New Trends in Optimal Control

Warwick in Venice Workshop, 15-17 May 2024

Book of Abstracts

Organisers:

Alessandro Alla (Università Ca' Foscari Venezia)

Radu Cimpeanu (University of Warwick)

Susana N. Gomes (University of Warwick)

Day 1 – Wednesday 15th of May

Time	Speaker	Title
10:30-11:00	Organisers	Registration, welcome and introductions
11:00-11:30	Dante Kalise	Global stabilization for global optimization: controlling interacting particle systems
11:30-12:00	Chiara Segala	Kernel methods for interacting particle systems: mean-field limit and control surrogate modelling
12:00-12:30	Sara Bicego	Steady states finding via deflation and controlled stabilization for the Fokker-Planck equation
12:30-14:30		Lunch break
14:30 – 15:00	Zhengang Zhong	Multi-level optimal control with neural surrogate models
15:00 - 15:30	Alessandro Scagliotti	Adversarial training as minimax optimal control problems
15:30 - 16:00	Jan Heiland	Polytopic autoencoders for higher-order series expansions of nonlinear feedback laws
16:00 - 18:00	Everyone	Coffee, free time, discussion
18:00-onward		Reception

11:00 – Dante Kalise: Optimal stabilization for global optimization: controlling interacting particle systems

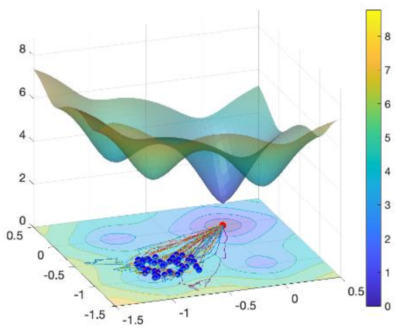

This talk explores the control of interacting particle systems to achieve desired stationary configurations, bridging the gap between microscopic particle dynamics and macroscopic mean-field descriptions. We focus on consensus dynamics in Consensus-Based Optimization (CBO) and introduce a controlled CBO framework that incorporates a feedback control term derived from the numerical solution of a Hamilton-Jacobi-Bellman equation. This control guides particles towards the global minimizer of the objective function, enhancing the performance of standard CBO methods.
We establish the well-posedness of the controlled CBO system and demonstrate its improved performance through numerical simulations. The controlled CBO framework offers a powerful tool for solving challenging optimization problems across various domains, such as machine learning, engineering design, and scientific computing.

11:30 – Chiara Segala: Kernel methods for interacting particle systems: mean-field limit and control surrogate modelling

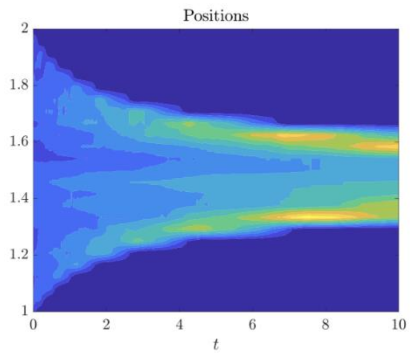

Interacting particle systems (IPS) are a very important class of dynamical systems, arising in different domains like biology, physics, sociology, and engineering. In many applications, these systems can be very large, making their simulation and control, as well as related numerical tasks, very challenging. Mean-field methods are an established approach to tackle large scale IPS, by considering only the distribution or density of all particles over the state space.

Very recently, the mean-field limit has been rigorously investigated in the context of kernel methods and their application to statistical machine learning, leading to the notion of the mean-field limit of a kernel and its reproducing kernel Hilbert space (RKHS). These developments open the path to kernel-based methods for large scale IPS, including learning and control of such systems. In this talk, numerical experiments on kernel methods for IPS will be shown, with a particular focus on control and learning problems on the large scale perspective.

