spline with a separate polynomial for each delay interval; the latter approximately follows the grey fourth order reference line. Both use Legendre bases.

Corollary 6.3.2. *For a strongly exponentially stable DDAE system (5.1) of index one with finite strong H^2 -norm, the method described in section 5.2 converges, i.e. $\|G_N^{\text{spl}}\|_{H^2} \rightarrow \|G\|_{H^2}$ for $N \rightarrow \infty$, when using discretization (6.1) with a spline based on a Legendre basis.*

Moreover, as alluded to before, the benefits of the spline discretization are not merely to be found in their theory. If we repeat the experiments involving multiple delays of section 5.3 again, we notice accelerated convergence on figs. 6.3 and 6.4. In both cases we gain a little over a full algebraic order, going from third to fourth order for RDDEs and from first to second order for NDDEs. Comparing fig. 6.5 to fig. 5.3, we also see that, as for a discretization based on a single polynomial, the delay in the endpoint does not impact the algebraic convergence rate.

The improvement in the RDDE case is consistent with our discussion in section 4.6. The knots of the spline coincide with the jumps in the second derivative of the solution of the delay-Lyapunov equation, hence only jumps in the third derivative remain on the polynomial segments. Similarly, for the same reasons as mentioned in section 5.3, we would not expect the type of the maximal delay to influence the algebraic convergence rate. More surprising is the improvement we see for NDDEs. The number of continuous derivatives is still only one on the polynomial segments. Note however, that the rate we now attain for




Figure 6.4: Convergence rate for a two delay NDDE configuration of system (5.11), with $\vec{\delta} = (0, 0; 1, 1)$. Like fig. 6.3, the dashed line uses a single polynomial for the entire interval and the solid line uses a spline with a separate polynomial for each delay interval; the latter approximately follows the grey second order reference line. Both use Legendre bases.




Figure 6.5: The same experiment as on fig. 5.3 but using a spline discretization using Legendre bases. The dashed curve has a neutral term on the maximal delay, the solid curve on the interior delay.

NDDES follows the same pattern as for RDDES, namely one higher than the convergence rate of the solution of the delay-Lyapunov equation for a spectral method. Our hypothesis is thus that splines undo whatever was reducing the convergence rate by one in the NDDE case before. This will be corroborated by the results of the next section, where adding additional knots at integer combinations of the delays improves the convergence rate further for RDDES but fails to do so for NDDES.

Finally, as solving the Lyapunov equation dominates the time complexity, the method requires $O(M^3)$ operations, with M the number of rows in the discretization (Hammarling 1982). Using a spline with its knots at the delays gives a dependence on the number of delays m , and is hence more expensive than the single polynomial, requiring $O(m^3 n^3 N^3)$ and $O(n^3 N^3)$ operations, respectively. However, as this is only a constant factor for a given system, a spline approach will generally need less operations for the same level of accuracy, especially when high accuracy is required or when there are only a moderate number of delays.

6.4 FURTHER ACCELERATION USING ADDITIONAL KNOTS

As mentioned in the previous section, the properties of the solution of the delay-Lyapunov equation suggest that we might obtain further improvements in convergence order for RDDES by placing knots not only at the delays but also at integer combinations of those delays.

From the discussion in section 4.6, we know that such a combination of delays results in a jump in the $(k+1)$ th derivative, with k the absolute sum of the coefficients. Thus, if we add knots at all combinations with $k \leq \nu$, we would expect $(\nu+2)$ th algebraic convergence. On fig. 6.6 we see the resulting convergence rates for $\nu = 2$. As expected, the convergence accelerates further for the RDDE. For NDDES there are jumps in the first derivative for all k and convergence thus remains of second order. Compared to fig. 6.4, we do see a constant gain in accuracy, likely due to the remaining jumps being smaller (see section 4.6). As the effect is fairly strong, small values for ν are generally worth it in the RDDE case if the number of combinations is not too high. Similarly, the constant gain for NDDES can also be larger than the cost of the larger discretization, such as between figs. 6.4 and 6.6.

Whilst it is impossible to place a knot at every jump for non-commensurate delays, i.e. there is no maximal k , this is possible in the commensurate case. We will see in the remainder of this section that this will allow us to recover the geometric convergence seen in the single delay case.

Remember that delays are commensurate when there is some $h > 0$ such that $\tau_k = d_k h$ with $d_k \in \mathbb{N}$ for $k = 1, \dots, m$. For the largest such h ,




Figure 6.6: Convergence rate for the RDDE example of fig. 6.3 (dashed) and the NDDE example of fig. 6.4 (solid) with the delay combinations for which $k \leq 2$ as knots of the spline, i.e. the points 0, $\tau_2 - \tau_1 = 0.9$, $\tau_1 = 1$, and $\tau_2 = 1.9$.




Figure 6.7: The same experiment as on fig. 6.5 but with $\tau_2 = 2$, i.e. with equidistant delays.

the jumps are located at the integer multiples of h . Thus, if we place a knot at all multiples of h , the solution of the delay-Lyapunov equation is infinitely differentiable on the interior of the polynomial segments. Under these conditions we would expect to attain geometric convergence. Indeed, on fig. 6.7 we see geometric convergence for both an RDDE and NDDE example with equidistant delays, i.e. with $h = \tau_1$ and $d_k = k$.

Given the dramatic increase in convergence rate, including all integer multiples of h is usually worth it. However, if the ratio between τ_m and h , i.e. d_m , is very large, the size of the resulting discretization can become prohibitive in practice.

Finally, note that the construction we propose is closely related to the polynomial discretization of the reformulation of the DDAE as one with only a single delay. This reformulation is often used to extend results in the single delay case to commensurate delays (see e.g. Unger 2020; Remark 1.2) and is given by

$$\begin{cases} E\dot{\mathbf{x}}_1(t) = A_0\mathbf{x}_1(t) + \sum_{k=1}^m A_k\mathbf{x}_{d_k}(t-h) + B\mathbf{u}(t), \\ \mathbf{x}_2(t) = \mathbf{x}_1(t-h), \\ \mathbf{x}_3(t) = \mathbf{x}_2(t-h), \\ \vdots \\ \mathbf{x}_{d_m}(t) = \mathbf{x}_{d_m-1}(t-h), \\ \mathbf{y}(t) = C\mathbf{x}_1(t). \end{cases}$$

Writing the additional algebraic constraints explicitly we find as polynomial Lanczos tau discretization (5.4) of this DDAE:

$$\begin{pmatrix} E\varepsilon_0^{(1)} \\ (\mathbf{0})_{j=1}^{d_m-1} \\ \mathcal{T}_{N-1} \end{pmatrix} \dot{\xi}_{tN} = \begin{pmatrix} A_0\varepsilon_0^{(1)} + \sum_{k=1}^m A_k\varepsilon_{-\tau}^{(d_k)} \\ \left(\varepsilon_{-\tau}^{(j)} - \varepsilon_0^{(j+1)} \right)_{j=1}^{d_m-1} \\ \mathcal{D} \end{pmatrix} \xi_{tN} + \begin{pmatrix} B \\ \mathbf{0} \\ \mathbf{0} \end{pmatrix} \mathbf{u}(t),$$

$$\mathbf{y}_N(t) = C\varepsilon_0^{(1)} \xi_{tN},$$

where ξ_{tN} maps $[-h, 0]$ to $\mathbb{C}^{d_m n}$ and, unlike equation (6.1), $\varepsilon_\theta^{(j)} \xi_{tN}$ selects the subvector from $\xi_{tN}(\theta)$ which corresponds to $\mathbf{x}_j$.

This slight abuse of notation helps to crystallize the connection between this discretization and the spline discretization (6.1) we would obtain by placing a knot at every integer multiple of h . Note first that we can regroup $\xi_{tN} : [-h, 0] \rightarrow \mathbb{C}^{d_m n}$ into d_m polynomials $\chi_{tN}^{(j)} : [-h, 0] \rightarrow \mathbb{C}^n$, each satisfying $\mathcal{T}_{N-1}\dot{\chi}_{tN}^{(j)} = \mathcal{D}\chi_{tN}^{(j)}$. Through a change of variables we can also think of these as the polynomials $\xi_{tN}^{(j)} : [-jh, -(j-1)h] \rightarrow \mathbb{C}^n$,

which, together with the algebraic constraints, means that they form a continuous spline Ξ_{tN} from $[-d_m h, 0] = [-\tau_m, 0]$ to $\mathbb{C}^n$. This way, we indeed recover (6.1) with knots at $\{0, h, 2h, \dots, d_m h\}$.

6.5 SUMMARY

We discussed how the use of a spline instead of a single polynomial in the Lanczos tau method has both theoretical and practical benefits for the approximation of the H^2 -norm.

In section 6.2, we showed that using a spline, with its knots at the delays, corresponds in the frequency domain to using distinct rational approximants of the exponential functions for each non-zero delay, which is unlike the Lanczos tau method with a single polynomial, where all approximants share the same poles and only the zeroes change with the delay.

