

Towards Real Time Control of Water Engineering with Nonlinear Hyperbolic Partial Differential Equations

Fabio Difonzo, Michael Holst, Morteza Kimiaei,
Vyacheslav Kungurtsev, Songqiang Qiu

December 2025

Abstract

This paper examines aspirational requirements for software addressing mixed-integer optimization problems constrained by the nonlinear Shallow Water partial differential equations (PDEs), motivated by applications such as river-flow management in hydropower cascades. Realistic deployment of such software would require the simultaneous treatment of nonlinear and potentially non-smooth PDE dynamics, limited theoretical guarantees on the existence and regularity of control-to-state mappings under varying boundary conditions, and computational performance compatible with operational decision-making. In addition, practical settings motivate consideration of uncertainty arising from forecasts of demand, inflows, and environmental conditions.

At present, the theoretical foundations, numerical optimization methods, and large-scale scientific computing tools required to address these challenges in a unified and tractable manner remain the subject of ongoing research across the associated research communities. Rather than proposing a complete solution, this work uses the problem as a case study to identify and organize the mathematical, algorithmic, and computational components that would be necessary for its realization. The resulting framework highlights open challenges and intermediate research directions, and may inform both more circumscribed related problems and the design of future large-scale collaborative efforts aimed at addressing such objectives.

Contents

1 Introduction

4

2	Formal Problem Statement and Assumptions	6
2.1	River Flow Modeling	7
2.2	Groundwater flow modeling and control	10
2.3	Relevant Background on PDE-Constrained Optimization	11
2.3.1	On Control of Hyperbolic PDEs	11
2.3.2	On Mixed Integer PDE-Constrained Optimization	14
2.4	Complete Stochastic Dynamic Programming Formulation	16
2.5	Background related to Stochastic PDE Dynamic Programming	17
2.5.1	Stochastic Discretization	17
2.5.2	Stochastic PDE Control	18
2.5.3	On Almost Sure State Constraints	20
2.5.4	Dynamic and Multistage Stochastic Programming with PDEs	22
2.5.5	Approximate Dynamic Programming of PDE Systems	22
3	On the Control to State Map	23
3.1	State Equation Solution Existence	23
3.2	Dual and Adjoint Considerations	26
4	Three-Time Scale Control - Summary	27
5	Offline Computation of a Catalog of Reference Solutions	31
5.1	Formulations	32
5.1.1	Multistage Stochastic Programming	32
5.1.2	Ergodic Optimization	33
5.1.3	Hierarchical Discrete-Continuous Formulation	34
5.2	Numerical Mixed Integer Optimization Methods	35
5.2.1	B&B	35
5.2.2	Other Offline Mixed-Integer Strategies	36

5.2.3	Comparison of Offline Mixed-Integer Methods for PDE-Constrained Optimization	38
5.3	Learning-based Approaches	38
5.3.1	Predicting Feasibility via Supervised Learning	39
5.3.2	Optimizing Discrete-Continuous Controls via Reinforcement Learning	39
5.3.3	Estimating Domains of Attraction for Continuous Controls	40
5.4	Catalog Construction	41
5.5	Open Challenges and Future Research	43
6	Meso Sequential Numerical Optimization	43
6.1	ADP-MPC Formulation	44
6.2	Sequential QCQP (SQCQP) Method	45
6.2.1	Normal and Tangential Step Computation	46
6.2.2	Filter and Trust Region Update	48
6.3	Open Challenges and Future Research	52
7	Offline-Meso Integration	52
7.1	Open Challenges and Future Research	56
8	Fast Newton Solutions of Reduced Order Models at Real Time	57
8.1	Real Time Control	57
8.2	Recent Advances in Mixed Integer Real Time Control	58
8.3	Piecewise Linear Approximation of the Shallow Water Equations	59
8.4	Open Challenges and Future Research	61
9	Meso Real-Time Integration	61
9.1	Open Challenges and Future Research	65
10	Discussion and Conclusion	66

1 Introduction

Consider that one wishes to control a system defined by a nonlinear time-dependent Partial Differential Equation (PDE). Moreover, the control decisions include both continuous and discrete variables; the presence of discrete variables induces a nonconvex feasible set and hence additional nonconvexity in the optimization problem. In addition, the control is intended to be performed in real-time, that is, with control decisions algorithmically computed within a short time frame. Finally, some quantities that influence the dynamics or the desired output of the system are uncertain. These uncertainties have probabilistic forecasts, and decisions must be made with this uncertainty in mind, whether by targeting an expectation of the desired outcome and/or a quantity representing risk aversion.

Such a problem arises in a number of applications. A prominent example is in hydrology, as in the operation of a dam or network of dams in service of generating hydropower to meet energy demand, which can subside or surge in real time. So far, nonlinear control algorithms have been developed applying a nonlinear hydropower production function (e.g., [53]) as well as empirical studies on, e.g., efficiency [127]. However, being able to successfully incorporate accurate models of the river flow, together with important physical (e.g., environmental) constraints, can facilitate a more realistic model and thus more simultaneously efficient and reliable operation.

One can also consider groundwater for the irrigation of a set of plots of land, with some growing agricultural produce and others being pasture for a range of livestock. By considering the porosity of the soil and modeling the water source input from precipitation and river flow at different times of the year, one can aim to direct it towards robust storage for possible drought periods, as well as the ideal irrigation for maximizing agricultural yield (see, e.g., [83, 2]).

These two examples will be the focus of this paper, with particular attention devoted to the first. However, there are many other circumstances and domains in which this problem can occur. For instance, in automated manufacturing and process control, a number of operations can be modeled with evolution PDEs. For instance, a chemical refinery plant could perform processing that can be modeled thermodynamically with a heat equation plus an additional nonlinear dissipative term.

In general, this is an intractable problem. That is, the combination of a nonlinear infinite-dimensional system, nonconvexities in the decisions, uncertainty, and fast computation for real-time solution is a set of requirements that far exceed the capability of state-of-the-art methods and software. Whereas an aspirational volume on real-time PDE-constrained optimization methods has appeared even in the early volume [14], in general, there are very few works in the literature with real implementation or even realistic simulations under the defined operational requirements. For a rare exception, see, e.g., [11], wherein while the sys-

tem is multi-component, Newton iterations are numerically reliable, facilitating real-time computation, although still uncertainty and combinatorial decisions are not considered in that work. Other recent attempts at real-time control of systems governed by PDEs incorporate neural network models (e.g., [122]); however, these often fail to generalize, and their fundamental black box nature and stochasticity present unacceptable hazards for any mission-critical system.

However, while both methods and software are a significant distance from being available to adequately solve nonlinear time-dependent PDE mixed-integer optimization, this paper presents a framework and overall methodology whereby this problem can be approached over the long run by researchers in the field. In particular, through a properly motivated and reasonable incorporation of existing tools, incremental extensions of available methods and analysis, together with a few larger research and development efforts to circumnavigate certain fundamental obstacles, one can attempt to tackle such a control problem.

This paper presents the following contributions in the direction of solving this problem. While it is not authoritatively comprehensive due to the significant research and computational developments necessary to realize the entire proposed operation, this manuscript can form the basis of a large, potentially multi-institutional formal project or serve as a template by which decentralized advances can yield, in the long run, a workable solution.

- We present a theoretical analysis of the PDE under consideration, in particular the appropriate function space to consider solutions. A novel local existence theory is presented for the nonlinear nonconservative form of the shallow water equations. The required challenges, as far as the lack of a global unique solution and a dual space amenable to computation, are defined, motivating the procedures in the rest of the paper.
- A novel three-tier computational architecture is developed. This includes an offline layer for exploiting HPC and machine learning, a meso layer for performing classical Approximate Dynamic Programming and Nonlinear-Optimization-based Model Predictive Control, and a real-time layer for online operation, together with desiderata and protocols for their integration.
- The details of offline layer are presented, including a novel use of machine learning tools to find a region of the control-to-state map on which local existence and regularity permits numerical optimization.
- The details of a meso layer that combines particle evolution strategies together with a novel Primal (due to the aforementioned dual space challenges) Sequential Quadratic Programming with Quadratic Constraints Algorithm for performing Nonlinear Model Predictive Control
- The appropriate tools for tractable and reliable real-time operation are proposed, with reduced order modeling used to define tractable optimiza-

tion problems for online operation, with feedback at reasonable latency across the layers to enable maximal instrumental utility of the more complex computations at the meso and offline layers.

The computational systems architecture presented is not only multilevel to account for limitations in real-time hardware latency, but the motivating application of dam networks naturally involves multiple, well-separated operational time scales. Long-term planning decisions, such as reservoir allocation and policy design, are typically performed over horizons of days to weeks. Intermediate supervisory adjustments, incorporating updated inflow forecasts and system measurements, occur on the scale of hours, for example through rolling-horizon updates every 1–6 hours. At the fastest level, hydraulic control actions such as gate movements and flow regulation must respond to system dynamics on time scales ranging from seconds to minutes. This separation of time scales directly motivates the three-layer computational architecture developed in this work, consisting of offline, mesoscopic, and real-time components.

2 Formal Problem Statement and Assumptions

To begin with, the problem is stated without considering noise and uncertainty, i.e., deterministic PDE-constrained optimization.

Formally,

$$\begin{aligned} \min_{u,v,z} \quad & \int_{t=0}^T (-\Pi(u, d) + C(v, z)) dt \\ \text{s.t.} \quad & e(u(x, t)) = f(v(t), z(t)), \\ & v \in \mathcal{V}, \\ & S(u) \geq 0. \end{aligned} \tag{1}$$

This defines a PDE-constrained optimization problem for a process u unrolling over time t and space x . In the standing example of interest, the process has multiple vector fields, the water level and the flow rate, however for simplicity we present the coupled PDE with $e(u(x, t)) = f(v(t), z(t))$, where $v(t) \in \mathbb{R}^{n_c}$ is a continuous finite dimensional control decision variable (e.g., the aperture of a turbine) and $z(t) \in \{0, 1\}^{n_b}$, e.g., the opening and closing of a dam. In practice, the influence of the decision variables on the state arises from the boundary conditions, the details of which are presented below.

The functional Π is an Economic Objective, typically the profit. Economic and more classical tracking objective functionals present important distinctions for Model Predictive Control, with stability guarantees requiring distinct problem formulations and analysis. A tracking objective, i.e., $\int \|u - u_T\|^2 dt$, for an off-line computed setpoint u_T is strongly convex, presenting a natural Lyapunov function facilitating stability, while Economic objective functionals do not.

The cost function C for the decisions v and z represents real costs of operating the water engineering, and the function also acts as a regularizer. Here, there is a fortunate correspondence between economic cost and regularization: both the total time and effort in operation, that is, the time-measure of an open dam and its wide opening, as well as its time rate of change, that is, the process of opening and closing, are penalized. The sum of the total norm of the two together is effectively Sobolev-like.

The constraint $\mathcal{V}$ is typically simple, e.g., lower and upper bounds on the components, for the continuous decisions. Fortunately, it is both simultaneously natural and physical, while also advantageous numerically, for the continuous decisions to be constrained to a compact set, e.g., $0 \leq v \leq \bar{v}$ for (almost) every x . Operationally, there are clear limits present for any driving force, while compactness facilitates the lower semicontinuity properties that are amenable to iterative procedures asymptotically converging to a minimum.

Finally, $S(u)$ is a *state constraint* that can require that the solution of the water flow exhibits certain physical requirements. This can include a) physical realizability, e.g., a nonnegative water height level, b) safety, e.g., non-flooding velocity, and c) externally mandated operational requirements, e.g., environmental constraints ensuring sufficient flow for fish migration. The state constraints present additional challenges as far as regularity; the dual for such a constraint is a Borel measure. In the sequel, the suggested strategy for the state constraint is to add a barrier term in the objective penalizing proximity of $S(u)$ to its bound. In practice, this would ensure a more conservative constraint adherence (that is, at a greater distance from the bound) than is necessary, which is reasonable for mission-critical systems.

2.1 River Flow Modeling

The shallow-water (Saint–Venant) family of systems considered in this work may be regarded as a hierarchy of reduced fluid models whose fidelity and numerical requirements vary with the physical processes retained [78, 118]. For the purposes of control and real-time optimization, it is useful to highlight the modeling choices that most directly affect numerical solvability, robustness, and cost. At the lowest complexity, the one-dimensional depth-averaged Saint–Venant (conservative) formulation provides a hyperbolic balance law for cross-section area (or depth) and discharge. Two-dimensional depth-integrated shallow-water equations extend this model to account for lateral velocity components and nontrivial bathymetry. When dispersion or non-hydrostatic pressure becomes important (for instance, during runup or for short waves interacting with bathymetry), non-hydrostatic or multilayer extensions should be considered and treated with corresponding numerical care. See [10] for more details.

Practical river and dam models include topographic source terms (bed slope),

frictional source terms (e.g., Manning or Darcy–Weisbach), and other nonconservative products (for example, parameterized exchange or porosity terms). Such source terms can be stiff and are often spatially nonsmooth (for example, step changes in bed elevation or wetting/drying fronts). These characteristics require numerical methods that carefully balance flux gradients against source terms and that preserve physically meaningful invariants such as mass and nonnegativity of the depth.

Robust numerical solvers for hydrological and river Computational Fluid Dynamics (hereafter, CFD) that are suitable for control integration generally satisfy three interrelated properties. First, they are well-balanced: they preserve steady states in which flux gradients are exactly balanced by source terms (for example, the lake-at-rest state), which prevents spurious waves and allows accurate long-time integration around controller operating points [6]. Second, they are positivity-preserving and handle wetting/drying in a stable manner so that water depth remains nonnegative even during transitions between inundated and dry states [76]. Third, they treat friction and other stiff sources in a stable way; common remedies include semi-implicit source discretizations or operator splitting with local implicit solves when tight stability constraints appear.

Finite volume (conservative) discretizations on structured or unstructured meshes are the dominant choice for hyperbolic shallow-water problems, since they provide local conservation, shock-capturing, and natural coupling with Riemann solver-based fluxes (e.g., see [78, 118]). High-order finite-volume or discontinuous Galerkin methods are used where higher accuracy is required in smooth regions, but they demand stronger limiters and increase computational cost. Mixed finite-element and stabilized continuous-Galerkin formulations are also employed in some operational hydrodynamic models because of their flexibility on complex geometries.

Hydrostatic reconstruction techniques and positivity-preserving limiters provide a robust way to treat bed-slope source terms and wetting/drying indicators of the form $H(\eta, z_b) = \max\{0, z_b + \eta\}$ (e.g., see [17, 18] and Section 8.3 below). Particular care is required at thin water layers near dry states: specialized reconstructions or local mesh refinement help avoid oscillations and preserve mass and nonnegativity [76].

When primitive (nonconservative) formulations are used, the numerical flux must be combined with a consistent discretization of nonconservative products; path-conservative schemes provide a rigorous framework for this purpose [91, 33]. A careful treatment of these terms is important for accuracy and for ensuring that discrete adjoint and linearized models (used by control and optimization routines) are consistent with the forward solver.

For gradient-based optimization and real-time Newton/SQP solvers, it is necessary to obtain consistent linearizations of the discrete forward solver or to use algorithmic differentiation of the numerical method to produce accurate

sensitivities.

Saint Venant System Consider the Saint Venant system, which depth averages the Shallow Water Equations to produce a one-dimensional nonlinear hyperbolic PDE, as follows

$$\begin{aligned} \frac{\partial A}{\partial t} + \frac{\partial Q}{\partial x} &= 0, \\ \frac{\partial Q}{\partial t} + \frac{\partial(Q^2/A + gA^2/2)}{\partial x} &= gA(S_o - S_f), \text{ with,} \\ S_o &:= -\frac{\partial z_b}{\partial x} \\ S_f &:= \frac{n_M^2 Q |Q|}{A^{10/3}}, \\ A(x, 0) &= A_0(x), \quad Q(x, 0) = Q_0(x). \end{aligned} \tag{2}$$

Above, the function $A(x, t)$ denotes the cross-sectional area of the flow, which in one-dimensional shallow-water modeling is typically interpreted as a proxy for water depth under suitable geometric assumptions; the quantity $Q(x, t)$ represents the discharge, i.e., the volumetric flow rate; the constant g denotes the gravitational acceleration; the bed elevation is given by $z_b(x)$; the friction term S_f is modeled using Manning's empirical law, where n_M denotes the Manning roughness coefficient; finally, $A_0(x)$ and $Q_0(x)$ specify the initial conditions for the cross-sectional area and discharge, respectively.

