Ph.D. in Information Technology
Thesis Defenses

March 3rd, 2025
At 3:00 p.m.
Room Alpha – Building 24

Raffaele Giuseppe CESTARI – XXXVII Cycle

ADVANCED MODEL PREDICTIVE CONTROL OF WATER RESERVOIR SYSTEMS

Supervisor: Prof. Simone Formentin

Abstract:

In this thesis project, we developed advanced MPC (Model Predictive Control) techniques in water resources management. Controlling one or more water basins means managing multiple objectives simultaneously, typically in conflict with each other. The controller must satisfy the water demand of downstream agricultural districts and protect coastal cities from damage caused by flooding during the rainiest periods. The regulator is also responsible for preventing dry conditions of the water basin itself as they can cause an impediment to the normal functioning of river and lake transport, which can lead to indirect economic losses due to lower tourist attractiveness. Furthermore, if the basin corresponds to a hydroelectric power production plant, it is the controller's responsibility to maximize its yield. The combined management of all these control objectives highlights the complexity of the problem. However, it is necessary to add a further complicating factor in regulating this type of system: the disturbance entering the system represented by the inflow to the basin. This quantity is difficult to measure due to the countless sources from which it comes. The inflow depends on meteorological phenomena, large-scale climatic phenomena, and the management of any upstream basins, which may be independent of the management of the basin for which the controller must be designed.

During the thesis project, we developed predictive control methodologies in this panorama that could include these delicate aspects in the design phase. Starting from an MPC control architecture with hourly resolution and assuming we have perfect knowledge of the inflows, we demonstrate the approach's effectiveness in the control of lake Como by outperforming the historical regulation and being able to track the ideal performance of a DDP controller. Subsequently, still working on lake Como, we removed the hypothesis of perfect knowledge of the inflows and concentrated on designing control architectures that could manage them. Initially, we proposed a Scenario MPC architecture, in which we optimize the control action based on a set of possible system evolutions derived from scenarios identification on historical data. This solution provides robust control that statistically outperforms deterministic counterparts relying on a single prediction model. Subsequently, we treated the same problem by changing perspective: inserting the best possible model into the prediction model, which we demonstrate is not necessarily the best model identified on the data but is undoubtedly the best model to use in a closed loop. To find this model, we select within a chosen class the one that minimizes a cost given by the combination of the identification of historical data and the performance obtainable using that model on a backtesting window within the control scheme.

We demonstrate that the methodology is particularly effective compared to traditional identification and control methods. We validate the performance on lake Como based on six years of actual data, demonstrating its particular effectiveness in dry periods. The thesis project continues by trying to answer the question: Is it possible to integrate an optimization of the cost function parameters from a multi-objective perspective into the MPC formulation? To answer this question, we designed a control architecture that automatically calibrates the combination of weights that constitute the cost function through Bayesian Optimization to maximize the quality of the Pareto frontier. At each interaction, we simulate the effect of the control action of a set of identical controllers except for the combination of weights of their cost function. This allows us to obtain a Pareto frontier in simulation, evaluated via multi-objective metrics and progressively improved according to the convergence of Bayesian Optimization. This solution can overcome the limitations of equivalent approaches in the literature since it guarantees convergence of the calibration algorithm. The case study of this methodology is the Red River system, a complex combination of interconnected water reservoirs and irrigation channels.

The system's complexity requires using an internal model for the MPC, which can capture the input/output relations of the original system while maintaining the possibility of expressing the Jacobians explicitly to be included in the prediction model effectively. The architecture identified that respects these constraints is the Neural States Space. As far as we know, this represents a unique example in the literature for predictive control in water resources. Finally, in this thesis, we discuss a further development of an advanced predictive control technique in the context of water resources, applied to lake Como, aiming to respond to a need: guaranteeing optimal control of the regulated basin even in the event of cyber threats, specifically Denial-of-Service attacks. This need arises from the awareness that these infrastructures represent a privileged target in an international cyber warfare context. Our proposed architecture consists of 2 MPC controllers tackling 2 different optimization problems at the same time. The first primary controller determines an optimal control action that balances system performance with the deviation from the control action calculated by the second controller. The second controller assumes that system control will be handed over to a hacker at the next time step for the remainder of the prediction horizon. Specifically, the aggressor is assumed to maximize the system's damage (i.e., in a cyber warfare context, its goal is sabotaging the targeted infrastructure). Consequently, the attacker's optimization problem is adversarial, seeking to achieve the opposite of the optimal control regulator solution at each time step. Recognizing that it is not trivial to correctly balance the relative weight of the safe component concerning the performance component within the main controller, we find that the best solution is to calibrate this weight online, called the resilience factor, identifying it by training a Hawkes process on historical series of attack occurrences, improving the estimate of the return time of the next attack and adjusting resilience accordingly. We validate the architecture in two attack scenarios. We outperform our counterparts in both cases based on the exclusive focus on performance or security. In summary, this doctoral thesis offers the scientific community an in-depth study of the problems of water resource management and how these problems have been addressed through the development of MPC architectures, with increasing complexity, treating the management of disturbances from a robust and from a controloriented perspective, addressing the issue of multi-objective optimization and the critical role that regulated basins play as strategic infrastructures of national importance.