12:00 – Sara Bicego: Steady states finding via deflation and controlled stabilization for the Fokker-Planck equation

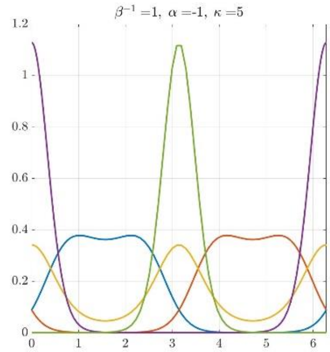

In the context of interacting particle systems, collective behavior is often described at the mean-field level by the time evolution of a probability density distribution governed by Fokker-Planck-type equations. The system's emergent patterns, based on the underlying microscopic interactions, are represented in terms of stationary states of the evolutionary partial differential equation. A common feature of these systems is the coexistence of various steady configurations and non-trivial phase transitions. The nature and number of steady states are linked to two key parameters: the noise amplitude and the interaction strength, measuring the interplay between diffusivity and drift, in addition to the modeling of the forces acting on the particle ensemble.
To capture the different steady states of the Fokker Planck equation, a Spectral Galerkin approximation is combined with a deflated Newton's method, factoring out roots as they are identified. Comparison with existing asymptotic analysis results allows for the verification of solutions and the classification between stable and unstable configurations. Once the steady states are found, an optimal control problem is designed to stabilize the system at a desired unstable steady state. The control action is computed via iterated open-loop solves in a receding horizon fashion.

12:30 – Lunch break

14:30 – Zhengang Zhong: Multi-level optimal control with neural surrogate models

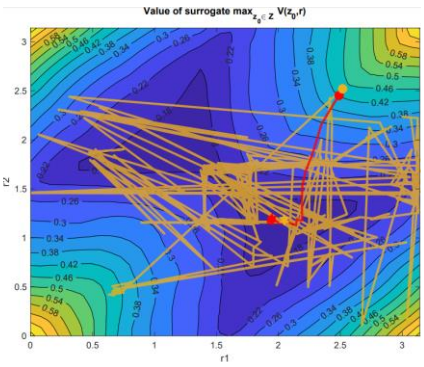

Optimal actuator and control design is studied as a multi-level optimization problem, where the actuator design is evaluated based on the performance of the associated optimal closed loop. The evaluation of the optimal closed loop for a given actuator realisation is a computationally demanding task, for which the use of a neural network surrogate is proposed. The use of neural network surrogates to replace the lower level of the optimization hierarchy enables the use of fast gradient-based and gradient-free consensus-based optimization methods to determine the optimal actuator design. The effectiveness of the proposed surrogate models and optimization methods is assessed in a test related to optimal actuator location for heat control.

15:00 – Alessandro Scagliotti: Adversarial training as minimax optimal control problems

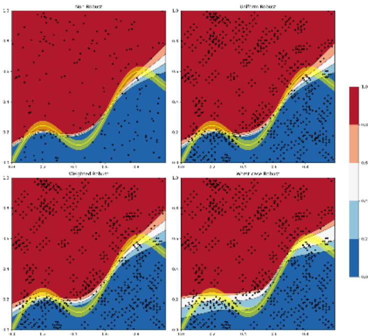

In this talk, we address the adversarial training of neural ODEs from a robust control perspective. This is an alternative to the classical training via empirical risk minimization, and it is widely used to enforce reliable outcomes for input perturbations. Neural ODEs allow the interpretation of deep neural networks as discretizations of control systems, unlocking powerful tools from control theory for the development and the understanding of machine learning. In this specific case, we formulate the adversarial training with perturbed data as a minimax optimal control problem, for which we derive first order optimality conditions in the form of Pontryagin's Maximum Principle. We provide a novel interpretation of robust training leading to an alternative weighted technique, which we test on a low-dimensional classification task.

15:30 – Jan Heiland: Polytopic autoencoders for higher-order series expansions of nonlinear feedback laws

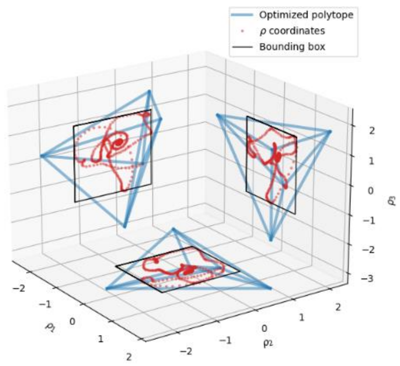

On the way to a computational and general-purpose approach to nonlinear controller design, the approximative embedding of nonlinear models in the class of quasi linear-parameter varying (LPV) systems seems a promising path. In this talk, we illustrate how the embedding and approximation works in general and highlight two recent research efforts for the efficient embedding and the controller design.