This system is an empirical approximation that is frequently used for modeling and simulations. It appears in [22]

Observe that the right-hand side has the nonlinear operator $-gAS_f = -\frac{gn_M^2 Q |Q|}{A^{7/3}}$. This presents significant challenges in deriving any weak form existence result, as any Sobolev space for A and Q will not be closed under these nonlinear operations.

Two-Dimensional Shallow Water System We present an alternative formulation of the Shallow Water equations, extending the consideration to two spatial dimensions. This presentation exhibits more non-smooth maximum operations but fewer degenerate exponents. Consider:

$$\begin{aligned} \frac{\partial u}{\partial t} + u \frac{\partial u}{\partial x} + g \frac{\partial \eta}{\partial x} &= \frac{1}{H(\eta, z_b)} \left(\nu H(\eta, z_b) \frac{\partial u}{\partial x} \right)_x + \frac{\gamma_T u_A - \gamma u}{H(\eta, z_b)} \\ \frac{\partial u}{\partial t} + u \frac{\partial u}{\partial y} + g \frac{\partial \eta}{\partial y} &= \frac{1}{H(\eta, z_b)} \left(\nu H(\eta, z_b) \frac{\partial u}{\partial y} \right)_y + \frac{\gamma_T u_A - \gamma u}{H(\eta, z_b)} \\ \frac{\partial \eta}{\partial t} + u \frac{\partial \eta}{\partial x} + H(\eta, z_b) \frac{\partial u}{\partial x} &= -u \frac{dz_b}{dx} \\ \frac{\partial \eta}{\partial t} + u \frac{\partial \eta}{\partial y} + H(\eta, z_b) \frac{\partial u}{\partial y} &= -u \frac{dz_b}{dy} \\ H(\eta, z_b) &= \max \{0, z_b + \eta\}, \end{aligned} \tag{3}$$

where here u is the velocity, η represents the height of the water level.

2.2 Groundwater flow modeling and control

In this case, the state constraint represents that under-saturation is undesirable, as it clearly leads to underwatering of agricultural production. Over-saturation is also undesirable as potentially damaging to flora; however, it is highly unlikely in an Arid region.

Groundwater dynamics in porous media are governed by diffusion-dominated parabolic partial differential equations derived from Darcy’s law and mass conservation [9]. For a domain $\Omega \subset \mathbb{R}^d$ and hydraulic head $h(x, t)$, the classical form is

$$S_s \frac{\partial h}{\partial t} - \nabla \cdot (K(x) \nabla h) = q(x, t),$$

where S_s is the specific storage coefficient, $K(x)$ is the hydraulic conductivity tensor, and $q(x, t)$ represents external sources or sinks such as pumping wells and recharge. Depending on the hydrogeological setting, one may consider confined flow (linear, single-valued head), unconfined flow with a free surface, or variably saturated flow described by the nonlinear Richards equation [4, 25, 56]:

$$\frac{\partial \theta(h)}{\partial t} - \nabla \cdot [K(h) \nabla (h + z)] = q(x, t),$$

where $\theta(h)$ denotes water content and $K(h)$ the unsaturated conductivity. The Richards equation is strongly nonlinear and degenerate parabolic, posing analytical and numerical challenges when saturation fronts move or dry zones appear.

For the linear confined-flow case, existence, uniqueness, and stability follow from standard variational arguments, while the Richards equation requires monotone operator techniques to ensure existence of weak solutions [4]. From a control perspective, the state equation is well-posed for bounded source controls, but the control-to-state map may lose differentiability near saturation or free-surface transitions, complicating gradient-based optimization. Control variables typically represent pumping or injection rates, recharge operations, or boundary fluxes, and must respect physical and environmental constraints.

Groundwater flow exhibits strong heterogeneity and parameter uncertainty, particularly in the hydraulic conductivity $K(x)$, which may vary over orders of magnitude. This uncertainty motivates stochastic formulations of the control problem [21], which are discussed in detail in Section 2.5 and subsequent sections.

Discretizations typically rely on finite-element or finite-volume schemes that ensure local mass conservation and robustness for heterogeneous $K(x)$ fields. Mixed formulations are often preferred when flux accuracy is critical. For variably saturated flow, nonlinear iterations such as Picard or Newton methods with adaptive time stepping are required to handle the strong nonlinearity in $K(h)$

and $\theta(h)$ [25, 31], while monotone schemes are necessary to ensure physically meaningful solutions.

For optimization and feedback design, adjoint equations for the linear case are classical parabolic PDEs, while nonlinear or unsaturated formulations require consistent linearizations of $K(h)$ and $\theta(h)$, achievable through algorithmic differentiation of the discrete solver. Reduced-order models (ROMs) based on Proper Orthogonal Decomposition (POD) or DEIM can accelerate repeated PDE evaluations [124].

2.3 Relevant Background on PDE-Constrained Optimization

Here, we review some background on PDE-constrained optimization that is particularly applicable to the problem of interest. We focus on deterministic problems, as formally defined above. In the sequel, we present and review the appropriate literature for problems that incorporate uncertainty.

2.3.1 On Control of Hyperbolic PDEs

Optimal control of hyperbolic conservation laws differs fundamentally from the elliptic and parabolic cases, primarily because even entropy solutions of hyperbolic PDEs exhibit finite-speed propagation and naturally develop discontinuities (shocks, rarefactions) in finite time. Even in the case of smooth data and continuous controls (and considering relaxations or switching interpretations of the discrete controls z), the mapping $(v, z) \mapsto u$ is not classically differentiable: the dominant effect of a control perturbation is frequently a shift of shock locations rather than a smooth variation of the state. Hence, the standard control-to-state derivative fails, and sensitivity analysis must be performed in weak or measure-valued topologies. This distinction underlies many of the algorithmic and analytical challenges discussed throughout Section 2, and we therefore focus here on the features that are specific to hyperbolic and parabolic-hyperbolic dynamics.

A general controlled hyperbolic conservation law in conservative form reads

$$\partial_t u(x, t) + \nabla_x \cdot F(u(x, t)) = S(u(x, t), v(t), z(t), x, t), \quad (x, t) \in \Omega \times [0, T],$$

where u is the conserved quantity, $F(u)$ is the nonlinear flux, and $S(u, v, z)$ encodes interior or boundary actuation through the continuous control v and binary control z . The conservative formulation is essential, as it provides the correct weak (entropy) solution concept and the Rankine–Hugoniot relations governing shock propagation.

We write

$$\Omega_T := \Omega \times (0, T) \quad \text{and} \quad L^\infty(\Omega_T)$$

for the space of measurable functions on the space-time cylinder that are essentially bounded, i.e.,

$$L^\infty(\Omega_T) = \left\{ u : \Omega_T \rightarrow \mathbb{R} \mid \operatorname{ess\,sup}_{(x,t) \in \Omega_T} |u(x,t)| < \infty \right\}.$$

Moreover, we write

$$C([0, T]; L^1_{\text{loc}}(\Omega))$$

for the space of functions that are continuous in time with values in $L^1_{\text{loc}}(\Omega)$; that is, for every compact $K \subset \Omega$,

$$\int_K |u(x, t) - u(x, t_0)| \, dx \rightarrow 0 \quad \text{as } t \rightarrow t_0.$$

$L^\infty(0, T)$ let denote the space of essentially bounded measurable functions on $(0, T)$, while

$$L^\infty(0, T; \{0, 1\}) = \{ z : (0, T) \rightarrow \{0, 1\} \text{ measurable} \}$$

denotes the space of measurable on/off (binary) switching controls.

Regularity of States, Controls, and Adjoints. Ulbrich's analysis [119] shows that entropy solutions of controlled hyperbolic conservation laws satisfy

$$u \in L^\infty(\Omega_T) \cap C([0, T]; L^1_{\text{loc}}(\Omega)),$$

and are functions of bounded variation (BV) in space for almost every time. Away from shock curves, the solutions are piecewise smooth, but discontinuities propagate along characteristic directions. Even if (v, z) are essentially bounded,

$$v \in L^\infty(0, T), \quad z \in L^\infty(0, T; \{0, 1\}),$$

the resulting state u depends on (v, z) in a highly nonsmooth way: small perturbations in the controls may lead to finite shifts in the shock locations. Hence, the mapping $(v, z) \mapsto u$ fails to be Gâteaux differentiable.

The adjoint equation corresponding to an optimal control problem with running cost $\ell(u, v, z)$ is a backward hyperbolic transport equation,

$$-\partial_t p - [\nabla_u F(u)]^\top \nabla_x p = \nabla_u S(u, v, z)^\top p + \nabla_u \ell(u, v, z),$$

derived in [119]. Ulbrich shows that the adjoint p inherits the BV regularity of u : it is transported backward along characteristics, remains piecewise smooth, and satisfies compatibility jump conditions across the shocks of the forward solution. These jump conditions are essential for obtaining a consistent sensitivity formula in the presence of discontinuities.

Shift-Differentiability and Sensitivity. Because classical derivatives do not exist, Ulbrich introduced shift-differentiability for entropy solutions [119]. A perturbation of the controls induces smooth variations in regions where u is continuous, and shifts of shock curves weighted by the jump magnitude. This decomposition enables a rigorous directional derivative of the reduced cost functional even when the solution contains shocks. Pfaff and Ulbrich extend this calculus to switched and hybrid controls in [94, 95], where sensitivities also incorporate the influence of switching times and discrete control actions.

Conservative vs. Nonconservative Form. Recall the distinction between a *conservative* and *nonconservative* force-driven evolution process. The latter incorporates forces that are dissipative, e.g., self-friction as well as friction with the bottom surface and air. Thus, a physically faithful approach to appropriately modeling the Shallow Water Equations as studied in hydrological engineering [22] requires consideration of the nonconservative hyperbolic form.

Ulbrich’s framework [119] applies strictly to conservation laws written in conservative form,

$$\partial_t u + \nabla \cdot F(u) = S,$$

because entropy solutions, shock speeds, and the associated adjoint jumps are intrinsically defined in this formulation. In contrast, the nonconservative form

$$\partial_t u + A(u) \nabla u = S$$

is not well defined across discontinuities unless an additional nonconservative product (in the sense of Dal Maso–LeFloch–Murat) is specified. For this reason, sensitivity and adjoint methods for control of hyperbolic PDEs have been formulated and analyzed in the conservative setting. Here, the matrix $A(u)$ denotes the Jacobian of the flux,

$$A(u) := \nabla_u F(u),$$

so that in one spatial dimension the conservative law $\partial_t u + \partial_x F(u) = S$ can be rewritten via the chain rule as the quasilinear form $\partial_t u + A(u) \partial_x u = S$.

Relevance for Hydrodynamic Control. The Saint–Venant (shallow-water) equations used in this work constitute a system of hyperbolic nonlinear PDEs with boundary and internal actuation. Adjoint-based sensitivity methods developed in [95, 119] and the switching-control extensions in [94] can provide a potential tool, or base of a tool, to construct appropriate numerical methods. However, because of the nonconservative features of physical river flow, including self-friction and bottom friction, theoretical guarantees are not available. They can serve as one tool for the multi-scale stochastic control framework developed in the subsequent sections; however, additional workarounds will be needed to consider the unknown topology of the control to state map and the

lack of a dual space that is amenable to numerical approximation, in the case of the complete nonconservative modeling.

2.3.2 On Mixed Integer PDE-Constrained Optimization

We describe the structural aspects of the mixed-integer PDE-constrained control problem arising from the hydropower cascade. The system consists of river segments $u_{(r)}$, $r = 1, \dots, R$, and between two successive segments the dam a is controlled by a set of continuous variables $\mathcal{N}_{a,v} \subset \{1, \dots, n_v\}$ and binary (on/off) switching variables $\mathcal{N}_{a,z} \subset \{1, \dots, n_z\}$, operating a boundary function

$$b_{(l)}(x) = B(v_{(j(l))}, z_{(k(l))}), \quad x \in \partial\Omega_{(l)}, \quad l = 1, \dots, R-1,$$

where $(j(l)) \subset \mathcal{N}_{a,v}$ and $(k(l)) \subset \mathcal{N}_{a,z}$. The resulting optimization problem links discrete operating modes of dams with continuous water-flow dynamics governed by the Saint–Venant or Shallow Water equations.

Stationary Mixed-Integer PDE-Constrained Optimization. Stationary mixed-integer PDE-constrained optimization extends classical finite-dimensional MINLP techniques to steady-state PDE systems of the form

$$e(u) = f(v, z),$$

where u denotes the PDE state, v is a continuous control, and z denotes binary decision variables. Each discrete assignment z requires solving a nonlinear stationary PDE.

Branch-and-bound strategies for stationary PDE problems, such as those developed by Hahn, Leyffer, and Zavala [52], embed relaxed continuous PDE-constrained subproblems at each node of the search tree. While the approach provides global guarantees on moderate-scale PDE models, evaluation of nonlinear PDE solves at every branch presents poor computational scaling.

Penalty and relaxation approaches relax $z \in \{0, 1\}$ to $z \in [0, 1]$ and enforce integrality using concave penalties,

$$\frac{1}{\varepsilon} \sum_i z_i(1 - z_i),$$

as in Garmatter, Porcelli, Rinaldi, and Stoll [40], who show that for sufficiently small ε the penalized problem becomes exact. These continuous relaxations are efficiently solved using matrix-free Newton–Krylov interior-point methods [55], enabling high-resolution PDE models.

Time-Dependent Mixed-Integer PDE-Constrained Optimization.

Time-dependent mixed-integer PDE optimization introduces substantial additional complexity: the binary variables become functions of time, $z = z(t)$, and the PDE evolves according to

$$\partial_t u(t) + \mathcal{A}(u(t)) = \mathcal{F}(v(t), z(t)), \quad t \in (0, T),$$

with discrete controls $z(t) \in \{0, 1\}^{n_z}$ that must be chosen at each time step. The combinatorial structure grows exponentially with the number of time intervals, since a control horizon with N_T time steps yields up to $2^{n_z N_T}$ admissible switching sequences. Moreover, switching can introduce nonsmooth and discontinuous behavior into the PDE solution map $t \mapsto u(t)$, especially for hyperbolic systems such as Saint–Venant.

Classical relaxation–rounding and sum-up rounding techniques developed by Sager [105] and extended in Zeile’s dissertation [126] form the foundation of many modern time-dependent mixed-integer approaches. These methods introduce a relaxed control $\bar{z}(t) \in [0, 1]^{n_z}$, solve the relaxed optimality system, and then construct integer-valued controls $z(t) \in \{0, 1\}^{n_z}$ satisfying integral matching constraints

$$\int_{t_k}^{t_{k+1}} z_i(t) dt = \int_{t_k}^{t_{k+1}} \bar{z}_i(t) dt \quad \forall i,$$

thereby preserving the structure of the relaxed solution.

Decomposition-based methods for dynamic MIPDECO, introduced by Hahn, Kirches, Manns, Sager, and Zeile [51], further reduce complexity by splitting the problem across space and time. Temporal decomposition replaces the full horizon problem with a sequence of smaller subproblems, while spatial decomposition separates integer-dependent and PDE-dependent components. Such techniques greatly improve scalability for high-dimensional, time-dependent systems, and are particularly effective for hyperbolic dynamics such as Saint–Venant flow, where repeated simulation of the PDE is required.

Time-parallelization and parareal-inspired decomposition methods developed by Ulbrich [121] add another layer of structure. By solving coarse and fine PDE propagations in parallel across multiple time slabs, these methods accelerate the evaluation of candidate switching sequences, and enable approximate real-time behavior when embedded inside real-time control frameworks.