We derived a closed form expression for these rational approximations. From the structure of these rational functions, we showed that the modulus along the imaginary axis equals one in the delays (see eq. (6.4)), which matches that of the exponential.

Because of these properties, such a spline based discretization tends to simplify theory and, when using a spline based on a Legendre basis, allows us to prove preservation of stability (Theorem 6.3.1) and convergence (Corollary 6.3.2) for DDAEs with multiple delays, extending the single delay results of chapter 5.

Furthermore, we showed that a spline based Lanczos tau method can partially extend the geometric convergence of the H^2 -norm seen for single delay systems in chapters 4 and 5 to multiple discrete delays. We saw convergence that is accelerated by a little over a full algebraic order for general multiple discrete delays if we take knots at the delays. In section 6.4, we saw that additionally including integer combinations of delays can yield further acceleration up to, in the commensurate case, geometric convergence. Whilst the spline discretization results in a larger discretization (by a factor equal to the number of knots), this cost is typically offset by the accelerated convergence.

7 H^2 -optimal synthesis

The goal of the past few chapters has been an effective method to compute the H^2 -norm of the DDAE (5.1). The aim of this chapter is to use this method to synthesize H^2 -optimal fixed-order feedback controllers and stable approximate models through direct optimization of the model's parameters.

In the first application, system (5.1) stems from the feedback interconnection of a plant and a dynamic or static feedback controller. In the second, the system describes the error between the dynamics of the full and the approximate model (see section 2.5 and example 7.2.3). As a consequence, some elements of the matrices A_0 , ..., A_m , B , and C may be equal to, or depend on, the controller or model's parameters. To efficiently optimize the H^2 -norm using a gradient based method, we will thus analyse the sensitivities of this norm with respect to changes of these matrices. Likewise, the application to delay based controllers (see e.g. Villafuerte et al. 2013), where intentional delays are used to improve the control performance, motivates the study of delay sensitivity as well.

The two existing methods to approximate the H^2 -norm of a time-delay system mentioned at the beginning of chapter 4, i.e. the discretize-first strategy of Vanbiervliet, Michiels, et al. (2011), that we have extended, and the delay-Lyapunov based method of Jarlebring, Vanbiervliet, et al. (2011), already include such expressions and give thus effective methods for H^2 -optimal synthesis, be it for a less general class of models.

Whilst H^2 -optimal model order reduction can be rephrased as a synthesis problem, specialized methods have also been developed where the targeted approximation is already a DDAE. However, to the best of our knowledge and unlike the approach we shall present, these have so far limited the types of systems under consideration. Pontes Duff et al. (2018) extend the celebrated iterative rational Krylov algorithm (IRKA) of Gugercin et al. (2008) to systems with input and output delays.

Data-driven approaches, using time-delay extensions of the Loewner framework, also exist, where, unlike the previous and our approach, only

The contents of this chapter previously appeared as the second part of Provoost, E., and Michiels, W. (2026), "Computing and optimizing the H^2 -norm of delay differential algebraic systems", arXiv: 2603.10734 [math.NA].

samples of the transfer function are needed instead of the entire state space representation. The TF-IRKA of Beattie and Gugercin (2012), the data-driven analogue of IRKA, is extended to systems whose dynamics only depend on the state at some fixed delay (and thus not the current state) in Pontes Duff et al. (2015). Such extensions of the Loewner framework can also be used directly to find reduced order models of systems with multiple delays, be it not necessarily H^2 -optimal, as in Schulze et al. (2018).

To keep our notation light, we shall stick to the polynomial based discretization of chapter 5. The exact same methodology would work for the spline based discretization of the previous chapter, merely requiring us to additionally keep track of which operator acts on which segment.

Overview. In section 7.1 we derive formulas for the gradient of the H^2 -norm with respect to changes of the system matrices $A_0, \dots, A_m, B$, and C , and the delays. In section 7.2 we illustrate the effectiveness of these expressions with a number of applications, both to robust control as to the synthesis of stable approximate models.

7.1 COMPUTING THE GRADIENT

In this section we discuss the computation of the derivative of the square of the H^2 -norm with respect to system parameters.

As a norm is always positive, minimizing its square also minimizes the norm. Some care must be taken as the H^2 -norm of a system of the form (5.1) may not be everywhere differentiable with respect to a parameter p . If the system matrices are smooth functions of p , we keep E fixed, and we assume the system to both have at most index one (see Definition 5.0.1) and a finite strong H^2 -norm (Theorem 5.1.6) in a neighbourhood of p , its square is a locally smooth function of p . Under the latter condition, its square is also a smooth function of the delays. In summary, if the conditions of Theorem 5.2.4, ensuring that the Lanczos tau discretization does not introduce feedthrough, are satisfied in a neighbourhood of the parameters, the required derivatives are well defined.

Approximating the H^2 -norm once, using the approach of section 5.2, takes $O(n^3 N^3)$ elementary operations when using a direct solver for the Lyapunov equation, such as the one by Hammarling (1982). Using finite differences we could then approximate the gradient with respect to the free parameters, for use in an optimization loop, in $O((r+1)n^3 N^3)$ time, where r is the number of parameters.

An alternative method, with a time complexity independent of r , obtains analytical expressions for the derivatives using an adjoint method,

by solving the dual Lyapunov equation. Hence, we can compute the gradient in a little over double the computational time of only computing the H^2 -norm itself. A disadvantage of this approach is that we compute the derivative of the discretization's H^2 -norm and not of the original system. When used in an optimization loop, some care must thus be taken to reject updates that increase the index beyond one or blow up the strong H^2 -norm.

The approach is similar to the one of Vanbiervliet, Vandereycken, et al. (2009), however we opt to use an alternative derivation using matrix differentials, as introduced by Magnus and Neudecker (1985), which makes the calculus a bit less tedious. Throughout, we will assume the convention $\left(\frac{df}{dX}\right)_{ij} = \frac{\partial f}{\partial x_{ij}}$.

Assume the notation of equation (5.6) and let $\mathcal{L}P = \tilde{A}P\mathcal{E}_{11}^T + \mathcal{E}_{11}P\tilde{A}^T$, then we can write the primary and dual Lyapunov equations (see section 4.2) as

$$\begin{aligned}\ell_P &:= \mathcal{L}P + \tilde{B}\tilde{B}^T = \mathbf{0}, \quad \text{and} \\ \ell_Q &:= \mathcal{L}^T Q + \tilde{C}^T \tilde{C} = \mathbf{0},\end{aligned}$$

respectively.¹

Let X be some variable independent of P and Q . Assuming everything but P and X fixed, we get from the top equation the following linear relation between infinitesimal perturbations of P and X and the resulting infinitesimal perturbation of ℓ_P ; written in terms of differentials this is

$$d\ell_P = \frac{\partial \ell_P}{\partial P} dP + \frac{\partial \ell_P}{\partial X} dX = \mathcal{L} dP + \frac{\partial \ell_P}{\partial X} dX = \mathbf{0},$$

which implies

$$dP = -\mathcal{L}^{-1} \frac{\partial \ell_P}{\partial X} dX.$$

We can assume the existence of $\mathcal{L}^{-1}$ as we only consider this relation when the solutions P and Q exist.

Rewriting $\ell_Q = \mathbf{0}$ we get $Q = -\mathcal{L}^{-T} \tilde{C}^T \tilde{C}$, thus, using $Q = Q^T$,

$$\tilde{C}^T \tilde{C} dP = Q \frac{\partial \ell_P}{\partial X} dX.$$

Now note that $d\|G_N\|_{H^2}^2 = \text{tr}(\tilde{C} dP \tilde{C}^T) = \text{tr}(\tilde{C}^T \tilde{C} dP)$, hence

$$d\|G_N\|_{H^2}^2 = \text{tr}\left(Q \frac{\partial \ell_P}{\partial X} dX\right).$$

¹Note that indeed $\mathcal{L}^T Q = \tilde{A}^T Q \mathcal{E}_{11} + \mathcal{E}_{11}^T Q \tilde{A}$ as $\mathcal{L}^T \cong (\mathcal{E}_{11} \otimes \tilde{A} + \tilde{A} \otimes \mathcal{E}_{11})^T$.