Firstly, we discuss autoencoders that provide low-dimensional parametrizations of states in a polytope (as opposed to a linear space). For nonlinear PDEs, this idea is readily used for low-dimensional linear parameter-varying approximations.

Secondly, we recall how LPV approximations can been the base for efficient nonlinear controller design via series expansions of the solution to the state-dependent Riccati equation.

Then, we adapt a general polytopic autoencoder for control applications and show how it outperforms standard linear approaches in view of LPV approximations of nonlinear systems and how the particular architecture enables higher order series expansions at little extra computational effort.

In a numerical study, we illustrate the procedure and how this combined approach can reliably outperform the standard linear-quadratic design.

16:00 – Discussion session (Topic TBC)

18:00 onwards – Reception

Day 2 – Thursday 16th of May

Time	Speaker	Title
09:30 - 10:00	Bosco García-Archilla	Optimal bounds for POD approximations of infinite horizon control problems based on time derivatives
10:00 - 10:30	Julia Novo	Optimal bounds for numerical approximations of infinite horizon problems based on dynamics programming approach
10:30 - 11:00		Coffee break
11:00 - 11:30	Luca Saluzzi	Decaying sensitivity for the resolution of high-dimensional optimal control problems
11:30 - 12:00	Vincent Liu	A Hamilton-Jacobi-Bellman approach to ellipsoidal approximations of reachable sets for linear time-varying systems
12:00 - 12:30	Tobias Ehring	Hermite kernel surrogates for the value function of high-dimensional nonlinear optimal control problems
12:30 - 14:30		Lunch break
14:30 - 15:00	Maria Strazzullo	Regularized full and reduced feedback control strategies for convection-dominated Navier-Stokes equations
15:00 - 15:30	Mattia Manucci	Model order reduction for large-scale switched differential-algebraic equations
15:30 - 16:00	Matteo Tomasetto	Real-time optimal control of parametrized systems by deep learning-based reduced order models
16:00 - onward	Everyone	Coffee, free time, discussion

09:30 – Bosco García-Archilla: Optimal bounds for POD approximations of infinite horizon control problems based on time derivatives

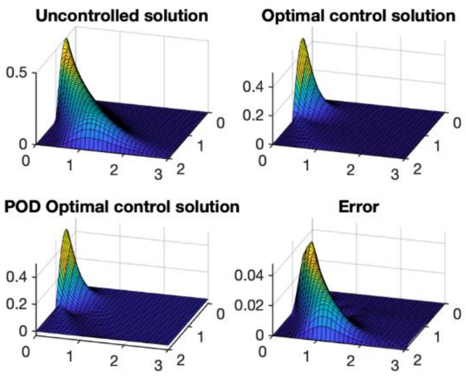

We consider the numerical approximation of infinite horizon problems via the dynamic programming approach. The value function of the problem solves a Hamilton-Jacobi-Bellman (HJB) equation that is approximated by a fully discrete method. It is known that the numerical problem is difficult to handle by the so called curse of dimensionality. To mitigate this issue we apply a reduction of the order by means of a new proper orthogonal decomposition (POD) method based on time derivatives.
We carry out the error analysis of the method using recently proved optimal bounds for the fully discrete approximations. Moreover, the use of snapshots based on time derivatives allow us to bound some terms of the error that could not be bounded in a standard POD approach. Some numerical experiments show the good performance of the method in practice.

10:00 – Julia Novo (& Javier de Frutos): Optimal bounds for numerical approximations of infinite horizon problems based on dynamic programming approach

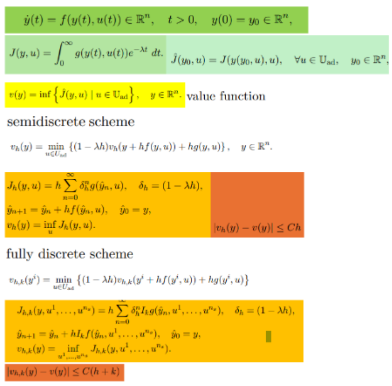

In this talk we show the error bounds that can be obtained for fully discrete approximations of infinite horizon problems via the dynamic programming approach. It is well known that considering only the time discretization with a positive step size h error bounds of size h are obtained for the difference between the value function (viscosity solution of the Hamilton-Jacobi-Bellman equation) and the value function of the discrete time problem. However, including also a spatial discretization based on elements of size k an error bound of size $O\left(\frac{k}{h}\right)$ can be found in the literature for the error in the fully discrete value function. In this talk, we revise the error bounds of the fully discrete method and prove that, under similar assumptions to those of the time discrete case, the error in the fully discrete case is $O(k+h)$ which gives first order in time and space for the method. This error bound matches the numerical experiments of many papers in the literature in which the behavior $\frac{1}{h}$ from the bound $O\left(\frac{k}{h}\right)$ had never been observed.