Finally, recent developments summarized in Fei, Brady, Larson, Leyffer, and Shen [36] emphasize that dynamic switching control remains computationally demanding, even outside PDE settings. Their observations on the need for relaxations, exact penalties, and structure-preserving rounding techniques extend directly to nonlinear PDE systems. In high-dimensional hyperbolic settings—such as hydropower control—achieving real-time switching performance requires combining these mixed-integer strategies with surrogate modeling, reduced-order modeling, and decomposition.

2.4 Complete Stochastic Dynamic Programming Formulation

Consider the idealized Dynamic Programming Problem under uncertainty for the operation of a hydropower cascade. Decisions include binary $z(t) \in \{0, 1\}$ and continuous $v(t) \in \mathbb{R}^{n_v}$ operation variables for all $t \in [0, T]$. The revenue Π depends on the current and future demand together with the generation partially satisfying the demand from other sources, expressed as a random variable $d(\xi(t))$. In general, since future decisions will be taken upon the realization of the uncertainty at that time, we can write them as a policy given the state.

However, consider that we also have a forecast of the future noise. This comes in the form of a probability density $\tilde{\rho}(\xi(t))$. Naturally, we expect the spread of the density to increase with t .

Consider the probability space $(\Xi, \mathcal{B}, \mathbb{P})$ with each realization a time-dependent stochastic process $\xi(t)$ over $[0, T]$. We will use $\hat{\xi}(t)$ to indicate the live, actual realizations associated with Model Predictive Control.

Generally, the control can be a function of the history of the state, as well as the previous controls up until the current time, as well as the forecasts for the future time:

$$(z^\xi(t), v^\xi(t)) = (z(t), v(t)) \left(\left\{ u(\hat{\xi}(\tau)), z(\hat{\xi}(\tau)), v(\hat{\xi}(\tau)) \right\}_{\tau \in [0, t]}, \{ \tilde{\rho}(\xi(\tau)) \}_{\tau \in (t, T]} \right).$$

Finally, the final Dynamic Programming problem is defined to be:

$$\begin{aligned} \min_{u, v, z} \quad & \int_{t=0}^T \mathcal{R}_\xi \left(-\Pi(u^\xi, d(\xi(t))) + C(v^\xi, z^\xi) \right) dt \\ \text{s.t.} \quad & e(u^\xi(x, t)) = f(v^\xi(t), z^\xi(t)), x \in \Omega, t \in [0, T] \\ & v^\xi \in \mathcal{V}, \\ & S(u^\xi) \geq 0, \text{ almost surely} \end{aligned} \tag{4}$$

for T large, theoretically even $T \rightarrow \infty$, although in practice the propagating uncertainty would make distant times unrealistic to control for.

Finally, $\mathcal{R}$ refers to a coherent risk measure. We note that the simple choice of an expectation (or risk-neutral stochastic optimization) formally satisfies the discussion of the problem while providing flexibility for incorporating, e.g., Conditional Value at Risk, which would in effect minimize the upper tail of the distribution of negative profit.

To solve (4) exactly, one would need to compute the solution of an infinite-dimensional primal-dual stochastic Hamilton-Jacobi-Bellman equation backwards in time given the entire stochastic process up to time t , and identify active and inactive measure and time for the state constraint (see, e.g., [92]).

This is entirely intractable in general for nonlinear problems, let alone ones in function space. Thus, we present an Approximate Dynamic Programming approach in which we attempt to integrate multiple techniques across different time scales. This permits simultaneously taking advantage of potential adaptivity to real-time changes, together with taking advantage of potential HPC-assisted offline solution computation to high precision.

2.5 Background related to Stochastic PDE Dynamic Programming

This subsection reviews methodological foundations relevant to stochastic dynamic programming for systems governed by partial differential equations. We summarize numerical approaches for discretizing uncertainty, algorithmic paradigms for stochastic PDE control, and techniques for handling probabilistic state constraints. The discussion emphasizes challenges specific to nonlinear and hyperbolic PDEs, including limited regularity, discontinuous solutions, and the computational intractability of classical adjoint-based methods. This background provides the context for the hierarchical, value-based control framework developed in the subsequent sections.

2.5.1 Stochastic Discretization

In PDE-constrained stochastic optimization, the propagation of uncertainty through time cannot generally be computed in closed form. Instead, the probability space is discretized through sampling or approximation, enabling numerical evaluation of statistical quantities such as expectations and risk measures. The choice of stochastic discretization depends on the regularity of the parametric solution map $\xi \mapsto u^\xi$, the dimension of the uncertainty, and the available computational resources. Here, ξ denotes the stochastic input parameter and $\tilde{\rho}$ denotes the associated probability measure (or probability density) governing the uncertainty.

Sampling-Based Methods. A classical approach is Monte Carlo (MC) sampling, in which independent realizations $\{\xi^{(i)}\}_{i=1}^M$ are drawn from the probability law $\tilde{\rho}$ and the corresponding PDEs are solved independently. The expectation of a quantity of interest $Q(u^\xi)$ is approximated by

$$\mathbb{E}[Q(u^\xi)] \approx \frac{1}{M} \sum_{i=1}^M Q(u^{\xi^{(i)}}).$$

MC estimators are unbiased and dimension-independent, but converge slowly at a rate $O(M^{-1/2})$. Quasi-Monte Carlo (QMC) methods replace random samples

by low-discrepancy sequences, such as Sobol or Halton points, improving convergence for sufficiently regular integrands. While straightforward and robust, both MC and QMC become computationally expensive when high-fidelity PDE solves are required.

Multilevel Monte Carlo. To mitigate the computational cost of naive sampling, multilevel Monte Carlo (MLMC) exploits a hierarchy of discretizations to reduce variance while controlling bias. Let $u^{\xi,\ell}$ denote the PDE solution at discretization level ℓ , with $\ell = 0$ corresponding to the coarsest level and $\ell = L$ to the finest. Using the telescoping identity

$$\mathbb{E}[Q(u^{\xi,L})] = \mathbb{E}[Q(u^{\xi,0})] + \sum_{\ell=1}^L \mathbb{E}[Q(u^{\xi,\ell}) - Q(u^{\xi,\ell-1})],$$

MLMC allows coarse levels to be sampled extensively at low cost, while fine levels are sampled sparsely. Correlated sampling across adjacent levels ensures variance reduction of the level differences.

Polynomial Approximation Methods. Alternatives to sampling include stochastic collocation and polynomial chaos expansions, which construct polynomial surrogates of the parametric solution map $\xi \mapsto u^{\xi}$. These methods can achieve high-order convergence when u^{ξ} depends smoothly on ξ . However, for nonlinear hyperbolic PDEs and switching dynamics, such as those arising in hydropower systems, parametric regularity is often lost due to shocks or regime changes. In these cases, polynomial methods may become unstable or inefficient.

Relevance for the Hydropower Control Problem. The stochastic inputs in the hydropower cascade induce nonlinear and potentially discontinuous effects in the forward PDE solutions. For this reason, sampling-based methods, and in particular MLMC, provide a robust and scalable approach to uncertainty propagation. In the offline layer, MLMC simulations generate the statistical data required for surrogate training and value approximation. The meso and real-time layers subsequently rely on these discretizations to evaluate stochastic expectations efficiently without imposing strong regularity assumptions on the underlying dynamics.

2.5.2 Stochastic PDE Control

Many applications involve uncertainties, such as uncertain coefficients, unknown boundary conditions, and initial conditions. These problems are commonly formulated as stochastic optimization problems constrained by partial differential equations (PDEs).

Sample Average Approximation (SAA) Sample Average Approximation (SAA) is one of the most widely used approaches for solving such problems numerically. It replaces expectations with empirical averages computed over a set of sampled realizations, resulting in a deterministic approximation that can be solved using standard PDE-constrained optimization techniques.

Depending on the governing PDE, problems can be broadly categorized into elliptic [69, 70, 103], parabolic [16, 20, 50], and hyperbolic or parabolic–hyperbolic systems [35]. While elliptic and parabolic problems generally exhibit smoother solution operators, hyperbolic systems introduce additional challenges due to discontinuities, finite-speed propagation, and strong sensitivity to perturbations, which fundamentally affect both analysis and numerical optimization.

Stochastic Approximation (SA) Stochastic Approximation (SA) methods, such as stochastic gradient descent, have been proposed as an alternative to SAA [85]. These methods rely solely on first-order information and avoid solving large deterministic systems, but they do not exploit the favorable second-order structure typically present in PDE-constrained optimization. As a result, they tend to converge more slowly than SAA-based approaches. Nevertheless, in settings where SAA becomes computationally prohibitive, SA methods may be useful in the offline phase, trading inexpensive iterations for longer overall runtimes.

Thus far¹, hybrid probabilistic models [7] have not demonstrated competitive performance relative to SAA or SA methods in stochastic PDE-constrained optimization.

Adjoint-based Methods. Adjoint-based techniques form the backbone of gradient computations in stochastic PDE control. For each realization ξ , one solves the forward PDE

$$e(u^\xi) = f(v, z, \xi),$$

followed by the corresponding adjoint equation. In the stochastic setting, gradients of expected objectives or risk measures are obtained by aggregating adjoint contributions across samples; see, e.g., [16, 41, 69, 70, 72, 87, 103]. This framework enables efficient computation of $\nabla_v J$ and $\nabla_z J$ in SAA, SA, and risk-averse formulations.

For elliptic and parabolic PDEs, adjoint equations are well posed and the control-to-state map is sufficiently smooth to permit classical sensitivity analysis. In contrast, for hyperbolic or parabolic–hyperbolic systems, adjoint calculus is significantly more delicate. Even in conservative settings, entropy solutions develop shocks and discontinuities, and perturbations in the control typically induce shifts in these discontinuities rather than smooth variations of the state. Consequently, classical differentiability fails, and more refined concepts such

¹By word of mouth in the optimization community.

as shift-differentiability and BV-based adjoint formulations must be employed [94, 95, 119].

Although these ideas are well developed in deterministic settings, their integration with stochastic sampling remains largely unexplored, with only limited results available for stochastic hyperbolic PDE control [35]. For non-conservative systems, the situation is even less understood, with no generally accepted adjoint framework. This motivates the use of primal and derivative-free approaches in parts of the methodology proposed in this work, where adjoint information may be unreliable or unavailable.

2.5.3 On Almost Sure State Constraints

Here the condition

$$S(u^\xi) \geq 0 \quad \text{a.s. in } \xi$$

represents a state constraint regarding water level or flow rate. Such constraints appear naturally in hydrodynamic control for several reasons:

1. physical feasibility, ensuring the computed free-surface remains above the riverbed,
2. operational safety, avoiding flooding or excessive discharge that could damage vessels or infrastructure,
3. environmental protection, guaranteeing ecologically mandated flow and preventing stagnation.

The almost-sure requirement is mathematically challenging because stochastic state constraints define a feasible region

$$\mathcal{F} = \{ (v, z) : S(u^\xi(v, z)) \geq 0 \text{ for a.e. } \xi \},$$

which is typically nonconvex and may have an empty interior. Recent developments provide three principal approaches for handling such constraints: relaxation methods, barrier regularization, and KKT multiplier-based formulations.

Optimality Conditions via Probability-Dependent Multipliers.

Geiersbach and Wollner [42] derive first-order necessary conditions for optimization problems with almost-sure state constraints in Banach spaces. Their analysis introduces a probability-dependent multiplier $\lambda(\xi) \geq 0$ supported on the active set

$$A = \{ \xi : S(u^\xi) = 0 \}.$$

The resulting optimality system involves the variational inequality

$$\mathbb{E} \left[\langle \nabla_u S(u^\xi), \delta u^\xi \rangle \lambda(\xi) + \langle \nabla_{(v,z)} J(u^\xi, v, z), (\delta v, \delta z) \rangle \right] \geq 0.$$

Such multipliers may be singular or highly concentrated, posing severe numerical difficulties in PDE settings.

Relaxation of Almost-Sure Constraints. Kouri, Staudigl, and Surowiec [71] propose a relaxation-based surrogate that replaces

$$S(u^\xi) \geq 0$$

by

$$S(u^\xi) + r_\epsilon(\xi) \geq 0,$$

where the relaxation term $r_\epsilon(\xi) \rightarrow 0$ in probability as $\epsilon \rightarrow 0$. The relaxed problems are differentiable, admit stable adjoints, and converge (in a probabilistic sense) to feasible solutions of the original almost-sure constrained problem. This approach is attractive when the number of uncertainty samples is large.

Barrier Regularization. Barrier methods [107] enforce strict feasibility via the augmented objective

$$\mathcal{J}_{\beta_B}(v, z) = \mathbb{E}[J(u^\xi, v, z)] - \beta_B \mathbb{E}[\log(S(u^\xi))], \quad \beta_B > 0,$$

ensuring $S(u^\xi) > 0$ for all samples. The barrier term preserves differentiability of the reduced problem and is compatible with adjoint-based gradient computation.

Adjoint Equation Under the Barrier Term. For each realization ξ , the adjoint p^ξ associated with the barrier-regularized functional satisfies

$$-e^*(p^\xi) = \nabla_u J(u^\xi, v, z) - \beta_B \frac{\nabla_u S(u^\xi)}{S(u^\xi)}.$$

The barrier forcing term grows as $S(u^\xi) \rightarrow 0$, preventing violation of the constraint and stabilizing the optimization.

KKT Multiplier Structure. As $\beta_B \rightarrow 0$, barrier multipliers converge to the almost-sure multipliers of [42]:

$$\lambda(\xi) = \lim_{\beta_B \rightarrow 0} \left(\frac{\beta_B}{S(u^\xi)} \right), \quad \lambda(\xi) S(u^\xi) = 0.$$

This formal correspondence explains why barrier methods serve as a practical numerical proxy for the theoretically exact but computationally intractable KKT system.

In the hydropower setting, barrier regularization offers the most practical balance between strict feasibility, differentiability, and computational efficiency. Relaxation approaches serve as an effective alternative for large sample sizes, whereas KKT multiplier formulations provide important theoretical insight, but are unsuitable as numerical tools for high-dimensional hyperbolic PDE-constrained stochastic optimization.

2.5.4 Dynamic and Multistage Stochastic Programming with PDEs

The paper [106] considers Stochastic MPC for a Hydropower Cascade, but does not include PDE systems for the water flow as state equations. It presents an efficient distributed scheme for solving the operation of performing MPC with scenarios for the economic operation of the cascade dams.

An important consideration for the operation of MPC is closed-loop stability—that the repeated implementation of the first-time control solution through the operation of MPC yields a trajectory that is physically stable and of approximately optimal performance. We do not attempt to prove closed-loop stability for the system under consideration, but can briefly remark that in the case of hyperbolic systems, the turnpike property, wherein we expect the solution to stay close to an action-minimizing trajectory, is the strategy for establishing stability, e.g., [49].

When the statistical integral $\mathcal{R}$ is a risk measure, rather than expectation, risk-consistency becomes a concern. This problem, discussed in [110, 96], considers the validity of backwards induction in multistage stochastic programming. That is, the risk optimal solution for time stage $t + 1$ may prove suboptimal when the risk metric of stages t and $t + 1$ taken together, in aggregate, is considered. At the time of this writing, this remains a largely open problem in the case of nonlinear problems.

2.5.5 Approximate Dynamic Programming of PDE Systems

The literature on the optimal control of PDEs using Approximate Dynamic Programming (ADP) remains limited. Talaei and co-authors appear to be the primary contributors addressing this class of problems, focusing on ADP-based boundary control for 1D and 2D parabolic equations [112, 113, 114]. The research in [112] considers the control of a coupled semi-linear parabolic PDE system with an unknown nonlinearity under Neumann boundary conditions, where the Hamilton-Jacobi-Bellman (HJB) equation is formulated directly in

the infinite-dimensional space, with the unknown nonlinearity estimated online by a neural network (NN) approximator. The framework was later extended to multi-dimensional nonlinear PDEs [113], specifically for the boundary control of an uncertain 2D Burgers equation. In a subsequent study [114], ADP is used to solve the problem with a linear uncertain parabolic PDE, characterized by representing the value function as a surface integral and a novel NN method with a new weight law.