Using analogous reasoning, we find from $\|G_N\|_{H^2}^2 = \text{tr}(\tilde{B}^T Q \tilde{B})$ that

$$d\|G_N\|_{H^2}^2 = \text{tr}\left(P \frac{\partial \ell_Q}{\partial X} dX\right).$$

The required derivatives are finally found from simple, but tedious, computations detailed in the next section and the fact that $df = \text{tr}(Y^T dX)$ implies $\frac{df}{dX} = Y$. We have

$$\begin{aligned} \frac{d\|G_N\|_{H^2}^2}{dA_k} &= 2\left(\begin{smallmatrix} I_n \\ \mathbf{0} \end{smallmatrix}\right)^T \mathcal{D}_S^T[\varepsilon - \tau_k]^T, \\ \frac{d\|G_N\|_{H^2}^2}{d\tau_k} &= -2 \text{tr}(A_k[\varepsilon - \tau_k] \mathcal{D}) \mathcal{D}_S\left(\begin{smallmatrix} I_n \\ \mathbf{0} \end{smallmatrix}\right) \quad \text{for } k < m, \\ \frac{d\|G_N\|_{H^2}^2}{d\tau_m} &= -2 \text{tr}\left(\frac{1}{\tau_m} [\mathcal{D}] \mathcal{D}_S\left(\begin{smallmatrix} \mathbf{0} \\ I_{n_N} \end{smallmatrix}\right)\right) - \sum_{k=0}^{m-1} \frac{\tau_k}{\tau_m} \frac{d\|G_N\|_{H^2}^2}{d\tau_k}, \\ \frac{d\|G_N\|_{H^2}^2}{dB} &= 2\left(\begin{smallmatrix} I_n \\ \mathbf{0} \end{smallmatrix}\right)^T Q_U^T \tilde{B}, \quad \text{and} \quad \frac{d\|G_N\|_{H^2}^2}{dC} = 2\tilde{C} P_V^T[\varepsilon_0]^T, \end{aligned}$$

where $Q_U = Q(U^\perp - \mathcal{A}_{12}\mathcal{A}_{22}^{-1}U)$, $P_V = (V^\perp - V\mathcal{A}_{22}^{-1}\mathcal{A}_{21})P$, $W_A = V\mathcal{A}_{22}^{-1}U$, and $\mathcal{D}_S = P_V \mathcal{E}_{11}^T Q_U - W_A \mathcal{B}_N \tilde{B}^T Q_U - P_V \tilde{C}^T \tilde{C}_N W_A$.

INTERMEDIATE STEPS TO DERIVE THE GRADIENT

This section summarizes the steps that were elided in the previous section and mainly consists of tedious derivations. Whilst illustrative of how to work with matrix differentials, one may safely move on to the next section on page 106.

To ease legibility, let us write $\gamma^2 = \|G_N\|_{H^2}^2$. Starting from the relations

$$d\gamma^2 = \text{tr}\left(Q \frac{\partial \ell_P}{\partial X} dX\right) \quad \text{and} \quad d\gamma^2 = \text{tr}\left(P \frac{\partial \ell_Q}{\partial X} dX\right),$$

obtained in the previous section, we straightforwardly get the expression

$$\begin{aligned} d\gamma^2 &= \text{tr}\left(Q(d\tilde{A} P \mathcal{E}_{11}^T + \mathcal{E}_{11} P d\tilde{A}^T)\right) + \text{tr}\left(Q(d\tilde{B} \tilde{B}^T + \tilde{B} d\tilde{B}^T)\right) \\ &\quad + \text{tr}\left(P(d\tilde{C}^T \tilde{C} + \tilde{C}^T d\tilde{C})\right) \\ &= 2 \text{tr}\left(P \mathcal{E}_{11}^T Q d\tilde{A} + Q d\tilde{B} \tilde{B}^T + P \tilde{C}^T d\tilde{C}\right). \end{aligned}$$

We continue by writing the differentials in terms of the original system, effectively retracing the steps of section 5.2. From the application of the

Schur complement formula we have

$$\begin{aligned}\tilde{A} &= dA_{11} - dA_{12} \mathcal{A}_{22}^{-1} A_{21} + A_{12} \mathcal{A}_{22}^{-1} dA_{22} \mathcal{A}_{22}^{-1} A_{21} - A_{12} \mathcal{A}_{22}^{-1} dA_{21}, \\ \tilde{B} &= dB_1 - dA_{12} \mathcal{A}_{22}^{-1} B_2 + A_{12} \mathcal{A}_{22}^{-1} dA_{22} \mathcal{A}_{22}^{-1} B_2 - A_{12} \mathcal{A}_{22}^{-1} dB_2, \quad \text{and} \\ \tilde{C} &= dC_1 - dC_2 \mathcal{A}_{22}^{-1} A_{21} + C_2 \mathcal{A}_{22}^{-1} dA_{22} \mathcal{A}_{22}^{-1} A_{21} - C_2 \mathcal{A}_{22}^{-1} dA_{21}.\end{aligned}$$

Differentiating through the construction of the standard form gives

$$\begin{aligned}\begin{pmatrix} dA_{11} & dA_{12} \\ dA_{21} & dA_{22} \end{pmatrix} &= \begin{pmatrix} U^\perp \\ U \end{pmatrix} dA_N \begin{pmatrix} V^\perp & V \end{pmatrix}, \quad \begin{pmatrix} dB_1 \\ dB_2 \end{pmatrix} = \begin{pmatrix} U^\perp \\ U \end{pmatrix} d\mathcal{B}_N, \quad \text{and} \\ (dC_1 & dC_2) &= d\mathcal{C}_N \begin{pmatrix} V^\perp & V \end{pmatrix}.\end{aligned}$$

Finally, the discretization (5.4) brings us back to the system's parameters

$$\begin{aligned}dA_N &= \begin{pmatrix} I_n \\ \mathbf{0} \end{pmatrix} \sum_{k=0}^m (dA_k [\varepsilon_{-\tau_k}] + A_k d[\varepsilon_{-\tau_k}]) + \begin{pmatrix} \mathbf{0} \\ I_{nN} \end{pmatrix} d[\mathcal{D}], \\ d\mathcal{B}_N &= \begin{pmatrix} I_n \\ \mathbf{0} \end{pmatrix} dB, \quad \text{and} \quad d\mathcal{C}_N = dC [\varepsilon_0],\end{aligned}$$

where $d[\varepsilon_{-\tau_k}] = -[\varepsilon_{-\tau_k} \mathcal{D}] \left(d\tau_k - \frac{\tau_k}{\tau_m} d\tau_m \right)$ and $d[\mathcal{D}] = -\frac{1}{\tau_m} [\mathcal{D}] d\tau_m$.

From this we can then derive $d\gamma^2$ as a function of dA_k , $d\tau_k$ for $k < m$, $d\tau_m$, dB , and dC , respectively

$$\begin{aligned}d\gamma^2 &= 2 \operatorname{tr} \left([\varepsilon_{-\tau_k}] \left((V^\perp - V \mathcal{A}_{22}^{-1} A_{21}) P \mathcal{E}_{11}^T Q (U^\perp - A_{12} \mathcal{A}_{22}^{-1} U) \right. \right. \\ &\quad \left. \left. - V \mathcal{A}_{22}^{-1} B_2 \tilde{B}^T Q (U^\perp - A_{12} \mathcal{A}_{22}^{-1} U) \right. \right. \\ &\quad \left. \left. - (V^\perp - V \mathcal{A}_{22}^{-1} A_{21}) P \tilde{C}^T C_2 \mathcal{A}_{22}^{-1} U \right) \begin{pmatrix} I_n \\ \mathbf{0} \end{pmatrix} dA_k \right), \\ d\gamma^2 &= -2 \operatorname{tr} \left(\left((V^\perp - V \mathcal{A}_{22}^{-1} A_{21}) P \mathcal{E}_{11}^T Q (U^\perp - A_{12} \mathcal{A}_{22}^{-1} U) \right. \right. \\ &\quad \left. \left. - V \mathcal{A}_{22}^{-1} B_2 \tilde{B}^T Q (U^\perp - A_{12} \mathcal{A}_{22}^{-1} U) \right. \right. \\ &\quad \left. \left. - (V^\perp - V \mathcal{A}_{22}^{-1} A_{21}) P \tilde{C}^T C_2 \mathcal{A}_{22}^{-1} U \right) \begin{pmatrix} I_n \\ \mathbf{0} \end{pmatrix} A_k [\varepsilon_{-\tau_k} \mathcal{D}] \right) d\tau_k, \\ &\quad \text{for } k < m, \\ d\gamma^2 &= -2 \operatorname{tr} \left(\left((V^\perp - V \mathcal{A}_{22}^{-1} A_{21}) P \mathcal{E}_{11}^T Q (U^\perp - A_{12} \mathcal{A}_{22}^{-1} U) \right. \right. \\ &\quad \left. \left. - V \mathcal{A}_{22}^{-1} B_2 \tilde{B}^T Q (U^\perp - A_{12} \mathcal{A}_{22}^{-1} U) \right. \right. \\ &\quad \left. \left. - (V^\perp - V \mathcal{A}_{22}^{-1} A_{21}) P \tilde{C}^T C_2 \mathcal{A}_{22}^{-1} U \right) \begin{pmatrix} \mathbf{0} \\ I_{nN} \end{pmatrix} \frac{1}{\tau_m} [\mathcal{D}] \right) d\tau_m \\ &\quad - \sum_{k=0}^{m-1} \frac{\tau_k}{\tau_m} \frac{d\gamma^2}{d\tau_k} d\tau_m, \\ d\gamma^2 &= 2 \operatorname{tr} \left(\tilde{B}^T Q (U^\perp - A_{12} \mathcal{A}_{22}^{-1} U) \begin{pmatrix} I_n \\ \mathbf{0} \end{pmatrix} dB \right), \quad \text{and} \\ d\gamma^2 &= 2 \operatorname{tr} \left([\varepsilon_0] (V^\perp - V \mathcal{A}_{22}^{-1} A_{21}) P \tilde{C}^T dC \right).\end{aligned}$$