10:30 – Coffee break

11:00 – Luca Saluzzi: Decaying sensitivity for the resolution of high-dimensional optimal control problems

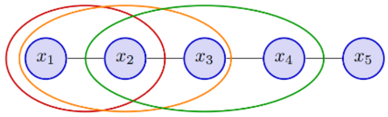

Optimal control problems arise in a variety of applications and one way to approach such problems is by computing the optimal value function by solving a Hamilton-Jacobi-Bellman (HJB)- PDE. However, the computational complexity of grid-based numerical methods for this problem grows exponentially with respect to the state space dimension, a phenomenon known as the curse of dimensionality.

To mitigate this problem, in this talk we describe the dynamical system as a set of subsystems interconn ected via a graph. It seems reasonable to expect that subsystems that are far away (in terms of the graph distance) only interact with each other very weakly. Based on this decaying sensitivity assumption between subsystems, we are able to build a separable approximation of the optimal value function based on neighborhoods in the graph. Under sufficient regularity assumptions, separable functions can be approximated by deep neural networks with a number of neurons that grows only polynomially in the state dimension, leading to a mitigation of the curse of dimensionality. We will then focus on the LQR case, where it is possible to prove this decaying sensitivity by studying the corresponding Algebraic Riccati Equation and we will investigate numerical approximations able to exploit the sparsity structure.

11:30 – Vincent Liu: A Hamilton-Jacobi-Bellman approach to ellipsoidal approximations of reachable sets for linear time-varying systems

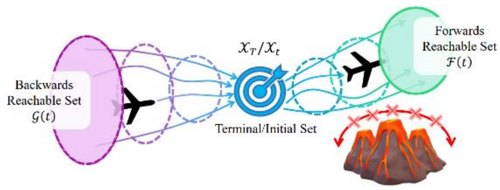

Society's ever-increasing integration of autonomous systems in day-to-day life has simultaneously brought forth concerns as to how their safety and reliability can be verified. To this end, reachable sets lend themselves well to this task. These sets describe collections of states that a dynamical system can reach in finite time, which can be used to guarantee goal satisfaction in controller design or to verify that unsafe regions will be avoided. However, general-purpose methods for computing these sets suffer from the curse-of-dimensionality, which typically prohibits their use for systems with more than a small number of states, even if they are linear. In this talk, we derive dynamics for a union and intersection of ellipsoidal sets that, respectively, under-approximate and over-approximate the reachable set for linear time-varying systems subject to an ellipsoidal input constraint and an ellipsoidal terminal (or initial) set. This result arises from the construction of a local viscosity supersolution and subsolution of a Hamilton-Jacobi-Bellman equation for the corresponding reachability problem. The proposed ellipsoidal sets can be generated with polynomial computational complexity in the number of states, making the approximation scheme computationally tractable for continuous-time linear time-varying systems of relatively high dimension.

12:00 – Tobias Ehring: Hermite kernel surrogates for the value function of high-dimensional nonlinear optimal control problems

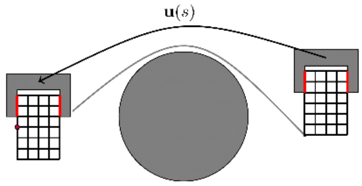

Numerical methods for the optimal feedback control of high-dimensional dynamical systems typically suffer from the curse of dimensionality. In the current presentation, we devise a mesh-free data-based approximation method for the value function of optimal control problems, which partially mitigates the dimensionality problem. The method is based on a greedy Hermite kern el interpolation scheme and incorporates context knowledge by its structure. Especially, the value function surrogate is enforced to be 0 in the target state, non-negative and constructed as a correction of a linearized model. The algorithm allows formulation in a matrix-free way which ensures efficient offline and online evaluation of the surrogate, circumventing the large-matrix problem for multivariate Hermite interpolation. Additionally, an incremental Cholesky factorization is utilized in the offline generation of the surrogate. For finite time horizons, both convergence of the surrogate to the value function and for the surrogate vs. the optimal controlled dynamical system are proven. Experiments support the effectiveness of the scheme, using among others a new academic model with an explicitly given value function. It may also be useful for the community to validate other optimal control approaches.