However, the references above concern deterministic systems. ADP for stochastic PDE systems has not yet been directly applied. We now mention some representative prior literature on ADP for stochastic ordinary differential equations. [19] proposed a new perspective for handling decision-making systems with random parameters, characterized by designing an approach to obtain a family of value functions by solving a family of ODEs. In [98], a differential-neural-network-based min-max robust control method was proposed for nonlinear disturbed systems with uncertain parameters.

3 On the Control to State Map

3.1 State Equation Solution Existence

A major challenge in hydrological modeling that presents limitations for the formulation and development of optimal control tools is the limited understanding of the basic solution and regularity properties for this class of functions. Namely, it has been generally observed that only a restricted set of initial and boundary conditions is consistent with the existence of solutions that can be stably and accurately simulated numerically. However, precise quantitative requirements are absent, and hydrologists either rely on catalogs of known acceptable conditions, or try various conditions until finding one suitable for simulation, or eschew physical modeling altogether and use empirical “HBV” models [13], which, however, generalize poorly across regimes and ambient conditions. For a general survey, see, e.g., [90] and also see [26] for a well-known representative comprehensive case study of hydrological modeling in the San Francisco Bay.

Formally, the nonconservative form indicates that even Young’s measures and statistical solutions are not applicable, and completely Generalized Function Spaces, corresponding to Colombeau algebras, see e.g., [27], become the most appropriate functional analysis toolbox.

Letting $\overline{\Omega} = [0, L] \times [0, T]$ for some $L, T > 0$, consider defining a hyperbolic

system differential operator on $\mathcal{G}(\bar{\Omega})$ by:

$$S_\epsilon U_\epsilon = \left(\frac{\partial A_\epsilon}{\partial t} + \frac{\partial Q_\epsilon}{\partial x} - \sum_{i=1}^K b_\epsilon^i (A_\epsilon)^{n_i^A} (Q_\epsilon)^{n_i^Q} \right) \quad (5)$$

Let us define a corresponding net of operator S_ϵ acting on $U_\epsilon = (A_\epsilon, Q_\epsilon)$ with,

$$S_\epsilon U_\epsilon = \left(\frac{\partial Q_\epsilon}{\partial t} + \frac{\partial(Q_\epsilon^2/A_\epsilon + gA^2/2)}{\partial x} - gA_\epsilon \left(-\frac{\partial z_b}{\partial x} + \frac{n_M^2 Q_\epsilon |Q_\epsilon|}{A_\epsilon^{10/3}} \right) \right) \quad (6)$$

Similarly, we must have for all $\epsilon \in (0, \epsilon_0)$, that $\inf |A_\epsilon(x)| \geq C\epsilon^m$ for some $m \in \mathbb{N}$.

Recall Schauder's Fixed Point Theorem (see, e.g., [125, Theorem 2.A and Corollary 2.13].

Theorem 3.1. *Let X be a Banach space and either*

- *M be a nonempty, closed, bounded and convex subset of X and suppose $T : M \rightarrow M$ is a compact operator, or*
- *M be a nonempty, compact, and convex subset of X and suppose $T : M \rightarrow M$ is a continuous operator,*

then T has a fixed point

Application of this Theorem to establish local existence for nonlinear PDEs by considering locally monotone operators is given in [81].

A Double Loop Fixed Point Scheme In [89], the existence of global Generalized solutions in the Colombeau algebra $\mathcal{G}(\mathbb{R})$ and $\mathcal{G}(\mathbb{R}^2)$ is established. Given that our systems of interest are nonlinear, we will consider a fixed-point iteration wherein a solution is first sought with components of the nonlinear PDE fixed.

Notably, the work in [89] considers the generic PDE:

$$M_1(x, t) \partial_t u + M_2(x, t) \partial_x u = F(x, t, u) \quad (7)$$

and studies existence guarantees under the standing assumption that $\partial_u F$ is bounded for every compact subset of (x, t) . It can be noted that this semilinear form does not fit the structure of the nonlinear PDE systems introduced as the focus of this work. This is what necessitates a double-loop fixed-point scheme to derive the Schauder argument.

In the cases we are considering, we are motivated by applications in design optimization and control with circumstances that necessitate bounds on the solutions. That is, we assume that in accordance with some external necessity, we are only concerned with solutions u satisfying,

$$l_A \leq |A| \leq u_A, \quad l_Q \leq |Q| \leq u_Q \quad (8)$$

in the pointwise almost sure sense. Thus, the field is both lower and upper bounded. This would correspond to a maximal allowable volume and velocity for fluid dynamics, which is reasonable for container limits and stability, as well as a lower bound, which can be of concern for maintenance of certain environmental and ecological properties.

For the first example, consider:

$$\frac{\partial Q}{\partial t} + \frac{\partial(Q^2/A + gA^2/2)}{\partial x} = gA(S_o - S_f(A, Q))$$

and let $Q^{(0)}, A^{(0)}$ be two functions satisfying initial and boundary conditions in (2). Then, by [89], let $Q^{(1)}, A^{(1)}$ be the solution to

$$\frac{\partial Q}{\partial t} + \frac{\partial(Q \frac{Q^{(0)}}{A^{(0)}} + gAA^{(0)}/2)}{\partial x} = gA(S_o - S_f(A^{(0)}, Q^{(0)})).$$

We proceed to consider, next, the solution $Q^{(2)}, A^{(2)}$ to

$$\frac{\partial Q}{\partial t} + \frac{\partial(Q \frac{Q^{(1)}}{A^{(1)}} + gAA^{(1)}/2)}{\partial x} = gA(S_o - S_f(A^{(1)}, Q^{(1)})).$$

and so on.

From (8), we know that there exists a K_l and K_u , as well as $\tilde{S}_l, \tilde{S}_u$, such that,

$$\begin{aligned} K_l &\leq \frac{l_Q}{u_A} \leq \frac{Q^{(0)}}{A^{(0)}} \leq \frac{u_Q}{l_A} \leq K_u \\ K_l &\leq \frac{l_Q}{u_A} \leq \frac{gA^{(0)}}{2} \leq \frac{u_Q}{l_A} \leq K_u \\ \tilde{S}_l &\leq \frac{n_M^2 l_Q^2}{u_A^{10/3}} \leq S_f(A^{(0)}, Q^{(0)}) = \frac{n_M^2 Q^{(0)} |Q^{(0)}|}{(A^{(0)})^{10/3}} \leq \frac{n_M^2 u_Q^2}{l_A^{10/3}} \leq \tilde{S}_u. \end{aligned} \quad (9)$$

We can obtain the following straightforward Theorem:

Theorem 3.2. *There exists some $\bar{T}$, such that for all $T \leq \bar{T}$, there exists a $\bar{B} > 0$ such that there exist $K_l, K_u, \tilde{S}_l, \tilde{S}_u$ such that with initial conditions satisfying (9), the following bound holds:*

$$\max \left\{ \|A\|_{L_\infty}, \|A^{-1}\|_{L_\infty}, \|A\|_{L_\infty}, \|A^{-1}\|_{L_\infty} \right\} \leq \bar{B} \quad (10)$$

for time $0 \leq t \leq T$. As such, the map defined by these iterations is a map between two compact spaces within the set of Colombeau Algebras $\mathcal{C}$, and thus a fixed point exists, which by definition is a solution to the PDE.

Proof. The bound follows immediately by integration with respect to t and x and general compactness and boundedness arguments. The result follows as a direct application of Schauder’s Fixed Point Theorem 3.1. $\square$

3.2 Dual and Adjoint Considerations

Most recently, analysis of optimal control with generalized function spaces has been explored in [68] and [38]. However, these define the Pontryagin principles and the corresponding Euler-Lagrange equations. In order to develop approximate numerical methods, duality will need to be explored. Duality theory for Colombeau algebras is discussed in [39].

More broadly, a Colombeau Algebra is a Banach Algebra, which is a complete normed Banach space (e.g., [111]). While the topological dual (with the weak* topology of convergence) is the set of bounded linear functions $\mathcal{L}(\mathcal{C}, \mathbb{R})$, we can recall from distribution theory that this includes the space of compactly supported analytic functions $\mathbb{C}_c^\infty$. In fact, by [39], this inclusion is a continuous embedding. Although this embedding is not dense, thus making approximation guarantees elusive at the present time, we can consider that a numerical implementation can incorporate multipliers from the space $\mathbb{C}_c^\infty$.

As the space of state solutions $\mathcal{C}$ is a Banach space, we can apply the literature on Optimization in Banach spaces [58]. From this work, we can see that standard arguments for the global convergence properties of the Armijo line search extend from optimization over Hilbert to Banach spaces.

The closest to our work is [121], which explored time-dependent problems and presented a trust region SQP method. The work [93] considers nonlinear hyperbolic conservation laws with switched boundary control, a setting that closely resembles ours. However, in practice, nonconservative shallow water equations are used by hydrologists to model river flow, as dissipative forces, including surface, air, and self-friction, are influential for the river flow dynamics. That is, in many practical river flow models trying to capture rapid variations (hydraulic jumps, rapid transients), practitioners often use non-conservative (also referred to as primitive) formulations of the shallow-water (Saint-Venant in 1d) equations (or mixed/combined formulations), because of their advantages in stability, handling of source terms, computational efficiency or matching observed behavior. See [1, 99, 123] for a discussion of some of the observations of numerical modeling and simulation of different forms of shallow water CFD systems.

4 Three-Time Scale Control - Summary

Notably, the problem described above is far from achieving a practical resolution. There are a number of theoretical and computational hurdles necessary to overcome to approach the problem with some level of clarity and precision.

In this work, we present a general framework for a three-time-scale approach that we believe constitutes a simultaneously necessary and probably sufficient set of operations to satisfactorily perform the control operation of river flow for hydrological engineering.

One of the most salient problems is the very partial existence results, with strict bounds in the Theorem above, presenting fundamental challenges in ensuring the existence of a reliable control to the state map. Because quantitative guarantees of solution existence may not be available in a number of boundary and initial condition regimes, one may have to spend considerable time offline finding the sets of conditions for which control solutions yield PDEs that exist and can be simulated.

To this end, the unlimited volume of computational power that can be theoretically applied in offline operation presents the necessity for defining HPC-aided procedures to solve very large control problems and fine PDE discretizations, and attempt to find a catalog of solutions that exist and direct the optimization criteria to favorable directions. From this catalog, a set of global solutions will be refined across different regimes of real-time exogenous noise as it occurs in practice.

At the meso layer, a set of candidate integer programming (IP) solutions is maintained. Approximate Dynamic Programming will be applied to optimize the continuous variables for problems defined by particles associated with each IP decision set. Heuristics will facilitate the execution of exploration in parallel, as well as potentially try a fresh draw from the solution catalog.

In real-time, Newton algorithms must be applied to be able to perform computation fast enough for real-time realization and modeling. Using techniques in bifurcation theory to handle solution kernel degeneracy, together with recent methods in Real Time Optimization for Mixed Integer Programs, along with Piecewise Newton Methods, one can define the appropriate toolkit to perform real-time control. The real-time formulations will arise from problems defined by the meso layer serving as a tracking term in a convex functional for the fast MPC problem. Moreover, the models used in real-time are Model-Order Reduced from the more precise PDE simulations in the meso layer.

Integration of the algorithms across the three layers presents a challenging and important mathematical and scientific computing exercise. We summarize the approach below, defining the properties of the three layers and their interfacing.

Recall that, operationally, in the motivating application of dam networks dis-

cussed in the Introduction, these components correspond respectively to planning, supervisory control, and real-time actuation tasks. Accordingly, the proposed framework is structured into the following three tiers:

1. **Offline:** This constitutes the running of computational algorithms without consideration of any contemporaneous operation, i.e., any cost-sensible and feasible running of computation. This permits the use of computing clusters with significant HPC resources to run over a long time period. Thus, this setting encourages formulations and discretizations that push the level of dimensionality of the problem to otherwise intractable orders of magnitude. This operation can be performed at the beginning to generate a catalog of a large set of feasible, sensible, stable, and well-performing solutions, with the attempt to arrive at, or get close to, global optimality for some formulations. Subsequently, this catalog will serve as a reliable reference that can feed the operation of meso-layered approximate dynamic programming. Offline solutions can be recommended through AI training and can serve as both potential solution paths to explore or a terminal state and/or value function for Model Predictive Control or Approximate Dynamic Programming, respectively. The offline computation can be periodically run to steadily accumulate more refined and better solutions, while taking feedback from the data observed in real-world operation. Table 1 describes the central features of this layer.

Property	Description	Challenge Addressed
Role	Find regions of control-to-state existence	Lack of a priori knowledge
Output	Construct Catalog of Solution Candidates	Intractability for low latency solvers
Hardware	HPC at Maximal Resource Utilization	Nonlinear High dimensional System
Latency	Unlimited	Extended computation time
Dimensionality	Millions	Complex system that requires high fidelity for accurate and precise modeling
Numerical Tools	CFD numerical solutions	
Data Driven Tools	Reinforcement Learning	Lack of a priori solution existence guarantees

Table 1: Operational purpose and resource requirements for Offline Computation

2. **Meso-Temporal:** The meso layer represents a reliable computational bridge between the offline and the real-time layer. While not pushing the

dimensionality to the extreme degree of the Offline layer, it still retains the full model faithfulness as far as incorporating the PDE formulation. It involves the steady evolution of a moderate number of particles representing potential solutions, and applies procedures that can be solved using Nonlinear Programming software. The offline catalog will feed candidates to the set of particles in a deliberate and informed manner, while the best (appropriately defined) particles will be the source of Model Order Reduction to define the Real-Time operation. The meso layer is meant to take advantage of the comprehensive toolkit available for solving nonlinear constrained optimization problems as commonly used for nonlinear model predictive control. These tools are meant to handle problems with amenable regularity - hence the necessity of the offline catalog to identify appropriate regimes, while also incorporate relatively high dimensional nonlinear models - which would be too difficult to solve in real time latency.

Property	Description	Challenge Addressed
Role	Perform reliable nonlinear numerical optimization	Complex nonlinear models
Output	Optimal and near optimal solution estimates for operation	Reliable model predictive control with appropriate forward horizon
Hardware	High end computing resources devoted entirely for this operation	Leveraging robust reliable numerical optimization solvers
Latency	0.1-2 hours	Nonlinear high dimensional problem with adaptation to changing environment
Dimensionality	Tens of Thousands	Balance of fidelity and tractability
Numerical Tools	Nonlinear Programming solvers	Reasonable model fidelity for optimal solution generation
Data Driven Tools	Adaptive tuning and particle management	Mixed binary-continuous problem and potentially disparate regimes

Table 2: Operational purpose and resource requirements for Meso-scale Computation

3. **Real-Time:** The real-time layer corresponds to the actual implemented control in the fast operation of the engineering system. It will use state-of-the-art fast Newton-based solvers, extended recently to problems with mixed-integer decisions. Linearized problems that approximate the cur-

rent state locally are fed from the meso layer, and the computation is performed locally on safe, fast-performing control code. Feedback from the real evolving state and any additional information will be continuously sent to the meso layer, closing the end-to-end operation of the entire algorithmic structure.

Property	Description	Challenge Addressed
Role	Choose real-time adaptive control law	Fast decisions for desired latency for power grid resilience (e.g. ancillary services)
Output	Controls to implement	Physical Operation
Hardware	Embedded controller	Fast desired latency with proximity to decision system
Latency	Seconds to Minutes	Rapid adaptation to requirements and new solution estimates
Dimensionality	Around 100	Tractability
Numerical Tools	Real Time Iteration and MIP solvers	Warm starting and optimizing reduced order models
Data Driven Tools	Solution and sensor information sent asynchronously up to the Meso hardware	Long run tuning of the architecture

Table 3: Operational purpose and resource requirements for Real-time Computation

Figure 1 presents the overall structure and flow of the computational architecture, together with the communication between the layers.