Using the abbreviations

$$\begin{aligned}Q_U &= Q(U^\perp - A_{12} \mathcal{A}_{22}^{-1} U), \quad P_V = (V^\perp - V \mathcal{A}_{22}^{-1} A_{21}) P, \\ W_A &= V \mathcal{A}_{22}^{-1} U, \quad \text{and} \quad \mathcal{D}_S = P_V \mathcal{E}_{11}^T Q_U - W_A \mathcal{B}_N \tilde{B}^T Q_U - P_V \tilde{C}^T C_N W_A,\end{aligned}$$

we get the compacter

$$\begin{aligned}
 d\gamma^2 &= 2 \operatorname{tr}([\varepsilon_{-\tau_k}] \mathcal{D}_S \begin{pmatrix} I_n \\ \mathbf{0} \end{pmatrix} dA_k), \\
 d\gamma^2 &= -2 \operatorname{tr}(\mathcal{D}_S \begin{pmatrix} I_n \\ \mathbf{0} \end{pmatrix} A_k [\varepsilon_{-\tau_k} \mathcal{D}]) d\tau_k, \quad k < m, \\
 d\gamma^2 &= -2 \operatorname{tr}(\mathcal{D}_S \begin{pmatrix} \mathbf{0} \\ I_{nN} \end{pmatrix} \frac{1}{\tau_m} [\mathcal{D}]) d\tau_m - \sum_{k=0}^{m-1} \frac{\tau_k}{\tau_m} \frac{d\gamma^2}{d\tau_k} d\tau_k, \\
 d\gamma^2 &= 2 \operatorname{tr}(\tilde{B}^T Q_U \begin{pmatrix} I_n \\ \mathbf{0} \end{pmatrix} dB), \quad \text{and} \quad d\gamma^2 = 2 \operatorname{tr}([\varepsilon_0] P_V \tilde{C}^T dC).
 \end{aligned}$$

From the fact that $df = \operatorname{tr}(Y^T dX)$ implies $\frac{df}{dX} = Y$, we find the final results of the previous section.

7.2 NUMERICAL EXAMPLES

We conclude this section by illustrating how the approach we developed so far can be used to synthesize robust controllers and H^2 -optimal reduced models. As optimization strategy we use the Broyden–Fletcher–Goldfarb–Shanno quasi-Newton method (see e.g. Nocedal and Wright 2006: section 6.1) with the line search algorithm by Hager and Zhang (2006). We compute the gradient using the formulas of section 7.1. For a step that would invalidate the conditions for the existence of the gradient outlined in section 7.1, we take the cost function to be infinite such that the line search algorithm will naturally guarantee those conditions. Throughout the past chapters we have seen that the choice $\phi_N = P_N$, the appropriately shifted and scaled Legendre polynomial of degree N , for the discretization (5.4) possesses the most favourable results. We shall thus use this basis for all examples. Finally, we take $N = 40$.

We first verify our method by reproducing examples 1–3 from Gomez, Egorov, et al. (2019), in order.

Example 7.2.1. Consider the following closed-loop system from Vanbiervliet, Verheyden, et al. (2008):

$$\begin{aligned}
 \dot{\mathbf{x}}(t) &= \begin{pmatrix} -0.08 & -0.03 & 0.2 \\ 0.2 & -0.04 & -0.005 \\ -0.06 & -0.2 & -0.07 \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} -0.1 \\ -0.2 \\ 0.1 \end{pmatrix} (p_1 \ p_2 \ p_3) \mathbf{x}(t-5) + \mathbf{u}(t), \\
 \mathbf{y}(t) &= \mathbf{x}(t).
 \end{aligned}$$

Introducing the slack variable x_4 , which corresponds to the control signal in the open-loop system, allows us to keep the parameters as

separate matrix elements, yielding the DDAE formulation

$$\begin{aligned} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \dot{\mathbf{x}}(t) &= \begin{pmatrix} -0.08 & -0.03 & 0.2 & 0 \\ 0.2 & -0.04 & -0.005 & 0 \\ -0.06 & -0.2 & -0.07 & 0 \\ p_1 & p_2 & p_3 & -1 \end{pmatrix} \mathbf{x}(t) \\ &+ \begin{pmatrix} 0 & 0 & 0 & -0.1 \\ 0 & 0 & 0 & -0.2 \\ 0 & 0 & 0 & 0.1 \\ 0 & 0 & 0 & 0 \end{pmatrix} \mathbf{x}(t-5) + \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix} \mathbf{u}(t), \\ \mathbf{y}(t) &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \mathbf{x}(t). \end{aligned}$$

Starting from the stabilizing parameters $p_1 = 0.472$, $p_2 = 0.505$, and $p_3 = 0.603$, which yield a H^2 -norm of approximately 8.91, we find $p_1^* \approx 0.538$, $p_2^* \approx 0.338$, and $p_3^* \approx 0.226$ giving a much lower H^2 -norm of about 5.70. All identical to the values reported by Gomez, Egorov, et al. (2019).

Example 7.2.2. The next example concerns the design of a proportional-retarded controller for a DC servomechanism (Villafuerte et al. 2013):

$$\begin{aligned} \dot{\mathbf{x}}(t) &= \begin{pmatrix} 0 & 1 \\ -\nu^2 & -bk_p - 2\delta\nu \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} 0 & 0 \\ bk_r & 0 \end{pmatrix} \mathbf{x}(t-\tau) + \begin{pmatrix} 0 \\ b \end{pmatrix} u(t), \\ y(t) &= (1 \ 0) \mathbf{x}(t), \end{aligned}$$

where $\nu = 17.6$, $\delta = 0.0128$, and $b = 31$, leaving τ , k_p , and k_r as control parameters.

Like the previous example, we can keep the parameters to optimize as separate matrix elements by introducing the control signal as the slack variable x_3 , giving

$$\begin{aligned} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \dot{\mathbf{x}}(t) &= \begin{pmatrix} 0 & 1 & 0 \\ -\nu^2 & -2\delta\nu & b \\ -k_p & 0 & -1 \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ k_r & 0 & 0 \end{pmatrix} \mathbf{x}(t-\tau) + \begin{pmatrix} 0 \\ b \\ 0 \end{pmatrix} u(t), \\ y(t) &= (1 \ 0 \ 0) \mathbf{x}(t). \end{aligned}$$

As in Gomez, Egorov, et al. (2019), optimizing for all three parameters yields no finite optimizer as we see $\tau \rightarrow 0$, $k_p \rightarrow -\infty$, and $k_r \rightarrow \infty$. Fixing $k_p = 22.57$ as suggested, and starting optimization from $\tau = 0.03$ and $k_r = 3$, yields $\tau^* \approx 0.0519$ and $k_r^* \approx 17.964$, improving the H^2 -norm from about 0.465 to 0.223, values identical to those obtained by Gomez, Egorov, et al. (2019).

Note that, since we can treat delay optimization directly, the formulation of the problem is significantly easier for our method as we did not need a change of variables to optimize for τ .

Example 7.2.3. Consider the following exponentially stable linearized model of a refinement plant by Ross (1971):

$$\begin{aligned} \dot{\mathbf{x}}(t) &= A_0 \mathbf{x}(t) + A_1 \mathbf{x}(t-0.1) + B \mathbf{u}(t), \\ y(t) &= C \mathbf{x}(t), \end{aligned}$$

where

$$A_0 = \begin{pmatrix} -4.93 & -1.01 & 0 & 0 \\ -3.20 & -5.30 & -12.8 & 0 \\ 6.40 & 0.347 & -32.5 & -1.04 \\ 0 & 0.833 & 11 & -3.96 \end{pmatrix}, \quad A_1 = \begin{pmatrix} 1.92 & 0 & 0 & 0 \\ 0 & 1.92 & 0 & 0 \\ 0 & 0 & 1.87 & 0 \\ 0 & 0 & 0 & 0.724 \end{pmatrix},$$

$$B = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}, \quad \text{and} \quad C = (1 \ 1 \ 1 \ 1).$$

We wish to find a stable approximate model of dimension two,

$$\begin{aligned} \dot{\mathbf{x}}_r(t) &= A_{0r}\mathbf{x}_r(t) + A_{1r}\mathbf{x}_r(t-0.1) + B_r\mathbf{u}(t), \\ y_r(t) &= C_r\mathbf{x}_r(t), \end{aligned}$$

with minimal L^2 -error on the transfer function along the imaginary axis. We can do so by minimizing the H^2 -norm of the following system, optimizing for A_{0r} , A_{1r} , B_r , and C_r ,

$$\begin{aligned} \dot{\mathbf{x}}_e(t) &= \begin{pmatrix} A_0 & \mathbf{0} \\ \mathbf{0} & A_{0r} \end{pmatrix} \mathbf{x}_e(t) + \begin{pmatrix} A_1 & \mathbf{0} \\ \mathbf{0} & A_{1r} \end{pmatrix} \mathbf{x}_e(t-0.1) + \begin{pmatrix} B \\ B_r \end{pmatrix} \mathbf{u}(t), \\ y_e(t) &= (C \ -C_r) \mathbf{x}_e(t). \end{aligned}$$

Starting from the stabilizing matrices

$$\begin{aligned} A_{0r} &= \begin{pmatrix} -3 & -1 \\ -3 & -2 \end{pmatrix}, \quad A_{1r} = \begin{pmatrix} 1 & 0 \\ 2 & 0 \end{pmatrix}, \\ B_r &= \begin{pmatrix} 1.6 & 0.3 \\ 0.15 & 0.7 \end{pmatrix}, \quad \text{and} \quad C_r = (0.7 \ -0.7), \end{aligned}$$

the H^2 -norm after optimization is approximately $5.91 \cdot 10^{-3}$ which is near identical to the error on the model obtained in Gomez, Egorov, et al. (2019).