12:30 – Lunch break

14:30 – Maria Strazzullo: Regularized full and reduced feedback control strategies for convection-dominated Navier-Stokes equations

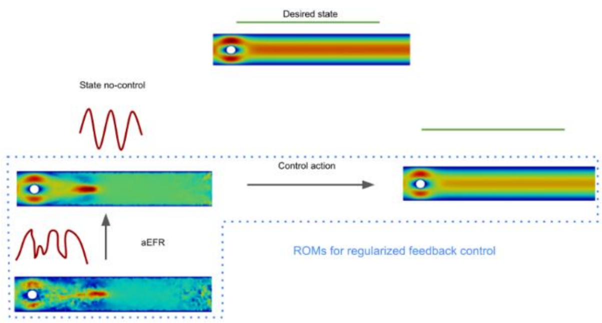

In numerous scientific and industrial contexts, controlling the flow regime is crucial. There is a growing interest in the subject with the main goal of steering the flow towards beneficial configurations, less turbulent and more laminar. To address this, we propose a novel linear feedback control method for the Navier-Stokes equations with high Reynolds numbers, ensuring exponential convergence to desired state. However, as Reynolds numbers increase, additional stabilization strategies become necessary. We exploit the Evolve-Filter-Relax (EFR) algorithm, showing i ts theoretical non-exponential convergence. Guided by these insights, we present an adaptive EFR (aEFR) approach, which mitigates numerical oscillations in controlled settings, restoring exponential convergence. Our theoretical framework is validated through numerical experiments on a 2D flow past cylinder with Reynolds 1000, both at full and reduced order model levels.

15:00 – Mattia Manucci (and B. Unger): Model order reduction for large-scale switched differential-algebraic equations

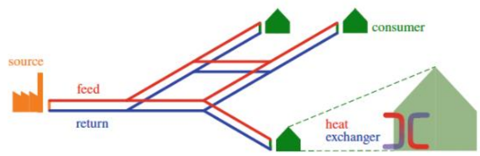

We discuss a projection-based model order reduction (MOR) for large-scale systems of switched differential-algebraic equations (sDAEs), i.e.,

$$(1) \qquad \Sigma_q \begin{cases} \boldsymbol{E}_{q(t)} \dfrac{dx(t)}{dt} = \boldsymbol{A}_{q(t)}\boldsymbol{x}(t) + \boldsymbol{B}_{q(t)}\boldsymbol{u}(t), & \boldsymbol{x}(t_0) = \boldsymbol{0}, \\ \boldsymbol{y}(t) = \boldsymbol{C}_{q(t)}\boldsymbol{x}(t), & \end{cases}$$

where $q\colon \mathbb{R} \to \mathcal{J} \coloneqq \{1, \dots, M\}$ is the switching signal, i.e., a piecewise constant function taking values in the index set $\mathcal{J}$, $\boldsymbol{x}(t) \in \mathbb{R}^n$, $\boldsymbol{u}(t) \in \mathbb{R}^m$, and $\boldsymbol{y}(t) \in \mathbb{R}^p$ denote respectively, the *state*, the controlled *input*, and the measured *output*. We emphasize that the matrices $\boldsymbol{E}_j \in \mathbb{R}^{n \times n}$ for $j \in \mathcal{J}$ may be singular. Control systems of sDAEs may arise in modelling physical systems with algebraic constraints and piecewise time-dependent parameters, like Stokes control system with piecewise time-dependent diffusion. If (1) has to be evaluated repeatedly, one can rely on MOR and replace (1) by the *reduced-order model*