Figure 1: Three time scale computational system architecture for real time control of complex nonlinear dynamics

5 Offline Computation of a Catalog of Reference Solutions

The goal of the offline computation phase is to apply brute force HPC computing power in order to obtain very strong and precise results. In this Section, we describe methods that attempt to solve the Dynamic Programming through a variety of approximating procedures and methods of solution computation. In the context of dam network operation, these problems correspond, for example, to the determination of admissible gate operation policies and flow regulation strategies under hydraulic and operational constraints. We can note immediately that the exact solution of (4) by way of defining and solving a Hamilton-Jacobi-Bellman equation is intractable, on account of the infinite dimensionality of the PDE system together with the presence of control constraints. Thus, we present a set of potential methods often used to solve Dynamic Programming approximately that are amenable to discretization and computation.

Before describing the individual components, we summarize the role of each offline method in Table 4. The purpose of Section 5 is not to propose each listed method as a standalone solver for the full stochastic mixed-integer PDE-constrained control problem. Rather, the offline layer combines numerical optimization, reduced switching-structure search, and data-driven learning to construct a catalog of feasible and high-quality reference trajectories. These catalog entries are later used for meso-layer initialization, terminal-value approximation,

feasibility guidance, and real-time warm starts.

Component	Purpose in the offline layer	Type	Techniques
Multistage stochastic mixed integer optimization	Approximate the stochastic dynamic programming problem and generate reference trajectories under uncertainty	Numerical	SAA, (period) ergodic B&B heuristic search
Switching-structure reduction	Replace the full binary trajectory with lower-dimensional switching variables (counts $\underline{z}$ and times $\bar{z}$)	Numerical	Bilevel and Black-Box Optimization
Offline learning	Learn feasible regions, policy recommendations π_ω , and continuous-control bucket classifier ψ_ω from high-fidelity simulations	Data-driven	Reinforcement & Supervised Learning with latent attraction domains $\mathcal{D}(v^*)$
Catalog construction	Reference planning catalog $\mathcal{C}_{\text{off}}$ for for meso-layer warm starts and value approximation	Hybrid	Supervised Learning HPC

Table 4: Summary of offline-layer methods and their roles in constructing the reference catalog.

5.1 Formulations

5.1.1 Multistage Stochastic Programming

Multistage Stochastic Programming (MSP) involves finding the solution of an optimization problem over a discrete-time stochastic process for a finite number of steps. One discretizes the stochasticity, typically through scenario generation. For a comprehensive monograph of scenario generation techniques applicable to the primary case study of interest, see, e.g., [28]. A scenario is a sequence of samples $\{\xi_i(1), \xi_i(2), \dots, \xi_i(T)\}$ for the realization of the noise over the T subsequent time steps. Typically, the scenario is an expanding tree, i.e., for $\xi_{i,1}$ one recursively branches to construct subsequent realizations

$\{\xi_{i+1}(1), \xi_{i+1}(2), \dots, \xi_{i+1}(T)\} = \{\xi_i(1), \xi_{i+1}(2), \dots, \xi_{i+1}(T)\}$. Thus, the infamous *curse of dimensionality* applies relative to the time horizon T . However, at this offline layer of solution generation, it is expected that the computational load of an HPC server will be pushed to capacity as far as the longest time horizon and most precise stochastic discretization possible.

Generically, we can construct an empirical sample average approximation to the risk functional $\mathcal{R}$; however, for simplicity of notation, assume that $\mathcal{R} = \mathbb{E}$, permitting the use of a simple sum expression.

The MSP problem with discretized time and stochasticity now reads as,

$$\begin{aligned}
\min_{u, v, z} \quad & \frac{1}{NT} \sum_{i=1}^N \sum_{t=0}^T [-\Pi(u^{\xi_i}(t), d(\xi_i(t))) + C(v^{\xi_i}(t), z^{\xi_i}(t))] \\
\text{s.t.} \quad & \tilde{e}(u^{\xi_i}(x, t), u^{\xi_i}(x, t+1)) = \tilde{f}(v^{\xi_i}(t), z^{\xi_i}(t)), \quad t \in [T], \quad \forall i \\
& v^{\xi_i}(t) \in \mathcal{V}, \quad t \in [T], \quad \forall i \\
& S(u^{\xi_i}(t)) \geq 0, \quad t \in [T], \quad \forall i \\
& u^{\xi_i}(0) = u_0, \quad \forall i \\
& \begin{cases} u^{\xi_i}(t) = u^{\xi_j}(t), \quad v^{\xi_i}(t) = v^{\xi_j}(t), \quad z^{\xi_i}(t) = z^{\xi_j}(t), \quad \forall t, i, j \\ \forall t' < t \text{ s.t. } \xi_i(t') = \xi_j(t') \end{cases}
\end{aligned} \tag{11}$$

observe that the decisions and state are parametrized by both time and stochastic realization. This models the fact that a decision is made at time t . The last constraint, defined as the non-anticipativity constraint, requires non-clarvoyance, i.e., a decision (as well as state) at time t cannot depend on noise realizations after t .

The state constraint is transformed into a constraint that holds for all samples. Of course, this can guarantee only some probability of the almost sure constraint holding. In the meso layer we instead use a barrier function to enforce the state constraint, whereas here, we can consider that the stochastic discretization can be fine enough such that, with a small robustness factor distancing $S(\cdot)$ from the exact break down value of desired/safe operation, we can more or less presume the almost sure state constraint will be satisfied.

Observe, however, that for the problem to be well-defined, an initial state u_0 must be set. Of course, in the general hydropower cascade operation, there is no offline initial state. Thus, one can simply take statistics of the history of the states through observations, and devote resources of computation towards solving (11) as according to the distribution of various states being set as u_0 .

5.1.2 Ergodic Optimization

In the case of ergodic optimization, one considers a stationary distribution that optimizes some long-run average over the history of the states. See, e.g., [61] for a generic exposition. Notably, the exogenous noise in the applications of interest,

including the standing example in this paper, can be considered to be multi-periodic. That is, precipitation and electricity demand tend to exhibit cyclical behavior, with periods spanning twenty-four hours, one week, and one year. Approaches towards explicitly considering this process structure are considered in [73]. Dynamic Programming, and a computational approach using Stochastic Dual Dynamic Programming, for linear dynamics is considered in [109].

In this case, problem (4) can be augmented with an additional set of periodicity constraints regarding the probability law of the solution:

$$\mathcal{P}(u^\xi(x, t)) = \mathcal{P}(u^\xi(x, t + p)), \forall t$$

for a single period p . With the multi-periodic, e.g., $\{p_1, p_2, p_3\}$, formalization, one can decompose the solution:

$$u^\xi(x, t) = u^{\xi,1}(x, t) + u^{\xi,2}(x, t) + u^{\xi,3}(x, t)$$

with

$$\begin{aligned} \mathcal{P}(u^{\xi,1}(x, t)) &= \mathcal{P}(u^{\xi,1}(x, t + p_1)), \mathcal{P}(u^{\xi,2}(x, t)) = \mathcal{P}(u^{\xi,2}(x, t + p_2)), \\ \mathcal{P}(u^{\xi,3}(x, t)) &= \mathcal{P}(u^{\xi,3}(x, t + p_3)), \forall t. \end{aligned}$$

This formalism permits stochastic discretization techniques incorporating spectral and wavelet bases.

5.1.3 Hierarchical Discrete-Continuous Formulation

In this section, we propose an alternative formulation of the decision variables that leverages the opportunity for HPC-enabled exhaustive computation for finding a global solution, or one that is approximately global.

In the procedures above, the continuous and integer decision variables are treated to exhibit optionality at every time instant. That is, an operation of a dam could, if it were to be optimal to a criterion, generate frequent alternation between an open and a closed state. Of course, the operation of opening or closing a dam is an engineering task that incurs a cost. In other works on integer-valued time-parametrized control, e.g., [79], total variation regularization is used to minimize the switching frequency. For the problems of interest in this paper, we consider that we can evaluate a real economic cost to the operation, and thus the regularization is a real economic quantity, rather than a computational tool.

We can go one step further, though, as far as considering an inductive bias towards less frequent changes in the binary and integer variables $z(t)$. To this

end, we replace this decision variable with an equivalent depiction:

$$\begin{aligned} \underline{z} \in \mathbb{N}^{n_b} &= \left\{ \underline{z}_i := \|z_i(t)\|_{TV} := \left| \left\{ \underline{t} : \lim_{t \rightarrow \underline{t}^-} z(t) \neq \lim_{t \rightarrow \underline{t}^+} z(t) \right\} \right| \right\} \\ \bar{z}(\underline{z}) \in \mathbb{R}^{n_b \times \|\underline{z}\|_1} &:= \left\{ \bar{z}_{i,j} = \underline{t} : \begin{array}{l} t > \bar{z}_{i,j-1}, \lim_{t \rightarrow \underline{t}^-} z(t) \neq \lim_{t \rightarrow \underline{t}^+} z(t), \\ \bar{z}_{i,0} = 0, j \leq \underline{z}_i. \end{array} \right\} \end{aligned} \quad (12)$$

That is, decompose the problem into 1) the *number of switches* $\underline{z}$ and 2) *how much time to let elapse before a switch* $\bar{z}$. Observe that the dimensionality of the latter quantity depends on the value of the former.

This presents a natural hierarchical optimization procedure, whose formulation we simplify for brevity,

$$\begin{aligned} \min_{\underline{z}} \quad & \hat{\mathcal{J}}(u(\underline{z}), v(\underline{z}), \bar{z}(\underline{z})), \\ \text{s.t.} \quad & v(\underline{z}) \in \underset{v}{\operatorname{argmin}} \mathcal{J}(u, v, \bar{z}(\underline{z})), \end{aligned} \quad (13)$$

i.e., find the optimal number of switches for all dams across the total operation horizon, such that the time between switches is chosen in conjunction with the continuous control variables to maximize the total (expected or negative risk) profit. Each evaluation of the integer problem triggers a continuous variate optimization problem. This facilitates the use of off-the-shelf solvers, as constrained continuous optimization tools can be employed as simulation code from a function evaluation. The integer programs, which have a nonlinear black box criterion in the form of this evaluation, can then be solved using the usual branch-and-bound, heuristic, or Bayesian approaches.

5.2 Numerical Mixed Integer Optimization Methods

The offline layer has unrestricted access to the full PDE solver, enabling the use of computationally intensive mixed-integer approaches. In this environment, we can run Branch-and-Bound (B&B), greedy and heuristic search strategies, and Black-Box or Derivative-Free Optimization (BBO/DFO) methods. These approaches explore the discrete decision space associated with dam operating modes and switching configurations. Because offline computation is unconstrained by real-time requirements, the objective is not merely to find a single optimal solution but to assemble a diverse, high-quality catalog of reference trajectories that span relevant hydrological and operational regimes.

5.2.1 B&B

Classical B&B dates back to Land and Doig [77]; see also Mitten [88] for a general formulation and properties. B&B is used in the offline layer as a structured

mixed-integer search mechanism, not as a new standalone algorithmic contribution. Its role is to explore moderate-dimensional subproblems in which some binary switching variables are fixed or relaxed, and to generate high-quality incumbent switching patterns for the reference catalog. At a node of the search tree, a relaxed PDE-constrained subproblem may be solved with $z \in [0, 1]$ or with a subset of binary variables fixed. Nodes whose relaxation lower bound cannot improve the incumbent are discarded, while promising nodes are further branched.

Because the number of switching configurations grows exponentially with the number of binary variables and time stages, exhaustive B&B is not intended for the full stochastic PDE-constrained control problem. This restricted use is consistent with the MIPDECO literature: direct discretization of PDE constraints can produce MILPs whose size makes standard branch-and-bound or branch-and-cut solvers ineffective without additional structure, preprocessing, or constraint management [46]. Related work also emphasizes the use of state elimination and problem-structure reduction for mixed-integer PDE-constrained optimal control [116], as well as lower bounds and convex relaxations for mixed-integer PDE-constrained problems as ingredients for global optimization and branch-and-bound-type procedures [80]. In the proposed architecture, B&B is therefore used offline on restricted switching structures, shortened horizons, or reduced discrete subspaces. Even when global optimality is not certified, the procedure can provide structured candidate trajectories that enrich the offline catalog and seed the meso-scale particle layer.

5.2.2 Other Offline Mixed-Integer Strategies

Greedy and heuristic search methods provide simple yet effective strategies for exploring mixed-integer decision spaces. These approaches iteratively improve a candidate solution through local modifications [62] or rule-based adjustments, including hill climbing [104], tabu search [43, 44, 45], and a broad family of metaheuristics [115]. Unlike globally oriented methods such as Branch-and-Bound, greedy and heuristic schemes prioritize speed and coverage over optimality. In the offline setting, they complement B&B by rapidly generating diverse feasible switching configurations through local search, greedy selection of switching variables, and trajectory-level perturbations guided by PDE-based cost evaluations. Their primary role is to populate the reference catalog with structurally varied high-quality solutions, rather than to provide global guarantees, thereby seeding the meso-scale evolutionary layer with a broad range of integer patterns.

Black Box Optimization methods form a broad family of algorithms designed for settings in which a function evaluation constitutes an expensive simulation. Included are Derivative-Free Optimization (DFO) methods for when gradients are unavailable, unreliable, or prohibitively expensive. Recall that in

the case study of interest, a priori reliable numerical adjoint calculation cannot be assured. Thus DFO methods can be used to poll and ultimately assess the existence and regularity of potential reference solutions. Classical DFO approaches include pattern search [117], mesh-adaptive direct search (MADS) [5], line-search variants [82, 64], trust-region models [29, 66], and surrogate-assisted schemes [102]. Black-box and population-based methods such as genetic algorithms [59], mixed-variable variants of CMA-ES [54, 65], and particle-swarm optimization [63] can be useful for exploratory search over nonsmooth or simulator-defined objective functions when the search variables are reduced or parameterized.

In the present offline PDE-constrained setting, DFO/BBO methods are included as auxiliary tools for reduced or parameterized search spaces, and also as possible initial diagnostic tools when the simulator is available only as a black box. They are not intended to solve the full space–time discretized PDE-constrained mixed-integer control problem directly. A direct application of generic DFO/BBO to the full control trajectory (v, z) would be computationally inappropriate, even in the offline layer, because the number of required high-fidelity PDE evaluations would scale poorly with the dimension of the control trajectory.

Instead, their main role is limited to low-dimensional supervisory or switching-structure variables arising in the offline catalog-construction phase. For example, after the binary control has been parameterized by the number of switches $\underline{z}$ and the switching times $\bar{z}$, DFO/BBO methods may be used to explore switching counts, switching-time patterns, policy parameters, or catalog-selection variables. In this case, each black-box evaluation may call a high-fidelity PDE simulation or a structured continuous PDE-constrained optimization subproblem, but the variables optimized by DFO/BBO remain low-dimensional.

A second restricted use case occurs at the beginning of offline exploration, when no reliable adjoint, differentiability information, relaxation, or exploitable PDE structure is yet available. In this case, DFO/BBO may be used as an initial black-box diagnostic step to sample a limited number of candidate controls or parameterized policies, identify feasible operating regimes, detect solver-failure regions, reveal promising switching structures, and estimate useful scalings or bounds for later optimization. Once such information has been obtained, the search should be transferred to reduced, parameterized, or structure-exploiting PDE-constrained optimization methods rather than continued as an unrestricted black-box search over the full trajectory.

This restricted use is particularly relevant when adjoints are unavailable, unreliable, or not meaningful because of nonsmooth switching behavior, wetting/drying transitions, solver failures, or changes in the numerical PDE solution regime. Thus, DFO/BBO contributes to catalog diversification, feasibility-region exploration, and early black-box diagnosis, but it is not used to provide global optimality certificates or to replace structure-exploiting PDE-constrained optimization methods.

5.2.3 Comparison of Offline Mixed-Integer Methods for PDE-Constrained Optimization

The three offline strategies play different roles in constructing the reference catalog. B&B is the most systematic approach and is appropriate when the number of active binary variables is moderate and useful relaxations are available. In such cases, it can provide high-quality incumbent solutions and, in favorable settings, useful bound information. Its limitation is the exponential growth of the search tree, especially when binary decisions are repeated over many time stages or many dams.