Next, we look at two neutral systems of increasing complexity.

Example 7.2.4. We consider the following example, inspired by example 6.2.2 of Byrnes et al. (1984):

$$\begin{aligned} \dot{x}(t) &= -x(t) + x(t-1) + p_1\dot{x}(t-1) + p_2x(t-1) + u(t), \\ y(t) &= x(t). \end{aligned}$$

For $p_1 = 0$ and $p_2 = -1$ the authors show this system to be stable.

To arrive at a system of the form of (5.1), we add the algebraic equation $x_2(t) = \dot{x}_1(t)$ to capture the neutral terms. The control signal is added as x_3 . This then gives us

$$\begin{aligned} \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix} \dot{\mathbf{x}}(t) &= \begin{pmatrix} -1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} 1 & 0 & 0 \\ 0 & p_2 & 0 \\ 0 & p_1 & 0 \end{pmatrix} \mathbf{x}(t-1) + \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} u(t), \\ y(t) &= (1 \ 0 \ 0) \mathbf{x}(t). \end{aligned}$$

For the aforementioned choice of parameters, the H^2 -norm is approximately 0.71. Using our method, we find an alternative controller with $p_1^* \approx -0.27$ and $p_2^* \approx -1.50$, yielding a H^2 -norm of about 0.66.

Example 7.2.5. A more interesting system—which also illustrates how one can use slack variables in a DDAE formulation to obtain higher order derivatives—stems from a linear oscillator under delayed acceleration-derivative feedback with additive noise on the measured acceleration and derivative:

$$\begin{aligned}\ddot{x}(t) + 2\xi\dot{x}(t) + x(t) &= p_1(\ddot{x}(t - \tau_1) + u_1(t)) + p_2(\dot{x}(t - \tau_2) + u_2(t)), \\ y(t) &= x(t).\end{aligned}$$

We can obtain the delayed derivative and acceleration from the algebraic equations $x_2(t) = \dot{x}_1(t)$ and $x_3(t) = \dot{x}_2(t) = \ddot{x}_1(t)$. The slack variables $x_4(t) = \ddot{x}_1(t - \tau_1) + u_1(t)$ and $x_5(t) = \dot{x}_1(t - \tau_2) + u_2(t)$ give us the measurements. Together this gives the DDAE formulation

$$\begin{aligned}\begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \dot{\mathbf{x}}(t) &= \begin{pmatrix} -1 & -2\xi & 0 & p_1 & p_2 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & -1 \end{pmatrix} \mathbf{x}(t) + \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix} \mathbf{x}(t - \tau_1) \\ &+ \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{pmatrix} \mathbf{x}(t - \tau_2) + \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} \mathbf{u}(t), \\ y(t) &= (1 \ 0 \ 0 \ 0 \ 0) \mathbf{x}(t).\end{aligned}$$

From the stability chart on figure 1 of Wang et al. (2017) we see that for $\xi = 0.2$, $p_1 = 0.5$, $p_2 = -20$, $\tau_1 = 0.2$, and $\tau_2 = 0.1$ the zero solution is exponentially stable, as $|p_1| < 1$ it is also strongly stable. We find a H^2 -norm of about 3.23 for this setup. Optimizing for p_2 we can improve this to about 0.57 by taking $p_2^* \approx -0.33$. Additionally optimizing for τ_1 , lower bounding it to $\tau_2 = 0.1$, yields $p_2^{**} \approx -0.28$ and $\tau_1^{**} \approx 0.1$ with a slightly improved H^2 -norm of approximately 0.53.

7.3 SUMMARY

In section 7.1 we derived analytical expressions for the gradient of the square of the H^2 -norm of the discretized system, with respect to the system matrices and the delays. These expressions, together with the H^2 -norm, can be evaluated for slightly more than double the computational cost of computing the H^2 -norm alone. This allows us to efficiently synthesize both H^2 -optimal fixed-order controllers and stable approximate models. We illustrated the effectiveness of this approach in section 7.2 by reproducing three known examples for retarded type system and two novel examples involving neutral terms. Whilst we did not do so here, this method is readily combined with existing techniques for H^∞ -norm optimization (see e.g. Appeltans 2024), allowing for mixed optimality criteria, as are often used in practice.

8 *Concluding remarks*

In addition to the summaries we provided at the ends of the past chapters, we want to use this final one to highlight some common threads throughout the preceding text and mention some questions left unanswered.

Our first main development was of the Lanczos tau method, of which we introduced an operator pencil formulation in chapter 3, which simplifies much of the later discussion, compared to trying to use a coefficient expression directly. Throughout the text we proved relations to rational approximation of the exponential in the transfer function (Propositions 3.4.1 and 5.3.1 and section 6.2). We could use these to show that the transfer function of the discretized system converges to that of the original system (Theorems 3.4.7 and 5.3.3) and that the Lanczos tau method is equivalent to the popular pseudospectral collocation method (Theorem 3.5.2). The latter correspondence allows us to reuse existing results on the super-geometric convergence of characteristic roots from Theorem 5.1 of Breda, Maset, et al. (2015).

When using the Legendre basis, we could prove preservation of stability for discrete delays (Theorems 3.8.1, 5.3.6 and 6.3.1), thanks to a connection to Padé approximation (Theorem 3.7.1 and Proposition 3.7.4). Whilst eventual preservation of stability is a useful property to have, in practice we often observe that a discretized system's characteristic roots approach the true ones from the left, i.e. is always stabler. It would be interesting to know if, and under which conditions, this is true. Additionally, our theorems only hold for discrete delays. As similar results exist for the distributed case (see Theorem 5.3 of Ito (1985) and section 7 of Ito and Kappel (1992)), it would be valuable to see if their arguments can be used in our setting as well. Furthermore, whilst we only have arguments for the Legendre basis, other bases also appear to preserve stability (though not all, see section 3.8) and similar developments could thus be possible there.

Finally, in the case of DDAEs of index one, we were able to show, under assumption of strong exponential stability, that discretizing the DDAE using a Lanczos tau method results in a DAE of which the algebraic constraints can be eliminated (Proposition 5.2.3) and, if the strong H^2 -norm is finite, that this discretized system does not have feedthrough

(Theorem 5.2.4). These results then allowed us to use such a discretization to approximate the H^2 -norm of these DDAEs in section 5.2.

This naturally leads us to our second aim, which was to further analyse and extend upon this method, which was initially introduced for RDDEs by Vanbiervliet, Michiels, et al. (2011). On the theoretical side we could prove for the first time, under reasonable assumptions, that the H^2 -norm of a Lanczos tau discretization converges to that of the original system (Theorems 4.1.3 and 5.3.5 and section 6.3). Whilst the current set of assumptions appears mostly sharp, some milder assumptions might still be possible, in particular Assumption 5.3.4 seems overly strict. As we can prove stability when using the Legendre basis for discrete delays, and most other assumptions are trivially satisfied for this choice, we can often guarantee these assumptions, and thus convergence, when using this basis (Corollaries 4.1.4, 5.3.7 and 6.3.2).

We saw empirically that symmetric bases, and in particular the Legendre basis (see fig. 4.3), yield better convergence rates for the H^2 -norm. We saw geometric convergence for both single delay RDDEs and NDDEs (figs. 4.1 and 5.2) and distributed delay (fig. 4.7). Using splines, we could improve both the third order and first order algebraic convergence of RDDEs and NDDEs, respectively, by a little over a full order when taking the delays as knots (figs. 6.3 and 6.4). By additionally placing knots at integer combinations of the delays, we could improve the convergence rate further, up to geometric convergence for commensurate delays (figs. 6.6 and 6.7 and section 6.4).