Kristoffer Fink LØWENSTEIN – XXXVII Cycle

PHYSICS-INFORMED ONLINE LEARNING OF GRAY-BOX MODELS BY MOVING HORIZON ESTIMATION AND EFFICIENT NUMERICAL METHODS

Supervisor: Prof. Daniele Bernardini

Abstract:

Advanced human-engineered systems will inevitably play a more dominant role in the future and are expected to operate with increased autonomy benefiting society as a whole. Major challenges lie ahead to ensure that such complex systems operate in a safe manner while satisfying strict performance requirements. A compelling idea is to use optimization-based methodologies such as Moving Horizon Estimation (MHE) and Model Predictive Control (MPC) for autonomous decision-making as they allow incorporating complex objectives and critical measures through a cost function and constraints on system states and inputs. A fundamental requirement for MHE and MPC is a reliable, but computationally light, model enabling online evaluation. Acquiring such a model and, especially, adapting it to temporal variations of the underlying system dynamics, surrounding environment, or part-to-part variations found in any system is a challenging endeavor. This issue is addressed by introducing parametric gray-box models: Relying on physics-based modeling facilitates the integration of function approximators such as Feed Forward Neural Networks (FFNN) or Recurrent Neural Networks (RNN) of rather limited size as the data-driven component in the gray-box models, and also benefits from increased model interpretability compared to pure black-box models. By deploying these gray-box models in an MHE scheme, a novel MHE framework for physics-informed learning is introduced. The learning is made safe through constraints consistent with physical laws seamlessly integrated into the optimization problem. The proposed MHE scheme is inherently suitable for online application and provides concurrent state estimation and model adaptation with uncertainty quantification. Thus, the proposed MHE framework can support any advanced decision-making process that relies on accurate state and parameter estimates, being autonomous or human-assisted, extending beyond just MPC. Furthermore, an offline training algorithm capitalizing on the proposed MHE scheme tailored for gray-box models is presented, which provides a viable alternative to classical training algorithms for such a class of models. The utility of the methodology is demonstrated through two numerical experiments, showing promising results for estimation, prediction, closed-loop Nonlinear MPC (NMPC), and online model adaptation.

For real-time implementation of MHE and MPC, the underlying optimization algorithms are of paramount importance. In linear MHE and MPC the problem boils down to a Quadratic Program (QP) while Sequential Quadratic Programming (SQP), a cherished solution procedure for Nonlinear MHE (NMHE) and NMPC, is based on solving a series of related QPs. To meet this demand, QPALM-OCP is introduced: QPALM-OCP is a Proximal Augmented Lagrangian Method (P-ALM) tailored for the multi-stage structure arising in MHE and MPC. The algorithm relies on a semi-smooth Newton method to solve the inner ALM problem allowing multiple active set changes between iterates and benefits from warm-starting. Due to specialized low-rank matrix modifications, the iterates remain cheap while providing fast execution times. Contrary to conventional Quadratic Programming solvers (QP solvers), QPALM-OCP comes with guarantees of R-linear convergence to a stationary point of nonconvex QPs, making it an ideal candidate for SQP methods, especially when using exact second-order information. A prototype implementation, despite being non-optimized, outperforms state-of-the-art general-purpose QP solvers and is slightly faster than state-of-the-art Optimal Control Problem (OCP)-specific solvers in numerical experiments, demonstrating great potential for further development.

PhD Committee

Prof. Gian Paolo Incremona, **Politecnico di Milano**
Prof. Mario Zanon, **ETH Zürich**
Prof. Antonio Ferramosca, **Scuola IMT Alti Studi Lucca**