Greedy and heuristic search methods have the opposite profile. They are inexpensive and useful for rapidly producing many feasible switching configurations, but they provide no optimality certificate and may become trapped in local patterns. Their primary value in the offline layer is coverage: they generate diverse switching structures and candidate trajectories that can later be refined, filtered, or used to initialize meso-scale computations.

DFO/BBO methods have two restricted roles. First, they may be used at the beginning of offline exploration as black-box diagnostic tools when no reliable adjoint, differentiability information, relaxation, or exploitable PDE structure is yet available. In this role, they sample a limited number of candidate controls or parameterized policies to identify feasible operating regimes, solver-failure regions, promising switching structures, and useful scalings or bounds for later optimization. Second, after the search space has been reduced or parameterized, they may be used for low-dimensional switching-time optimization, exploration over the number of switches, policy parameters, reduced supervisory variables, and catalog-selection variables. They are not intended as direct solvers over the full space-time discretization of the PDE control problem. In these restricted roles, DFO/BBO methods can contribute robustness to nonsmooth simulator behavior and help enrich the catalog in regimes where derivative-based information is unavailable or unreliable.

Thus, the offline methods are complementary not because they solve the same problem in different ways, but because they populate the catalog through different mechanisms: B&B supplies structured mixed-integer search, greedy and heuristic methods supply rapid coverage, and DFO/BBO supplies early black-box diagnosis together with reduced-dimensional exploration of switching structures and simulator-defined regimes.

5.3 Learning-based Approaches

This subsection presents the learning mechanisms at the offline layer of the three-tier control architecture. At this level, the complete high-fidelity CFD solver is available and used directly to explore the feasible control space and

learn probabilistic and policy structures for the PDE-constrained stochastic optimization problem. No model-order reduction or surrogate approximation is performed here; those tasks belong to the meso and real-time layers discussed later.

Role within the Three-Tier Control Principle. The offline layer serves as the computational foundation for the entire learning-control hierarchy. It (i) collects data from full CFD evaluations to characterize feasible and infeasible control configurations; (ii) trains classifiers that predict the existence of a stable numerical PDE solution; (iii) applies reinforcement learning to identify optimal discrete-continuous control combinations using CFD-based feedback; and (iv) estimates domains of attraction for continuous variables, enabling later discretization and policy transfer to the meso scale. In addition, the learned policy π_{ϕ^*} also serves as a recommender of promising offline catalog entries, guiding the meso layer toward high-value candidate trajectories that the learner has found to perform well under upstream stochastic conditions.

Offline Learning Objectives. Given the stochastic PDE system

$$e(u^\xi) = f(v^\xi, z^\xi), \quad \xi \in \Xi,$$

where v denotes continuous and z discrete control components, the offline learning process pursues three principal objectives, described below.

5.3.1 Predicting Feasibility via Supervised Learning

The first objective is to construct a classifier

$$\psi_\omega : (v, z, \xi) \mapsto [0, 1]$$

that approximates the probability of a stable PDE solution for each control-uncertainty triplet. For every CFD evaluation, a binary label $\chi_i \in \{0, 1\}$ records solver convergence or failure. The resulting model ψ_ω provides a feasibility map that restricts later search and prevents infeasible simulations.

5.3.2 Optimizing Discrete-Continuous Controls via Reinforcement Learning

By using PDE simulations as the model in a reinforcement learning (RL) loop, we can aim to perform RL with the actions set as the decision variables and the computed objective as the reward. This requires parametrizing the decisions and rewards with deep neural networks, in the case of Deep Reinforcement Learning (DRL).

The second objective formulates the stochastic control problem as an RL environment in which each CFD evaluation corresponds to an episode. For each feasible configuration (v, z) , the observed reward equals the negative risk-adjusted objective

$$J(v, z) = \mathbb{E}_{\xi} [\mathcal{R}_{\xi}(-\Pi(u^{\xi}(v, z), d(\xi)) + C(v, z))] .$$

Two RL formulations are supported at the offline layer:

- (a) **Bucketed-continuous formulation:** the continuous controls v are discretized into coarse buckets, inducing a finite action set and converting the RL problem to a discrete-action MDP.
- (b) **Discrete-action formulation:** the RL agent outputs only the discrete controls z , while the continuous component v is obtained by solving the continuous PDE control subproblem with z fixed. This matches the logic of the real-time layer.

The policy π_{ϕ} maps current estimates of the system or uncertainty to control actions (v, z) (or just z in the discrete-action formulation). RL updates improve π_{ϕ} over repeated high-fidelity evaluations, effectively learning decision rules for the mixed-integer PDE-constrained control problem.

5.3.3 Estimating Domains of Attraction for Continuous Controls

For each local minimizer v^* obtained by PDE optimization, nearby sampled controls are clustered according to which minimizer they converge to. This yields characteristic regions $\mathcal{D}(v^*) \subset \mathcal{V}$ representing domains of attraction.

These regions play two crucial roles: (1) they define coarse discrete lattices for continuous-control discretization at the meso scale, and (2) they provide an optimal bucketing scheme for the RL formulation in which continuous variables are discretized. Each domain of attraction corresponds to a stable, behaviorally coherent bucket in the action space, leading to more effective RL training and more robust policy generalization. Both conceptually and operationally, they can be considered as latent variables that assist the overall solution catalog identification pipeline.

Offline Data Generation. A comprehensive dataset is assembled by sampling control configurations and stochastic realizations

$$\mathcal{D}_{\text{off}} = \{(v_i, z_i, \xi_i, u_i, J_i, \chi_i)\}_{i=1}^{N_{\text{off}}} ,$$

where χ_i marks CFD feasibility, and J_i is the computed objective. The dataset supports both supervised learning for ψ_{ω} and policy-gradient or Q-learning updates for π_{ϕ} . Since full CFD runs are used, the offline layer provides a high-fidelity basis for all subsequent surrogate construction.

Offline Reinforcement Procedure. Algorithm 1 summarizes the offline RL workflow consistent with the paper’s algorithmic structure. In **(S0₁)**, the policy π_{ϕ_0} , the feasibility model ψ_{ω_0} , and the dataset $\mathcal{D}_{\text{off}}$ are initialized. Each episode begins with sampling $(v^{(k)}, z^{(k)})$ and $\xi^{(k)}$, followed by running a high-fidelity CFD simulation. If the simulation converges, a risk-adjusted objective is computed, and the policy is updated; if it diverges, the feasibility model is updated instead. After K_{off} episodes, both models are retrained on the accumulated dataset to obtain the final offline policy π_{ϕ^*} and feasibility map ψ_{ω^*} .

Algorithm 1 Offline RL with High-Fidelity CFD Evaluations

Require: CFD solver for $e(u) = f(v, z)$; initial policy parameters ϕ_0 ; feasibility model ψ_{ω_0} ; learning rates $\eta_{\phi}, \eta_{\omega} > 0$; number of episodes K_{off}

Ensure: Trained policy π_{ϕ^*} and feasibility predictor ψ_{ω^*}

(S0₁) Initialize π_{ϕ_0} and ψ_{ω_0} ; set dataset $\mathcal{D}_{\text{off}} = \emptyset$.

for $k = 0, 1, \dots, K_{\text{off}} - 1$ **do**

(S1₁) Sample candidate controls $(v^{(k)}, z^{(k)}) \sim \pi_{\phi_k}$ and stochastic realization $\xi^{(k)} \sim \tilde{\rho}(\xi)$.

(S2₁) Run high-fidelity CFD simulation $e(u^{\xi^{(k)}}) = f(v^{(k)}, z^{(k)})$.

(S3₁) Determine feasibility $\chi^{(k)} \in \{0, 1\}$.

if $\chi^{(k)} = 1$ **then**

(S4₁) Compute risk-adjusted objective $J^{(k)}$.

(S5₁) Update policy: $\phi_{k+1} = \phi_k - \eta_{\phi} \nabla_{\phi} J^{(k)}$.

(S6₁) Append $(v^{(k)}, z^{(k)}, \xi^{(k)}, J^{(k)}, \chi^{(k)})$ to $\mathcal{D}_{\text{off}}$.

else

(S7₁) Update feasibility predictor.

end if

end for

(S8₁) Train final models π_{ϕ^*} and ψ_{ω^*} on $\mathcal{D}_{\text{off}}$.

Offline Outputs and Interface. The offline stage provides: (i) the feasibility map ψ_{ω^*} defining regions of numerical solvability; (ii) a control policy π_{ϕ^*} together with a curated catalog of CFD-validated control solutions that the learner recommends for meso-layer initialization and replacement; and (iii) statistical information on attraction domains in $\mathcal{V}$, which define optimal bucketing schemes for continuous-action discretization in RL and the meso-scale surrogate model. These components form the input knowledge base for surrogate management at the meso layer and feedback control at the real-time layer.

5.4 Catalog Construction

The overall purpose of the offline phase of the computational control operation is to establish a catalog of candidate solutions. By leveraging the resources, in

time and computation, that could be applied to solve the problems formulated in this Section, the limits of scale, accuracy, and precision could be reached, as far as available methods and hardware.

Note that the *initial* state and decision variable regime are inputs to all of the problems above, and the algorithms are meant to define the optimal sequence of decisions from that state in some forward horizon. Of course, in real-time operation, there is no initial state; the state continuously changes. And so the practical engineering initiation of Hydropower control must be agnostic with respect to this initial state.

Thus, the appropriate course of action is to construct the catalog of solutions across a variety of initial states $u(x, t_0)$ and control orientations $\{z(t_0), v(x, t_0)\}$. Their choice can be defined based on data observed in the actual operation of the hydropower cascade. Through sensors measuring water height and velocity in river banks, while defining a database of historical control actions, a large dataset of initial conditions can be defined.

Then, which initial conditions to use can be chosen through:

- Sampling, with importance/Thompson sampling used to target rare events of critical importance to Water Engineering authorities
- Lattice grid construction across the discretized dimensionality of solutions, carefully tuning to the fact that this may differ across solution forms.

The database of offline solutions can then serve as a useful tool for investigating particular scenarios of interest during the course of hydrocascade operation. Meanwhile, development can proceed as far as formulating, analyzing, and implementing the methods described below in the meso and real-time layers, while developing tools to integrate the offline solution catalog to be fed to the meso layer.

The critical importance of the component of using ML to learn the domain of the feasible control-to-state map must be emphasized. Given the lack of a priori knowledge and the lack of the availability of this information for the Shallow Water Equations, finding the appropriate set of control valued in some connected subdomain on which to perform optimization is critical. The use of ML permits navigating around the lack of a priori knowledge in this case, presenting a novel synergy of learning with operations research.

Given the nonconvexities in the optimization problem as well as the lack of a priori global regularity guarantees, the catalog of offline solutions will serve as valuable warm starts to the faster, “meso” level operation. By biasing warm starts to solutions that are in the interior of a region identified to exhibit a reliable control-to-state map, this provides a landscape over which optimization can take place without encountering unforeseen inoperability.

5.5 Open Challenges and Future Research

A number of novel methods have been formulated in this Section - specifically the branching-switching formulation and domain-of-attraction learning. As these are entirely new in the literature, to the best of the authors' understanding, research work formalizing them, studying their theoretical learning, convergence and robustness properties, and validating their performance numerically should be performed before including them into a complete water engineering operational pipeline. Other methods, including learning feasible and regular control-to-state regions of the control domain and the use of mixed-integer optimization tools, are applications of existing methods towards a significantly higher dimensional and nonlinear problem relative to conventional baselines. This stress-testing should be investigated numerically and, quite possibly, require additional algorithmic advances in the techniques recommended.

6 Meso Sequential Numerical Optimization

Nonlinear Model Predictive Control (MPC) (e.g. [3]) is a lookahead approach towards computationally solving Dynamic Programming by repeatedly solving an optimal control problem over a succeeding T time steps. Sequentially, after the problem is solved, the decision corresponding to the first time step is implemented, then the state is evolved and measured, and the optimal control problem at the next time step with the horizon at the same length, thus extending one more time period, is set up and solved. With stochastic MPC, one performs the same operation; however, by solving a multistage stochastic program at each iteration [86]. We will consider this as occurring in real meso-time, which is slower than the true real-time operation, but develops a steady stream of control sequences that feeds data into the real-time layer.

The incorporation of Terminal Costs or Terminal Constraints is critical for guaranteeing that an NMPC scheme exhibits stability and suboptimality guarantees with respect to the desired Dynamic Programming solution. To this end, the use of Approximate Dynamic Programming (ADP) [97] can provide a data-driven approach to define a terminal cost that well approximates the true value function at an ideal optimal policy. With ADP, one maintains an empirical model for the Terminal Cost Value function. In general, it can be difficult to develop a reliable scheme for computing such a function; however, the volume of information available across both the catalog of solutions as well as the real-time operation provides hope for being able to have sufficient data of enough scope to be able to properly inform this function.

6.1 ADP-MPC Formulation

Finally, the ADP program becomes:

$$\begin{aligned}
\min_{u, b, v, z} \quad & \sum_{t=t_0}^{t_{N_T}} \mathcal{R}_{\xi(t)} \left(-\Pi(u^\xi(t), d(\xi(t))) + C(v^\xi(t), z^\xi(t)) \right) \\
& + \mathcal{R}_{\xi(t_1:N_T)} V(u^\xi(t_{N_T}), v^\xi(t_{N_T}), z^\xi(t_{N_T}), d(\xi, t_{N_T}), \dots, d(\xi, t_{N_T} - \bar{\tau}_d)) \\
& - \beta_B \mathbb{E} [\log S(u^\xi)] \\
\text{s.t.} \quad & e(u^\xi(x, t_\tau)) = 0, \quad x \in \Omega \\
& b_{(l)}^\xi(x, t_\tau) = B(v_{(j(l))}^\xi, z_{(k(l))}^\xi, t_\tau), \quad x \in \partial\Omega_{(l)}, \quad \xi \text{ a.s.} \\
& \mathcal{B}(u_{(r)}^\xi(x, t_\tau), u_{(r-1)}^\xi(x, t_\tau), b_{(l)}^\xi(x, t_\tau)) = 0, \quad x \in \partial\Omega_{(l)}, \quad \xi \text{ a.s.} \\
& v^\xi \in \mathcal{V}, \quad \xi \text{ a.s.},
\end{aligned} \tag{14}$$

where $\mathcal{B}$ defines a generic (partial Dirichlet/Neumann/Robin) boundary condition operator.

The lookahead Approximate Dynamic Programming [97], which can also be described as closed-loop Model Predictive Control (MPC) operation, proceeds by the real succession of the current time $\hat{t}$. At each time instant, the problem (14) is formed with $t_0 = \hat{t}$, and computationally solved. The control corresponding to the initial time step $(\hat{v}, \hat{z})$ is performed, and the system is allowed to evolve to the next time step, at which point this entire procedure is repeated.

Decision Variables Recall that generically, the controls include both discrete (open-close) and continuous (dam width/aperture) decisions. In the Meso time layer of computation, however, we consider the application of algorithms for (faster) continuous variate decisions. In particular the focus is to develop procedures that can use off the shelf Nonlinear Programming solvers. To this end, the discrete decisions are eliminated by the following heuristics:

1. We maintain a set of solutions, each of which corresponds to specific values for the discrete variables. Specifically we consider a set of particles $\mathcal{P}$ which at time $\hat{t}$ defines a set of binary decisions as fixed from $t_0 = \hat{t}$ to $t_{N_T} - 1$,

$$\{z_t^{*,p}\}, \quad p \in \mathcal{P}, \quad t \in \{t_0, \dots, t_{N_T} - 1\}.$$

2. The final, receding horizon adjusted, time instance will have the decision variables defined as probability measures, that is,

$$z_{t_{N_T}}^p \in \mu(\{0, 1\}),$$

yielding a continuous optimization problem for all particles p .

3. The total horizon length will be limited to facilitate maintaining a tractably small collection of optimization problems to solve at each time instant. That is, N_T is not particularly large, less than ten. The cost-to-go is meant to define an estimate for the cost value of the system after time N_T . This is by design for the ADP approach for defining V using data from both model simulations and real operation.

From an application perspective, these procedures form the computational backbone supporting supervisory and real-time decision-making in hydraulic systems, where rapid and reliable control updates are required.