We seem able to predict the convergence rate through a link with the delay-Lyapunov equation (Proposition 4.4.1). The H^2 -norm depends on $\lambda(0)$, the value of its solution in zero, and we observe one order faster convergence than the corresponding spectral method would converge to λ (sections 4.6 and 6.4). The one exception are NDDEs when discretized using a polynomial instead of a spline, which display one order lower convergence that we would expect from this pattern (section 5.3). We assume that the mechanism that prevents the additional order from being attained disappears for a spline discretization.

However, this is no proof of these rates. Furthermore, except for the link to the delay-Lyapunov equation, the modulus on the imaginary axis of the rational approximation of the exponential (Proposition 4.1.2 and eq. (6.4)) likely also plays a role, as this is closely related to the symmetry condition we needed on the basis.

Comparatively, we have said little about which discretization degree N results in a what error. As the convergence rates are generally fast and monotone enough to test accuracy by comparing two different discretization degrees, an iterative procedure is possible, however this can quickly become expensive. Ideally, we would have a reasonably

cheap heuristic, like Wu and Michiels (2012) provide for the rightmost characteristic roots.

Both the extensive theoretical results available and the superior convergence strongly point at the Lanczos tau method based on the Legendre basis as preferred choice when approximating the H^2 -norm.

Furthermore, there is one possible additional advantage that we have not investigated in this text. Whilst we saw that a Lanczos tau method using a Legendre basis admits a highly sparse matrix representation (see section 3.6), which is readily adapted to Arnoldi iteration in the style of Jarlebring, Meerbergen, et al. (2010) to approximate the characteristic roots, we have yet to exploit this for the H^2 -norm of large scale systems. One possible starting point is the projection based method of Michiels and B. Zhou (2019) to approximate the delay-Lyapunov matrix of an RDDE.

In the penultimate chapter, we concluded our discussion by deriving the gradient of the square of the approximate H^2 -norm with respect to system matrices and delays (section 7.1). This allowed us to move from analysis to synthesis, as it gives us an effective method to find H^2 -optimal controllers and reduced order models as illustrated in section 7.2. Whilst we did not do this, the resulting method is readily combined with existing methods for H^∞ -norm optimization (see e.g. Appeltans 2024), allowing for the mixed optimality criteria often used in practice.

Finally, it is important to contextualize these results somewhat. To reiterate the epigraph of this thesis: “Essentially, all models are wrong, but some are useful.”—a remark generally attributed to George E. P. Box; here as it appears in Box and Draper (1987: p. 424). In the book, this remark excuses the use of a polynomial to locally approximate a statistical model. The argument goes that, as the statistical model itself is bound to be at least a bit wrong compared to reality, the additional error due to the polynomial approximation need not be detrimental to the intended application. If anything, it makes the application feasible. In an analogous fashion, we used polynomials to approximate a—now dynamical—model, such that we can obtain approximate values for the characteristic roots and H^2 -norm. Whilst we saw great convergence for both, this fast decay of relative numerical error should itself be interpreted relative to the error of the model, i.e. the error with respect to the real world. Remember, a delay is generally itself already a—very useful—simplification of reality.

Acknowledgements

COLLABORATORS

The articles on which most chapters of this thesis are based (i.e. Provoost and Michiels 2024; 2025; 2026) were coauthored by Wim Michiels, who provided valuable scientific and authorial input.

The author wishes to thank Nick Dewaele for providing the proof of Lemma 3.4.6 and Jan Sieber for suggesting that placing knots at combinations of delays might increase convergence further in chapter 6.

FUNDING

This work was supported by KU Leuven project C14/22/092 and by FWO-Flanders under grant number G.092.721N.

USED TOOLS

Given the recent rise in the abilities of generative artificial intelligence (AI), the Arenberg Doctoral School (to which this thesis is presented) stipulates that one should acknowledge the ways in which one uses such tools. For this thesis, the following of their standardized statements applies:

I did not use generative AI assistance tools during the research/writing process of my thesis.

The text, code, and images in this thesis are my own (unless otherwise specified) and generative AI has only been used in accordance with the KU Leuven guidelines and appropriate references have been added. I have reviewed and edited the content as needed and I take full responsibility for the content of the thesis.

This work represents a significant part of the last four years of my life. I would find it unsettling to allow anyone or anything but myself to convey those results. This thesis is mine, the words chosen to convey its message—including possible idiosyncrasies or even errors—should hence also be mine.

However, I would like to use this opportunity to acknowledge a number of tools that were invaluable to the completion of this thesis, yet generally go unmentioned. Of course, such a list is always bound to be incomplete, but it should at the very least include, in no particular order: my text editor of choice Emacs with the Aspell and AUCT_EX packages, all maintained by the GNU project;¹ the Julia programming language,² which was used to produce all the graphs and experiments presented in this thesis; the Maxima computer algebra system,³ which verified some of the more tedious mathematical derivations; and the Linux kernel,⁴ which, together with the GNU core utilities,¹ provides the required environment for the aforementioned tools to run. This thesis itself was typeset by Lua_T_EX and laid out using the macros provided by the L^AT_EX project and the wonderful memoir document class, with a number of additional figures prepared using the Asymptote vector graphics language.⁵ Finally, where automated tools fell short, I could rely upon *New Hart's Rules*, edited by A. Waddingham (2nd edn, Oxford, 2014), and the *Concise Oxford English Dictionary*, edited by A. Stevenson and M. Waite (11th edn, Oxford, 2011), for linguistic guidance.

¹Source code of these tools and many more available at <https://www.gnu.org/>.

²Documentation and links to the source can be found at <https://julialang.org/>.

³Source code available via <https://maxima.sourceforge.io/>.

⁴Further information and source available from <https://www.kernel.org/>.

⁵Source code and documentation of all four are available from the Comprehensive T_EX Archive Network at <https://ctan.org/>.

Code availability

The source code needed to reproduce the experiments and figures in this text is available at <https://doi.org/10.5281/zenodo.22940620>.

This deposit includes implementations of the presented methods that closely mimic the notation of the text, all used scripts, a machine-readable definition of the required dependencies and their versions, and instructions on how to generate the figures and numerical results.

Curriculum vitae

EDUCATION

2022–now Doctor of Engineering Science, Computer Science

Thesis: *Lanczos tau methods for time-delay systems: Approximation of system and its H^2 -norm*

2020–’22 Master of Science, Mathematical Engineering

Thesis: *Fast and accurate computation of eigenvalue problems for the numerical spectroscopy of fluids and plasmas*

2017–’20 Bachelor of Science, Computer Science

All at KU Leuven.

PUBLICATIONS AND PREPRINTS

- Provoost, E., and Michiels, W. (2026), “Computing and optimizing the H^2 -norm of delay differential algebraic systems”, arXiv: 2603.10734 [math.NA]
- Provoost, E., and Michiels, W. (2025), “Accelerated H^2 -norm approximation for time-delay systems with discrete delays”, *IFAC-PapersOnLine*, 59/13, Presented at the *19th IFAC Workshop on Time-Delay Systems*: 93–8
- Provoost, E., and Michiels, W. (2024), “The Lanczos tau framework for time-delay systems: Padé approximation and collocation revisited”, *SIAM Journal on Numerical Analysis*, 62/6: 2529–48

TALKS AND SEMINARS

- Contributed talk, *15th AIMS Conference*, Athens, Greece (July 2026)
- Invited talk, *BeNeLux Mathematical Congress 2026*, Antwerp, Belgium (April 2026)

- Internal seminar, *NUMA Seminar*, KU Leuven, Leuven, Belgium (November 2025)
- Invited talk, *19th IFAC Workshop on Time-Delay Systems*, Gif-sur-Yvette, France (July 2025)
- Internal seminar, *NUMA Seminar*, KU Leuven, Leuven, Belgium (December 2024)
- Invited talk, *18th IFAC Workshop on Time-Delay Systems*, Udine, Italy (September 2024)
- Poster, *6e Journées Approximation*, Lille, France (May 2024)
- Internal seminar, *NUMA Seminar*, KU Leuven, Leuven, Belgium (March 2024)
- Contributed talk, *42nd Benelux Meeting on Systems and Control*, Elspeet, The Netherlands (March 2023)

RESEARCH VISITS

August 2025 Visiting researcher, School of Mathematics, University of Leeds. To work on numerical methods for the basins of attraction of nonlinear delay differential equations. Hosted by Dr Francesca Scarabel through a grant of the London Mathematical Society.