$$(2) \qquad \tilde{\Sigma}_q \begin{cases} \widetilde{\boldsymbol{E}}_{q(t)} \frac{d\tilde{\boldsymbol{x}}(t)}{dt} = \widetilde{\boldsymbol{A}}_{q(t)} \tilde{\boldsymbol{x}}(t) + \widetilde{\boldsymbol{B}}_{q(t)} \boldsymbol{u}(t), \qquad \tilde{\boldsymbol{x}}(t_0) = \boldsymbol{0}, \\ \qquad \tilde{\boldsymbol{y}}(t) = \widetilde{\boldsymbol{C}}_{q(t)} \tilde{\boldsymbol{x}}(t), \end{cases}$$

with $r \ll n$, $\widetilde{\boldsymbol{E}}_j, \widetilde{\boldsymbol{A}}_j \in \mathbb{R}^{r \times r}$, $\widetilde{\boldsymbol{B}}_j \in \mathbb{R}^{r \times m}$, and $\widetilde{\boldsymbol{C}}_j \in \mathbb{R}^{p \times r}$ for $j \in \mathcal{J}$. Model reduction for systems of switched ordinary differential equations (sODEs) has been addressed in many works, see for instance [1,3,4] and reference therein, while, to the best of our knowledge, MOR of sDAEs only appears in [2] for a known switching signal and with several limitations to the large-scale setting, see [2, Sec. 5]. For our MOR scheme, we rely on a reformulation of (1) as a system of sODEs with state jumps at the switching times [2] and successively show that the method proposed in [3] can be successfully applied, in this generalized settings, to derive a reduced-order model for the set of generic admissible switching signals.

References:

[1] I. V. Gosea, M. Petreczky, A. C. Anntoulas, and C. Fiter. *Balanced truncation for linear switched systems*. Adv. Comput. Math., 44(6):1845–1886, 2018.

[2] M. S. Hossain and S. Trenn. *Model reduction for switched differential-algebraic equations with known switching signal*. Technical report, 2023.

[3] I. Pontes Duff, S. Grundel, and P. Benner. *New gramians for switched linear systems: Reachability, observability, and model reduction*. IEEE Trans. Automat. Control, 65(6):2526–2535, 2020.

[4] P. Schulze and B. Unger. *Model reduction for linear systems with low-rank switching*. SIAM J. Cont. Optim., 56(6):4365–4384, 2018.

15:30 – Matteo Tomasetto (& Francesco Braghin, Andrea Manzoni): Real-time optimal control of parameterized systems by deep learning-based reduced order models

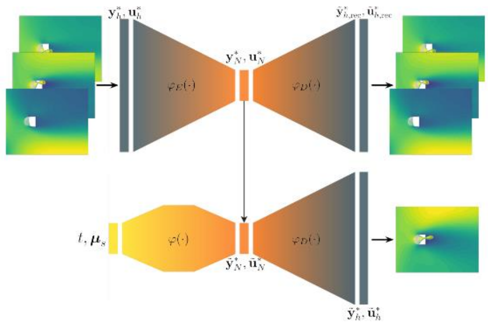

Many optimal control problems require suitable strategies in order to steer instantaneously the considered dynamics. Moreover, the control action needs to be updated whenever the underlying scenario undergoes variations, as often happens in applications. Fullorder models based on, e.g., Finite Element Method, are not suitable for these settings due to the computational burden. In addition, conventional reduced order modeling techniques, such as the Reduced Basis method, are linear, intrusive, and usually not efficient in addressing nonlinear time-dependent dynamics. We thus propose a nonlinear, non-intrusive Deep Learning-based Reduced Order Modeling (DL-ROM) technique to control rapidly parametrized PDEs under different scenarios. After optimal snapshots generation, dimensionality reduction and neural networks training in the offline phase, optimal control strategies can be retrieved online in real-time for all the scenarios of interest. The speedup and the high accuracy of the proposed approach have been assessed on different PDE-constrained optimization problems, ranging from the minimization of energy dissipation under Navier-Stokes equations to thermal active cloaking.