6.2 Sequential QCQP (SQCQP) Method

In this Section, we describe a Trust Region Filter Composite Step Sequential Quadratically Constrained Quadratic Programming method for solving stochastically discretized ADP problems. This is the most comprehensive novel algorithm presented in the manuscript, meant to circumnavigate the challenges in the proposed application. As it presents a contribution to mathematical programming in its own right, relative to other methods in this paper, we will describe the components with significantly greater detail.

Because of the challenge of defining Lagrange multipliers described above, we incorporate the Filter technique (e.g., [37]) that obviates the need to track the progress in minimizing a Lagrangian penalty function, as well as use QCQP subproblems to define convex subproblems that incorporate second-order information of both the constraints and the objective function, thus harnessing their potential convergence benefits. Finally, the composite step framework, e.g., [55] obviates the challenge of potentially infeasible linearized constraints. The composite step Trust Region method first computes a direction that reduces the infeasibility measure, and then seeks to decrease the objective while maintaining the new predicted constraint value.

A filter used to perform globalization involves maintaining the set:

$$\mathcal{F} := \left\{ (u^{*,(1)}, b^{*,(1)}, v^{*,(1)}, z^{*,(1)}, f^{*,(1)}, c^{*,(1)}, \dots, \right. \\ \left. u^{*,(F)}, b^{*,(F)}, v^{*,(F)}, z^{*,(F)}, f^{*,(F)}, c^{*,(F)}) \right\} \quad (15)$$

of points that are Pareto optimal with respect to infeasibility and objective value.

We assume that there is a discretization of the probability space $(\Xi, \mathcal{B}, \mathbb{P})$ denoted by a sample set ξ_n with $n \in [N]$. These can constitute, in the standard case, Sample Average Approximation oriented Monte Carlo samples of Ξ for the entire time interval $[0, T]$. We leave time continuous in the formulation, such

that discretizing the optimality conditions of the subproblems can proceed with adaptive time-stepping. We will assume that within the trust region radius Δ , the linear and the quadratic approximation to the PDE exhibit local regularity that permits a primal-dual solution for the subproblems.

We proceed to now define the operations associated with each major iteration of the primary Algorithm, which we later state formally, discuss implementation details, and then derive some theoretical guarantees regarding convergence.

At the start of each iteration k , consider that we have a current estimate $(u^{(k)}(\{\xi_n\}), b^{(k)}(\{\xi_n\}), v^{(k)}(\{\xi_n\}), z^{(k)}(\{\xi_n\}))$ and a trust region radius Δ_k . We trivially consider that the non-anticipativity constraints are satisfied in the definitions of the estimates, and implicitly assume they are enforced throughout the definitions below.

6.2.1 Normal and Tangential Step Computation

Let:

$$\begin{aligned} e^{(k)} &:= e(u^{(k)}(\xi_n)(x, t), t) \\ \mathcal{B}_{(r)}^{(k)} &:= \mathcal{B}(u_{(r)}^{(k)}(\xi_n)(x, t), u_{(r-1)}^{(k)}(\xi_n)(x, t), b_{(l)}^{(k)}(\xi_n)) \\ B_{(l)}^{(k)} &:= B(v_{(l)}^{(k)}(\xi_n)(x, t), z_{(l)}^{(k)}(\xi_n)(x, t), t) \\ S^{(k)} &:= S(u^{(k)}(\xi_n)(x, t), t) \end{aligned}$$

and their corresponding directional derivatives $e'^{(k)}d$ as applied to d , etc. Here we parametrize the control-to-state map e and the state constraint S by t while considering the dams fixed. This can correspond, in the application of dam control, to changes in the precipitation and vessel traffic affecting the river flow e and operational requirements at different times of day for S .

The normal step $(\tilde{n}^{u,(k)}(\{\xi_n\}), \tilde{n}^{b,(k)}(\{\xi_n\}), \tilde{n}^{v,(k)}(\{\xi_n\}), \tilde{n}^{z,(k)}(\{\xi_n\}))$ is defined as the solution to:

$$\begin{aligned} \min_{u,v,z} \quad & \sum_{n=1}^N \int_0^T \left(\int_{\Omega} (e^{(k)}(\xi_n) + e'^{(k)}(\xi_n) \tilde{n}^{u,(k)}(\xi_n))^2 dx \right) dt \\ & + \sum_{n=1}^N \int_0^T \int_{\partial\Omega} \left(\mathcal{B}_{(r)}^{(k)}(\xi_n) + \mathcal{B}'_{(r)}^{(k)}(\xi_n) \right. \\ & \quad \left. (\tilde{n}^{u_{(r)},(k)}(\xi_n), \tilde{n}^{u_{(r-1)},(k)}(\xi_n), \tilde{n}^{b_{(l)},(k)}(\xi_n)) \right)^2 dx dt \\ & + \sum_{n=1}^N \int_0^T \int_{\partial\Omega} \left(b_{(l)}^{(k)}(\xi_n) + \tilde{n}_{(l)}^{b,(k)}(\xi_n) - B_{(l)}^{(k)}(\xi_n) \right. \\ & \quad \left. - B'_{(l)}^{(k)}(\xi_n) (\tilde{n}_{(j(l))}^{v,(k)}(\xi_n), \tilde{n}_{(k(l))}^{z,(k)}(\xi_n)) \right)^2 dx dt \\ \text{s.t.} \quad & v^{(k)}(\xi_n) + \tilde{n}^{v,(k)} \in \mathcal{V}, \quad \forall n \\ & \|(\tilde{n}^{u,(k)}(\{\xi_n\}), \tilde{n}^{b,(k)}(\{\xi_n\}), \tilde{n}^{v,(k)}(\{\xi_n\}), \tilde{n}^{z,(k)}(\{\xi_n\}))\| \leq \Delta_k. \end{aligned} \tag{16}$$

Now, in contrast to the classical variant, we not only compute the new candidate values for the constraints $e(u^{(k)} + \tilde{n}^{u,(k)})$, etc. (which are conventionally used for actual to predicted constraint infeasibility reduction comparison for adaptive adjustment of the trust region and step acceptance criteria), but the gradients and the Hessian (or practically, Hessian-vector product operators) $e'(u(k) + \tilde{n}^{u,(k)})[v]$ and $\nabla^2 e(u(k) + \tilde{n}^{u,(k)})[v]$.

Let us define

$$\begin{aligned}\tilde{e}^{(k)}(\xi_n) &:= e^{\xi_n}(u^{(k)} + \tilde{n}^{u,(k)}), \\ \tilde{e}'^{(k)}(\xi_n) &:= e'^{\xi_n}(u^{(k)} + \tilde{n}^{u,(k)}), \\ \tilde{e}''^{(k)}(\xi_n) &:= e''^{\xi_n}(u^{(k)} + \tilde{n}^{u,(k)})\end{aligned}$$

and similarly for the other constraints as well as the objective functions:

$$\tilde{F}^{(k)}(\xi_n) = \{\tilde{\Pi}^{(k)}(\xi_n), \tilde{C}^{(k)}(\xi_n), \tilde{V}^{(k)}(\xi_n), \tilde{G}^{(k)}(\xi_n)\}$$

and their derivatives, where $\tilde{G} = -\beta_B \log S(\cdot)$ defines the log barrier term, and the first two terms implicitly sum the quantities across $t \in \{t_0, \dots, t_{N_T}\}$.

As in the offline optimization problem formulations, we take $\mathcal{R} = \mathbb{E}$ on account of the challenges of risk consistency and the notation complications of SAA as applied to risk measures. See [72] for a description of risk-averse PDE-constrained problems. Observe that with the regularity of (16), this problem can be solved by computing the adjoint, and subsequently perform (potential adaptive) discretization methods. For comprehensive description without excess volume, the continuous time formulation is defined for the normalized step, and the discretized time formulation of the tangential direction subproblem, defined below.

Let η be a forcing factor. The tangential direction $\tilde{t}^{(k)} = \{\tilde{t}^{u,(k)}, \tilde{t}^{b,(k)}, \tilde{t}^{v,(k)}, \tilde{t}^{z,(k)}\}$ is obtained by solving a QCQP with constraints within new values and minimizing quadratic of objective:

$$\begin{aligned}\min_{u,v,z} \quad & \frac{1}{N} \sum_{n=1}^N \left[\tilde{F}'^{(k)}(\xi_n) \tilde{t}^{(k)} + \frac{1}{2} (\tilde{t}^{(k)})^T \tilde{F}''^{(k)}(\xi_n) \tilde{t}^{(k)} \right] \\ \text{s.t.} \quad & -\eta(e^{(k)} - \tilde{e}^{(k)}) \leq \tilde{e}'^{(k)} \tilde{t}^{u,(k)} + \frac{1}{2} (\tilde{t}^{u,(k)})^T \tilde{e}''^{(k)} \tilde{t}^{u,(k)} \leq \eta(e^{(k)} - \tilde{e}^{(k)}), \forall \xi_n \\ & -\eta(B^{(k)} - \tilde{B}^{(k)}) \leq \tilde{B}'^{(k)} \tilde{t}^{(k)} + \frac{1}{2} (\tilde{t}^{(k)})^T \tilde{B}''^{(k)} \tilde{t}^{u,(k)} \leq \eta(B^{(k)} - \tilde{B}^{(k)}), \forall \xi_n \\ & -\eta(B^{(k)} - \tilde{B}^{(k)} - \tilde{n}^{b,(k)}) \leq \\ & \quad \tilde{B}'^{(k)} \tilde{t}^{(u,v,z),(k)} + \frac{1}{2} (\tilde{t}^{(u,v,z),(k)})^T \tilde{B}''^{(k)} \tilde{t}^{(u,v,z),(k)} + \tilde{t}^{b,(k)} \\ & \quad \leq \eta(B^{(k)} - \tilde{B}^{(k)} - \tilde{n}^{b,(k)}), \forall \xi_n \\ & v^{(k)}(\xi_n) + \tilde{n}^{v,(k)} + \tilde{t}^{v,(k)} \in \mathcal{V}, \forall n \\ & \|\tilde{t}^{(k)}(\{\xi_n\})\| \leq \Delta_k.\end{aligned}\tag{17}$$

6.2.2 Filter and Trust Region Update

Algorithm 2 summarizes the trust-region, filter-based Sequential QCCP (SQCCP) procedure used at the meso scale to solve the stochastically discretized ADP-MPC problem. All mathematical objects required by the algorithm—constraint residuals, their first- and second-order derivatives, quadratic feasibility models, and the quadratic model of the ADP objective—are defined in the preparatory formulas preceding the algorithm.

After computing the normal and tangential steps, the new objective and constraint values are evaluated via

$$\begin{cases} \frac{1}{N} \sum_{n=1}^N F\left((u^{(k)}, b^{(k)}, v^{(k)}, z^{(k)}) + \tilde{n}^{(k)} + \tilde{t}^{(k)}\right), \\ \left\{ e\left((u^{(k)}, b^{(k)}, v^{(k)}, z^{(k)}) + \tilde{n}^{(k)} + \tilde{t}^{(k)}\right) \right\}, \end{cases}$$

together with the corresponding boundary and state-constraint residuals.

To make the filter update precise, we first define the infeasibility measure. Let

$$x := (u, b, v, z)$$

denote the full vector of state and control variables at a candidate iterate, and let

$$e(x), \quad \mathcal{B}(x), \quad B(x)$$

be, respectively, the PDE residual, the internal coupling/boundary residual, and the control-induced boundary residual defined in the ADP-MPC constraint system. The state-inequality constraint is

$$S(u) \geq 0,$$

and we denote its negative part by

$$S^-(u) := \max\{-S(u), 0\}.$$

Using these components, the infeasibility measure is defined by

$$\phi(x) := \|e(x)\| + \|\mathcal{B}(x)\| + \|B(x)\| + \|S^-(x)\|,$$

where $\|\cdot\|$ denotes any norm consistent with the discretization (e.g., an ℓ^2 norm of the discrete residual vector or the corresponding L^2 norm of its continuous representation).

A trial step produces a new point

$$x_{\text{trial}} := x^{(k)} + d^{(k)}, \quad d^{(k)} := \tilde{n}^{(k)} + \tilde{t}^{(k)},$$

with objective value

$$F_{\text{trial}} := F(x_{\text{trial}})$$

and infeasibility

$$\phi_{\text{trial}} := \phi(x_{\text{trial}}).$$

The initial filter $\mathcal{F}_0$ is initialized with the objective and infeasibility values of the initial iterate $x^{(0)}$. The filter at iteration k is a set of ordered pairs

$$\mathcal{F}_k = \{(F^j, \phi^j)\}_{j=1}^{m_k},$$

each representing an objective–infeasibility pair from previously accepted iterates. The filter acceptance criterion used in (S4₂) is:

$$(F_{\text{trial}}, \phi_{\text{trial}}) \text{ is acceptable if for all } (F^j, \phi^j) \in \mathcal{F}_k : F_{\text{trial}} < F^j \text{ or } \phi_{\text{trial}} < \phi^j.$$

That is, the trial point must not be Pareto dominated by any entry in the current filter.

If the trial point is acceptable, the filter is updated by

$$\mathcal{F}_{k+1} := (\mathcal{F}_k \cup \{(F_{\text{trial}}, \phi_{\text{trial}})\}) \setminus \{(F^j, \phi^j) \in \mathcal{F}_k : F_{\text{trial}} \leq F^j \text{ and } \phi_{\text{trial}} \leq \phi^j\}.$$

Thus, any dominated points are removed, and the new point is added. If both the feasibility-only step and the composite step satisfy the filter acceptance test, the algorithm may proceed in two parallel branches or select one according to a heuristic.

To make the algorithm fully explicit, we recall the standard definitions that govern the trust-region updates. The *predicted reduction* associated with the combined step $d^{(k)} := \tilde{n}^{(k)} + \tilde{t}^{(k)}$ is the quadratic model

$$\text{pred}^{(k)} := -\left(F'^{(k)}d^{(k)} + \frac{1}{2}d^{(k)\top}F''^{(k)}d^{(k)}\right),$$

while the *actual reduction* is

$$\text{act}^{(k)} := F(x^{(k)}) - F(x^{(k)} + d^{(k)}).$$

The trust-region ratio

$$\rho^{(k)} := \frac{\text{act}^{(k)}}{\text{pred}^{(k)}}$$

controls the radius update in (S5₂) as

$$\Delta_{k+1} = \begin{cases} \min\{\gamma_{\text{inc}}\Delta_k, \Delta_{\text{max}}\}, & \rho^{(k)} \geq \rho_{\text{good}}, \\ \max\{\gamma_{\text{dec}}\Delta_k, \Delta_{\text{min}}\}, & \rho^{(k)} < \rho_{\text{bad}}, \\ \Delta_k, & \text{otherwise.} \end{cases}$$

where $0 < \rho_{\text{bad}} < \rho_{\text{good}} < 1$, $0 < \gamma_{\text{dec}} < 1 < \gamma_{\text{inc}}$, and $\Delta_{\text{max}} > \Delta_{\text{min}} \in (0, 1)$ are tuning parameters.

The iteration begins with the computation of the *normal step* in (S1₂), which reduces constraint infeasibility by solving a feasibility-oriented QCQP within the trust region. Once the feasibility-improving direction is obtained, the *tangential step* is computed in (S2₂) by minimizing the quadratic model of the objective while remaining in the approximate feasible manifold generated by the normal step. Here, the normal and tangential steps are defined by (16) and (17), respectively. The combined step is tested against the filter in (S4₂), and the trust-region radius is updated in (S5₂) based on the ratio of actual to predicted reduction. If insufficient progress is detected and the radius becomes too small, a feasibility restoration phase is triggered in (S6₂), where only normal steps are taken to recover feasibility. The resulting framework combines feasibility enforcement, second-order curvature information, and global convergence through the filter, yielding a robust SQCQP solver for the stochastic ADP-MPC problem.