INSTITUTIONAL SERVICE AT KU LEUVEN

2024–'25 Member of the Academic Council and University Council

2024–'25 Member of the Executive Committee of the Group Science, Engineering and Technology

2024–'25 Member of the Commission on Research Integrity

2024–'25 Member of the Library Council

2020–'26 Member of the Permanent Education Commission Mathematical Engineering

2020–'22 Member of the Departmental Council Computer Science

2019–'20 Observing Member of the Permanent Education Commission Computer Science

TEACHING AT KU LEUVEN

2022–'25 Co-organizer of a half-yearly internal presentation skills workshop

2023 Guest minilecture on ‘Visual understanding of complex functions’ for *Complex Analysis with Applications*

2022 Co-organizer of an internal reading course on *Functional Analysis*

Mentoring and teaching assistance:

2025 Teaching assistant for *Numerical Analysis*

2024–'25 Mentor of a master thesis on *Automatic accuracy and complexity optimization of matrix expressions*

2022 & '23 Teaching assistant for *Problem solving and design, part 3*

2023–'24 Mentor of a master thesis on *Rational approximation methods for evaluating and approximating matrix functions*

Bibliography

- Ahmad, F. (1998), “The orthogonal polynomials and the Padé approximation”, *Applied Mathematics and Mechanics*, 19/7: 663–8.
- Ahmed, N., and Fisher, P. S. (1970), “Study of algorithmic properties of Chebyshev coefficients”, *International Journal of Computer Mathematics*, 2: 307–17.
- Álvarez de Morales, M., Pérez, T. E., and Piñar, M. A. (1998), “Sobolev orthogonality for the Gegenbauer polynomials $\{C_n^{(-N+1/2)}\}_{n \geq 0}$ ”, *Journal of Computational and Applied Mathematics*, 100: 111–20.
- Appeltans, P. (2024), “Robust Control of Time-Delay Systems: Analysis, Computational Methods and Software”, PhD thesis (KU Leuven).
- Beattie, C., and Gugercin, S. (2012), “Realization-independent $\mathcal{H}_2$ -approximation”, in *Proceedings of the 51st IEEE Conference on Decision and Control*, 4953–8.
- Bellen, A., and Zennaro, M. (2003), *Numerical Methods for Delay Differential Equations* (Oxford).
- Bellman, R., and Cooke, K. L. (1963), *Differential-Difference Equations* (New York).
- Bellman, R., and Danskin, J. M. (1954), *A Survey of the Mathematical Theory of Time-Lag, Retarded Control, and Hereditary Processes* (Santa Monica: The RAND Corporation).
- Bernier, C., and Manitius, A. (1978), “On semigroups in $R^n \times L^p$ corresponding to differential equations with delays”, *Canadian Journal of Mathematics*, 30/5: 897–914.
- Bernstein, S. (1912), “Sur l’ordre de la meilleure approximation des fonctions continues par des polynômes de degré donné”, *Académie Royale de Belgique, Mémoires de la Classe des Sciences, Collection in-4°*, 2nd ser., 4: 1–103.
- Birkhoff, G., and Varga, R. S. (1965), “Discretization errors for well-set Cauchy problems I”, *Journal of Mathematics and Physics*, 44/1: 1–23.

- Box, G. E. P., and Draper, N. R. (1987), *Empirical Model-Building and Response Surfaces* (New York).
- Boyd, J. P. (2001), *Chebyshev and Fourier Spectral Methods* (rev. 2nd edn., Mineola).
- Breda, D., Diekmann, O., Gyllenberg, M., Scarabel, F., and Vermiglio, R. (2016), “Pseudospectral discretization of nonlinear delay equations: New prospects for numerical bifurcation analysis”, *SIAM Journal on Applied Dynamical Systems*, 15/1: 1–23.
- Breda, D., Maset, S., and Vermiglio, R. (2005), “Pseudospectral differencing methods for characteristic roots of delay differential equations”, *SIAM Journal on Scientific Computing*, 27/2: 482–95.
- (2015), *Stability of Linear Delay Differential Equations: A Numerical Approach with MATLAB* (New York).
- Brockett, R. W. (1970), *Finite Dimensional Linear Systems* (New York).
- Butcher, E. A., Ma, H., Bueler, E., Averina, V., and Szabo, Z. (2004), “Stability of linear time-periodic delay-differential equations via Chebyshev polynomials”, *International Journal for Numerical Methods in Engineering*, 59/7: 895–922.
- Byrnes, C. I., Spong, M. W., and Tarn, T. (1984), “A several complex variables approach to feedback stabilization of linear neutral delay-differential systems”, *Mathematical Systems Theory*, 17/1: 97–133.
- Conni, G., Naets, F., and Meerbergen, K. (2024), “AAA rational approximation for time domain model order reduction”, *Applied Mathematics Letters*, 157: 109188.
- Curtain, R. F., and Zwart, H. (1995), *An Introduction to Infinite-Dimensional Linear Systems Theory* (New York).
- Davis, P. J. (1963), *Interpolation and Approximation* (New York).
- Diekmann, O., Verduyn Lunel, S. M., Gils, S. A., and Walther, H.-O. (1995), *Delay Equations: Functional-, Complex-, and Nonlinear Analysis* (New York).
- Engelborghs, K., Luzyanina, T., and Roose, D. (2002), “Numerical bifurcation analysis of delay differential equations using DDE-BIFTOOL”, *ACM Transactions on Mathematical Software*, 28/1: 1–21.
- Erneux, T. (2009), *Applied Delay Differential Equations* (New York).
- Farris, F. A. (1998), review of T. Needham (1997), *Visual Complex Analysis* (Oxford), in *The American Mathematical Monthly*, 105/6: 570–76.

- Foias, C., Özbay, H., and Tannenbaum, A. (1996), *Robust Control of Infinite Dimensional Systems: Frequency Domain Methods* (London).
- Fridman, E. (2002), “Stability of linear descriptor systems with delay: A Lyapunov-based approach”, *Journal of Mathematical Analysis and Applications*, 273/1: 24–44.
- Glass, D. S., Jin, X., and Riedel-Kruse, I. H. (2021), “Nonlinear delay differential equations and their application to modeling biological network motifs”, *Nature Communications*, 12: 1788.
- Glover, K., Lam, J., and Partington, J. R. (1991), “Rational approximation of a class of infinite-dimensional systems II: optimal convergence rates of L_∞ approximants”, *Mathematics of Control, Signals, and Systems*, 4: 233–46.
- Gomez, M. A., Egorov, A. V., Mondié, S., and Michiels, W. (2019), “Optimization of the $\mathcal{H}_2$ norm for single-delay systems, with application to control design and model approximation”, *IEEE Transactions on Automatic Control*, 64/2: 804–11.
- Gomez, M. A., Jungers, R. M., and Michiels, W. (2020), “On the m -dimensional Cayley–Hamilton theorem and its application to an algebraic decision problem inferred from the $\mathcal{H}_2$ norm analysis of delay systems”, *Automatica*, 113: 108761.
- (2022), “On the strong $\mathcal{H}_2$ norm of differential algebraic systems with multiple delays: Finiteness criteria, regularization, and computation”, *IEEE Transactions on Automatic Control*, 67/1: 121–33.
- Gugercin, S., Antoulas, A. C., and Beattie, C. (2008), “ $\mathcal{H}_2$ model reduction for large-scale linear dynamical systems”, *SIAM Journal on Matrix Analysis and Applications*, 30/2: 609–38.
- Gumussoy, S., and Michiels, W. (2010), “A predictor–corrector type algorithm for the pseudospectral abscissa computation of time-delay systems”, *Automatica*, 46/4: 2212–38.
- Hager, W. W., and Zhang, H. (2006), “Algorithm 851: CG_DESCENT, a conjugate gradient method with guaranteed descent”, *ACM Transactions on Mathematical Software*, 32/1: 113–37.
- Hale, J. K. (1977), *Theory of Functional Differential Equations* (New York).
- Hale, J. K., and Verduyn Lunel, S. M. (1993), *Introduction to Functional Differential Equations* (New York).
- (2002), “Strong stabilization of neutral functional differential equations”, *IMA Journal of Mathematical Control and Information*, 19/1: 5–23.

- Hammarling, S. J. (1982), “Numerical solution of the stable, non-negative definite Lyapunov equation”, *IMA Journal of Numerical Analysis*, 2/3: 303–23.
- Hayes, N. D. (1950), “Roots of the transcendental equation associated with a certain difference-differential equation”, *Journal of the London Mathematical Society*, 1st ser., 25/3: 226–32.
- IEC 61966-2-1:1999, *Multimedia systems and equipment – Colour measurement and management – Part 2-1: Colour management – Default RGB colour space - sRGB* (International Electrotechnical Commission).
- Ito, K. (1985), “Legendre–tau approximation for functional differential equations part III: Eigenvalue approximations and uniform stability”, in *Proceedings of the 2nd International Conference on Distributed Parameter Systems*, 191–212.
- Ito, K., and Kappel, F. (1992), “Two families of approximation schemes for delay systems”, *Results in Mathematics*, 21/1: 93–137.
- Ito, K., and Teglas, R. (1986), “Legendre–tau approximations for functional-differential equations”, *SIAM Journal on Control and Optimization*, 24/4: 737–59.
- (1987), “Legendre–tau approximation for functional differential equations part II: The linear quadratic optimal control problem”, *SIAM Journal on Control and Optimization*, 25/6: 1379–408.
- Jarlebring, E., Meerbergen, K., and Michiels, W. (2010), “A Krylov method for the delay eigenvalue problem”, *SIAM Journal on Scientific Computing*, 32/6: 3278–300.
- Jarlebring, E., Vanbiervliet, J., and Michiels, W. (2011), “Characterizing and computing the $\mathcal{H}_2$ norm of time-delay systems by solving the delay Lyapunov equation”, *IEEE Transactions on Automatic Control*, 56/4: 814–25.
- Keane, A., Krauskopf, B., and Postlethwaite, C. M. (2017), “Climate models with delay differential equations”, *Chaos*, 27/11: 114309.
- Kovesi, P. (2015), “Good colour maps: How to design them”, arXiv: 1509.03700 [cs.GR].
- Krasovskii, N. N. (1964), “The approximation of a problem of analytic design of controls in a system with time-lag”, *Journal of Applied Mathematics and Mechanics*, 28/4: 876–85.
- Lanczos, C. (1938), “Trigonometric interpolation of empirical and analytical functions”, *Journal of Mathematics and Physics*, 17/1–4: 123–99.