16:00 – Discussion session (Topic TBC)

Day 3 – Friday 17th of May

Time	Speaker	Title
10:00-10:30	Giacomo Albi	Feedback control strategies for magnetically confined fusion plasma
10:30-11:00		Coffee break
11:00-11:30	Oscar Holroyd	Feedback control of thin liquid films
11:30-12:00	Alexander Wray	Controlling Navier-Stokes via reduced order models
12:00-12:30	Organisers	Closing and discussions
12:30-14:30		Lunch

10:00 – Giacomo Albi: Feedback control strategies for magnetically confined fusion plasma

In this talk we address the challenge of confining high-temperature plasma in magnetic fusion devices. Our method focuses on employing instantaneous control techniques to guide the plasma within a specified spatial region. We model the process using the Vlasov-Poisson equation in a bounded domain, considering self-induced electric fields and external magnetic fields. Our control strategy incorporates real-time feedback, allowing for the direct minimization of a specified cost function within the particle interactions of the Vlasov model. Numerical simulations validate the effectiveness of our approach, demonstrating the ability of external magnetic fields to steer the plasma away from device boundaries.

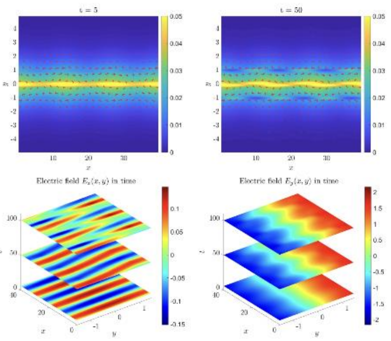

10:30 – Coffee break

11:00 – Oscar Holroyd: Feedback control of thin liquid films

We propose a method to stabilise an unstable solution to equations describing the interface of thin liquid films falling under gravity with a finite number of actuators and restricted observations. As for many complex systems, full observation of the system state is challenging in physical settings, so methods able to take this into account are important. The Navier-Stokes equations modelling this interfacial flow are a complex, highly nonlinear set of PDEs, so standard control theoretical results are not applicable. Instead, we chain together a hierarchy of increasingly idealised approximations, developing a control strategy for the simplified model which is shown to be successfully applicable to direct numerical simulations of the full system.

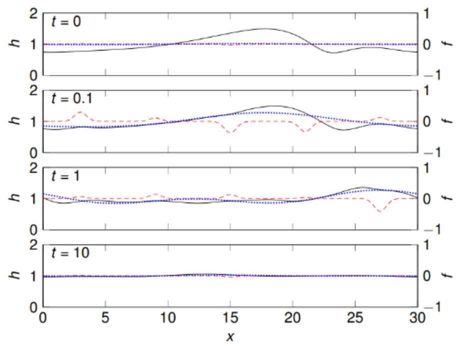

11:30 – Alexander Wray: Controlling Navier-Stokes via reduced order models

As in many fields, control has long been a key focus in fluid dynamics. Unfortunately, for multiphase flows (such as interfacial flows), the governing equations are the highly nonlinear Navier-Stokes equations, which must be solved on a time-evolving, non-rectilinear domain. Even for modest problem sizes, this presents a significant computational challenge. As a consequence, most inverse/control problems have only been examined in single phase contexts. For multiphase problems, the forward problem has often been treated via physics-driven reduced order modelling, but even these can be challenging to control due to issues of stiffness.

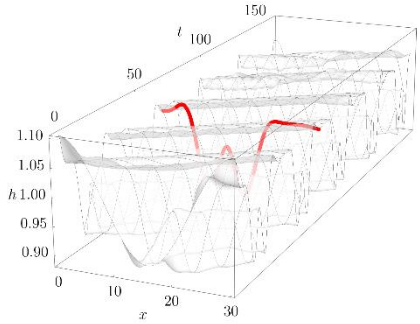

Here we resolve these issues in order to allow for control of the full Navier-Stokes equations for a multiphase system. In particular, we examine the paradigmatic case of free-surface flow down an inclined plane. The system is simulated using the open-source volume-of-fluid solver Basilisk. Control is implemented via development of an extremely high-fidelity reduced order model using a projection method akin to the Method of Weighted Residuals, and the use of a Model Predictive Control loop to control the direct numerical simulation with judicious use of the model. Actuation is achieved via imposition of a spatiotemporally varying electric potential on an electrode parallel to the substrate.

The model is investigated in detail, demonstrating a high degree of accuracy even into the short-wave regime. The control mechanism is shown to be applicable to both uniform and non-uniform target states, and the efficacy of the model predictive control loop is investigated across a wide variety of parameters.

12:00 – Closing, discussions, and lunch