- Magnus, J. R., and Neudecker, H. (1985), “Matrix differential calculus with applications to simple, Hadamard, and Kronecker products”, *Journal of Mathematical Psychology*, 29/4: 474–92.
- Mastroianni, G., and Szabados, J. (1995), “Jackson order of approximation by Lagrange interpolation. II”, *Acta Mathematica Hungarica*, 69/1–2: 73–82.
- Mattenet, S. M., De Iuliis, V., Gomez, M. A., Michiels, W., and Jungers, R. M. (2022), “An improved finiteness test and a systematic procedure to compute the strong $\mathcal{H}_2$ norm of differential algebraic systems with multiple delays”, *Automatica*, 144: 110495.
- Michiels, W. (2011), “Spectrum-based stability analysis and stabilisation of systems described by delay differential algebraic equations”, *IET Control Theory & Applications*, 5/16: 1829–42.
- Michiels, W., and Niculescu, S.-I. (2014), *Stability, Control, and Computation for Time-Delay Systems: An Eigenvalue-Based Approach* (2nd edn., Philadelphia).
- Michiels, W., and Zhou, B. (2019), “Computing delay Lyapunov matrices and $\mathcal{H}_2$ norms for large-scale problems”, *SIAM Journal on Matrix Analysis and Applications*, 40/3: 845–69.
- Minorsky, N. (1942), “Self-excited oscillations in dynamical systems possessing retarded actions”, *Journal of Applied Mechanics*, 9/2: A65–71.
- Mishkis, A. D. (1951), *Lineinye differentsialnye uravneniya s zapazdyvayushchim argumentom* (Moscow); trans. *Lineare Differentialgleichungen mit nachteilendem Argument* (Berlin, 1955).
- Munoa, J., Beudaert, X., Dombovari, Z., Altintas, Y., Budak, E., Brecher, C., and Stepan, G. (2016), “Chatter suppression techniques in metal cutting”, *CIRP Annals*, 65/2: 785–808.
- Neumann, C. (1862), *Ueber die Entwicklung einer Function mit imaginärem Argument nach den Kugelfunctionen erster und zweiter Art* (Halle).
- Nocedal, J., and Wright, S. J. (2006), *Numerical Optimization* (2nd edn., New York).
- Olver, S., Slevinsky, R. M., and Townsend, A. (2020), “Fast algorithms using orthogonal polynomials”, *Acta Numerica*, 29: 573–699.
- Olver, S., and Townsend, A. (2013), “A fast and well-conditioned spectral method”, *SIAM Review*, 55/3: 462–89.
- Orosz, G., and Molnár, T. G. (2025), *Dynamics and Control of Connected Vehicles* (Cham).

- Perron, O. (1913), *Die Lehre von den Kettenbrüchen* (Leipzig).
- Pontes Duff, I., Poussot-Vassal, C., and Seren, C. (2015), “Realization independent single time-delay dynamical model interpolation and $\mathcal{H}_2$ -optimal approximation”, in *Proceedings of the 54th IEEE Conference on Decision and Control*, 4662–7.
- (2018), “ $\mathcal{H}_2$ -optimal model approximation by input/output-delay structured reduced order models”, *Systems & Control Letters*, 117: 60–7.
- Provoost, E., and Michiels, W. (2024), “The Lanczos tau framework for time-delay systems: Padé approximation and collocation revisited”, *SIAM Journal on Numerical Analysis*, 62/6: 2529–48.
- (2025), “Accelerated H^2 -norm approximation for time-delay systems with discrete delays”, *IFAC-PapersOnLine*, 59/13, Presented at the *19th IFAC Workshop on Time-Delay Systems*: 93–8.
- (2026), “Computing and optimizing the H^2 -norm of delay differential algebraic systems”, arXiv: 2603.10734 [math.NA].
- Ross, D. (1971), “Controller design for time lag systems via a quadratic criterion”, *IEEE Transactions on Automatic Control*, 16/6: 664–72.
- Scholl, T. H., Hagenmeyer, V., and Gröll, L. (2024a), “Lyapunov–Krasovskii functionals of robust type and their Legendre-tau-based approximation”, *IFAC-PapersOnLine*, 58/27: 219–24.
- (2024b), “What ODE-approximation schemes of time-delay systems reveal about Lyapunov–Krasovskii functionals”, *IEEE Transactions on Automatic Control*, 69/7: 4614–29.
- Schulze, P., Unger, B., Beattie, C., and Gugercin, S. (2018), “Data-driven structured realization”, *Linear Algebra and its Applications*, 537: 250–86.
- Silviu-Ioan, F., Nakatsukasa, Y., Trefethen, L. N., and Beckermann, B. (2018), “Rational minimax approximation via adaptive barycentric representations”, *SIAM Journal on Scientific Computing*, 40/4: A2427–55.
- Sipahi, R., Niculescu, S.-I., Abdallah, C. T., Michiels, W., and Gu, K. (2011), “Stability and stabilization of systems with time delay”, *IEEE Control Systems Magazine*, 31/1: 38–65.
- Stewart, G. W., and Sun, J. (1990), *Matrix Perturbation Theory* (Boston).
- Szegő, G. (1975), *Orthogonal Polynomials* (4th edn., Providence).
- Tadmor, E. (1986), “The exponential accuracy of Fourier and Chebyshev differencing methods”, *SIAM Journal on Numerical Analysis*, 23/1: 1–10.

- Trefethen, L. N. (2013), *Approximation Theory and Approximation Practice* (Philadelphia).
- Unger, B. (2020), “Well-Posedness and Realization Theory for Delay Differential-Algebraic Equations”, PhD thesis (Technischen Universität Berlin).
- Vanbiervliet, J., Michiels, W., and Jarlebring, E. (2011), “Using spectral discretisation for the optimal $\mathcal{H}_2$ design of time-delay systems”, *International Journal of Control*, 84/2: 228–41.
- Vanbiervliet, J., Vandereycken, B., Michiels, W., Vandewalle, S., and Diehl, M. (2009), “The smoothed spectral abscissa for robust stability optimization”, *SIAM Journal on Optimization*, 20/1: 156–71.
- Vanbiervliet, J., Verheyden, K., Michiels, W., and Vandewalle, S. (2008), “A nonsmooth optimisation approach for the stabilisation of time-delay systems”, *ESAIM: Control, Optimisation and Calculus of Variations*, 14/3: 478–93.
- Villafuerte, R., Mondié, S., and Garrido, R. (2013), “Tuning of proportional retarded controllers: Theory and experiments”, *IEEE Transactions on Control Systems Technology*, 21/3: 983–90.
- Volterra, V. (1928), “Sur la théorie mathématique des phénomènes héréditaires”, *Journal de Mathématiques Pures et Appliquées*, 9th ser., 7: 249–98.
- Vyasarayani, C. P., Subhash, S., and Kalmár-Nagy, T. (2014), “Spectral approximations for characteristic roots of delay differential equations”, *International Journal of Dynamics and Control*, 2: 126–32.
- Wang, Z., Hu, H., Xu, Q., and Stepan, G. (2017), “Effect of delay combinations on stability and Hopf bifurcation of an oscillator with acceleration-derivative feedback”, *International Journal of Non-Linear Mechanics*, 94: 392–99.
- Wegert, E. (2012), *Visual Complex Functions* (Basel).
- Whittaker, E. T. (1902), “On a new connexion of Bessel functions with Legendre functions”, *Proceedings of the London Mathematical Society*, 1st ser., 35: 198–206.
- Wu, Z., and Michiels, W. (2012), “Reliably computing all characteristic roots of delay differential equations in a given right half plane using a spectral method”, *Journal of Computational and Applied Mathematics*, 236/9: 2499–514.
- Zhou, K., Doyle, J. C., and Glover, K. (1995), *Robust and Optimal Control* (Englewood Cliffs).

FACULTY OF ENGINEERING SCIENCE
DEPARTMENT OF COMPUTER SCIENCE
NUMERICAL ANALYSIS AND APPLIED MATHEMATICS
Celestijnenlaan 200A box 2402
B-3001 Leuven