Algorithm 2 Sequential QCQP Filter Trust-Region Method for Meso-Scale ADP-MPC

Require: Initial iterate $x^{(0)}$, samples $\{\xi_n\}_{n=1}^N$, initial radius Δ_0 , filter $\mathcal{F}_0$, forcing factor η , trust-region update factors $(\gamma_{\text{inc}}, \gamma_{\text{dec}})$, minimum and maximum radii Δ_{min} and Δ_{max} , and user-prescribed tolerance ε for constraint infeasibility.

Ensure: Sequence $\{x^{(k)}\}$ approaching a stationary point.

(S0₂) Initialize $k = 0$. Evaluate objective and constraint infeasibility at $x^{(0)}$ and insert into $\mathcal{F}_0$.

while infeasibility $> \varepsilon$ **do**

(S1₂) Compute the *normal step* $\tilde{n}^{(k)}$ using the feasibility QCQP.

(S2₂) Compute the *tangential step* $\tilde{t}^{(k)}$ using the quadratic objective model and feasibility windows.

(S3₂) Form the trial point $x_{\text{trial}}^{(k)} = x^{(k)} + \tilde{n}^{(k)} + \tilde{t}^{(k)}$, evaluate objective and infeasibility, and compute the actual/predicted reduction ratio $\rho^{(k)}$.

(S4₂) Apply the *filter acceptance test*. If acceptable, set $x^{(k+1)} = x_{\text{trial}}^{(k)}$ and update $\mathcal{F}_{k+1}$. Otherwise reject the step and keep $x^{(k+1)} = x^{(k)}$.

(S5₂) Update the trust-region radius Δ_{k+1} using $\rho^{(k)}$ and the standard expansion/shrinkage rules.

(S6₂) If the step is rejected and $\Delta_{k+1} \leq \Delta_{\text{min}}$ for a prescribed $\Delta_{\text{min}} > 0$, invoke a *feasibility restoration* step (pure normal-step QCQP).

(S7₂) Increase $k \leftarrow k + 1$.

end while

(S8₂) Return final iterate $x^{(k)}$.

Feasibility Restoration Phase A standard feature of Filter methods is to incorporate a secondary backup feasibility restoration phase. This procedure is activated upon a sequence of steps yielding an unsatisfactory decline in the value of infeasibility, while the trust region is decreased to a small size, precluding significant escape from the infeasible region. In this case, a sequence of purely normal steps (16) is computed, and feasibility is uniquely targeted for the trust region and filter update.

Computing Approximate Subproblem Solutions The subproblems are formed in function space with respect to t and x . Given their structure, the Robinson Constraint Qualification holds for the two subproblems, and the necessary optimality conditions can be solved to obtain a solution. Due to convexity, there is one unique solution. Since the subproblem is a QCQP, the optimality conditions define a linear PDE.

However, given the existence of state inequality constraints, even with linear problems, multipliers, which become Borel measures in the case of state constraints [58], again lose regularity. This presents an open research problem that needs to be addressed. Otherwise, one can pursue discretizing the subproblems directly, but this faces the usual issues of mesh-dependence and the physicality of solutions [108]. The standard approach to state constraints with PDE control includes a Moreau-Yosida regularization, that is, an L^2 penalty in the objective for the constraints [57]. This, however, adds another important parameter, the corresponding regularization multiplicative factor, that needs to be chosen. This complicates any potential convergence analysis, given multiple co-interacting approximations and discretizations, including the barrier term for the original state constraints and any potential spatial (finite difference/volume/element) discretization of the PDEs, as well as the chosen stochastic discretization in formulating the SAA problem.

Potential Theoretical Guarantees The method presented can functionally navigate around the problems associated with the challenge of the absence of well-defined Lagrange multipliers. A filter is a purely primal approach to enforce global convergence, that is, convergence to a stationary point for the original problem from any starting point.

The use of QCQP permits the use of second-order information for accelerating convergence. This is known to be particularly useful in the case of PDE-constrained optimization, given the structure in the matrices defining the Newton iterations.

However, a number of technical details need to be handled before convergence can be established with mathematical certainty. At the MPC-ADP level, the estimate of the Terminal cost should be an accurate-enough estimate of a value function in the Dynamic Programming formulation. Stability and Suboptimality

would have to be established, possibly by applying turnpike methods [49]. For the optimization algorithm convergence guarantees for QCQP would have to be established, for which little is known in the literature (a recent and useful work is [47]). Finally, the barrier parameter considered as both theoretical weak convergence as well as, potentially, a concomitant scheme of adjusting the parameter, should be considered to properly understand the guarantees associated with the proposed algorithm.

6.3 Open Challenges and Future Research

This Section introduced a novel numerical optimization algorithm for performing Nonlinear Model Predictive Control within an Approximate Dynamic Programming framework for the meso-scale operation. The method is new in the literature, meant to circumnavigate the challenges of a control problem without a numerically suitable dual space together with uncertainty. As such a number of important details must be developed before it can be implemented and tested at scale. These include theoretical guarantees - including the stability and suboptimality guarantees needed for trustworthy model predictive control operation [3] and convergence theory for the algorithm itself as far as generating iterates whose asymptotic limits solve the optimization problem. The numerical implementation, as far as the appropriate discretization used to even exhibit approximation properties for functions which can only be guaranteed to lie in a Colombeau Algebra, become the next significant hurdle. Altogether this would constitute a serious research program towards reliable nonlinear optimization tools for solving this class of problems, insofar as a meso-scale intermediary between high-fidelity HPC-guided simulations and real-time operation.

7 Offline–Meso Integration

This section develops the meso-scale evolution mechanism that refines the discrete–continuous control trajectories between the offline and meso-scale tiers. As emphasized earlier, the meso layer does not perform any mixed-integer or continuous optimization beyond the ADP problem that is solved for each particle. Integer programming (IP/MIP) is reserved for the offline layer, where long computation times are acceptable, and for the real-time layer via reduced-order, linearized models. The meso layer instead operates through a structured evolutionary procedure that acts on a population of discrete trajectories.

In practical water management systems, groundwater models are often coupled with surface hydrodynamics through exchange fluxes at riverbeds, reservoirs, or recharge zones. These coupled surface–subsurface systems require mass-conservative interface conditions and may be solved using either partitioned or monolithic strategies. In turn, this must take into account uncertainty in

the exogenous evolution in quantities inflows, boundary conditions, and system response, and their impact on the selection and robustness of control actions.

Offline Catalog as Terminal Value Surrogate. The offline catalog $\mathcal{C}_{\text{off}}$ consists of high-quality, PDE-validated control trajectories collected through Branch-and-Bound, greedy heuristics, DFO/BBO methods, and offline reinforcement learning. Each trajectory includes the full continuous-discrete control sequence, the associated PDE state evolution, and its realized cumulative cost. Thus $\mathcal{C}_{\text{off}}$ provides a data-driven approximation of the downstream cost-to-go that appears in the ADP terminal term of (14),

$$V(u^\xi(t_{N_T}), v^\xi(t_{N_T}), z^\xi(t_{N_T}), d(\xi, t_{N_T}), \dots).$$

Let

$$x := (u^\xi(t_{N_T}), v^\xi(t_{N_T}), z^\xi(t_{N_T}), d(\xi, t_{N_T}), \dots)$$

denote the terminal augmented state. For each catalog entry $c \in \mathcal{C}_{\text{off}}$, let $x_{\text{terminal}}^{(c)}$ be its terminal state and $J_{\text{tail}}^{(c)}$ its realized continuation cost. A non-parametric ADP terminal-value approximation takes the form

$$V_{\text{ADP}}(x) \approx \min_{c \in \mathcal{C}_{\text{off}}} \left(J_{\text{tail}}^{(c)} + \omega(\|x - x_{\text{terminal}}^{(c)}\|) \right),$$

where $\omega(\cdot)$ is a stabilization weight. Alternatively, a parametric surrogate

$$V_\theta(x) = \text{NN}_\theta(x)$$

is trained on the terminal pairs $(x_{\text{terminal}}^{(c)}, J_{\text{tail}}^{(c)})$. Thus, the offline catalog directly informs the ADP terminal cost, providing a computationally efficient proxy for the value function and stabilizing the rolling-horizon MPC behavior at the meso layer.

Particle Representation. Each particle corresponds to a mixed control trajectory for the next T time steps. The binary variables are fixed for the first $T - 1$ steps and relaxed in the final step:

$$z^{(p)}(t) \in \{0, 1\}, \quad t = \hat{t}_c, \dots, \hat{t}_c + T - 2, \quad z^{(p)}(\hat{t}_c + T - 1) \in [0, 1].$$

This preserves discrete feasibility while allowing a low-dimensional continuous relaxation for the ADP solver. No additional optimization of v or z is performed at the meso scale.

Evolution Mechanism. Each particle p is independently evaluated by solving its ADP subproblem, yielding a value estimate $J^{(p)}$. The meso evolution then proceeds as:

- (i) **ADP evaluation:** compute $J^{(p)}$ for each particle.
- (ii) **Selection:** rank particles by $J^{(p)}$ and identify underperformers.
- (iii) **Replacement:** inject high-quality trajectories from $\mathcal{C}_{\text{off}}$.
- (iv) **Mutation:** flip 1–2 binary entries of $z^{(p)}$ following the refined integer uniform sequence [66].
- (v) **Retention:** keep the best particles unchanged.

This purely evolutionary mechanism maintains population diversity without invoking expensive mixed-integer or continuous optimization.

Meso-Layer Master Algorithm. Algorithm 3 summarizes the evolutionary procedure used to propagate and refine the population of discrete–continuous control trajectories at the meso scale. The algorithm operates on a fixed prediction window of length T and treats each particle as an independent candidate policy that is evaluated solely through its ADP value $J^{(p)}$. Because the meso layer performs no discrete optimization of its own, the ADP solver functions as the only source of objective information, and all population updates are driven by the relative ordering of the ADP values.

Step (S0₃) initializes the particle set either from the offline catalog $\mathcal{C}_{\text{off}}$ or by sampling from the learned policy π_{ϕ^*} , thereby ensuring that the initial population spans both historically optimal patterns and recently learned, policy-guided structures. This initialization phase is critical: it seeds the particle ensemble with high-quality discrete configurations that reflect the structure of the true dynamic programming solution.

In each iteration, the algorithm performs three core operations. First, step (S1₃) performs a full ADP evaluation for every particle. This means solving the reduced ADP optimization problem with the particle’s fixed discrete schedule $\{z^{(p)}(t)\}$, obtaining a continuous control refinement $\{v^{(p)}(t)\}$ and a numerical value $J^{(p)}$. Because all particles can be evaluated in parallel, this stage constitutes the main computational workload but scales efficiently on modern computing architectures.

Second, step (S2₃) ranks particles according to their ADP values. This creates a strictly ordered performance hierarchy that drives every subsequent evolutionary

action. Underperformers are flagged for replacement or mutation, while the top-ranked particles form a stable “elite set” that anchors the distribution of the population.

Third, steps (S3₃)-(S5₃) apply the evolutionary operators: replacement injects high-quality discrete schedules from the offline catalog, mutation introduces local variation by flipping one or two components of $z^{(p)}$, and retention preserves the best-performing particles. This yields a balance between exploitation (preserving and refining strong discrete structures) and exploration (introducing diversity). Mutations follow the refined integer uniform sequence [66], ensuring that variation is spatially well-distributed across the binary vector rather than randomly clustered.

Step (S6₃) performs periodic CFD validation of the best particles. If the surrogate $\hat{\mathcal{S}}_\theta$ drifts significantly from the high-fidelity PDE solver, the surrogate is retrained using the discrepant CFD snapshots. This maintains consistency between the meso layer’s ADP evaluations and the true PDE dynamics without imposing CFD costs at every iteration.

Finally, after K_{meso} iterations, step (S7₃) returns the refined population and the single particle with the best value estimate $J^{(p)}$. This best-performing discrete-continuous control pair $(\hat{v}(\hat{t}_c), \hat{z}(\hat{t}_c))$ is passed directly to the real-time layer, where it is linearized and used as the initial guess for rapid model-predictive optimization.

Overall, the algorithm functions as a dynamic bridge between the offline catalog and the real-time MPC solver: it propagates discrete structural information forward in time, injects offline knowledge when necessary, maintains population diversity through mutation, and ensures surrogate fidelity through selective PDE validation. These properties make the meso layer a robust intermediate stage capable of tracking, adapting, and refining discrete hydropower operating modes in fluctuating stochastic regimes.

Algorithm 3 Meso-Scale Particle Evolution

Require: Particle set $\mathcal{P}_{\hat{t}_c}$; ADP solver; offline catalog $\mathcal{C}_{\text{off}}$; mutation rate λ ; K_{meso}

Ensure: Updated population $\mathcal{P}_{\hat{t}_c + K_{\text{meso}}}$ and the best control pair

(S0₃) Initialize $\mathcal{P}_{\hat{t}_c}$ from $\mathcal{C}_{\text{off}}$ or policy π_{ϕ^*} .

for $k = 0, 1, \dots, K_{\text{meso}} - 1$ **do**

for each particle $p \in \mathcal{P}_{\hat{t}_c + k}$ **do**

 (S1₃) Solve particle’s ADP problem to obtain $J^{(p)}$.

end for

(S2₃) Rank particles; identify underperformers.

(S3₃) Replace worst-performing particles with samples from $\mathcal{C}_{\text{off}}$.

(S4₃) Mutate selected particles:

 flip 1–2 entries of $z^{(p)}$ following the uniform integer sequence.

(S5₃) Retain top-performing particles unchanged.

(S6₃) Periodically validate top particles with CFD and update surrogate.

end for

(S7₃) Return updated population and best-performing particle.

Discussion on Mutation and Surrogate Feedback. No gradients or adjoint sensitivities are used at the meso-scale particle evolution master algorithm scale; the only signal guiding evolution is the ADP value estimate $J^{(p)}$. Binary mutation uses the refined integer sequence [66] to ensure non-clustered, well-distributed flips. Solving the ADP problem is performed by applying off-the-shelf Nonlinear Programming solvers in parallel for all particles p . The meso layer then passes the best control pair $(\hat{v}(\hat{t}_c), \hat{z}(\hat{t}_c))$ to the real-time layer for linearization and execution. Meanwhile the information gained from the ADP solution computation is added in a database for learning at both the Meso layer (for the terminal cost V) and the offline layer.

7.1 Open Challenges and Future Research

There are a number of details to ensure the operation described in this Section is successful. Notably, the optimal, or possibly satisficing set of optimal solution must change at a sufficiently slow latency, relative to the computationally required load. Conversely, the number of particles must evolve, through mutation to new solutions and selection of the leading solution to feed to the lower layer, with sufficiently frequent latency for the solution pathfollowing to proceed adequately.

8 Fast Newton Solutions of Reduced Order Models at Real Time

Ultimately, the practical goal of the problem(s) discussed in this paper is to perform operations in real-time, that is, sequential decisions that make changes within a short time interval. The shorter the time interval, the more quickly the system can adapt to changes in the environment. At the same time, short time intervals prevent the computation of complex and sophisticated procedures. In a real-time control context, coupled complex PDE systems must typically be approximated using surrogate models or reduced-order interface representations in order to maintain computational efficiency within operational time constraints.

In practical terms, for dam networks, this corresponds to constructing and selecting from a set of admissible control policies that encode feasible gate operations under varying hydrological conditions. These models are integrated with the meso layer by Model Order Reduction (MOR). MOR, as described in the well-known three-volume text [12], is a comprehensive toolkit for obtaining simpler, e.g., linear models that approximate a more complicated nonlinear model with some degree of instrumental accuracy. The specific details of performing MOR for meso-to-real-time optimization would constitute a comprehensive research program, and in this paper, some general suggestions and options are presented, without claiming authority on their optimality.

8.1 Real Time Control

The implementation of MPC in real-time, that is, with short durations between measuring changes in the environment and performing the next control action, depends on fast computational procedures that take advantage of the fact that drastic changes in the environment typically only occur over long time intervals. A principal component is the preferential application of Newton methods. Notably, Newton iterations are quadratically convergent to a solution from starting points sufficiently close, and so if a method traces a trajectory of solutions while maintaining proximity within the radius of fast convergence, real-time operation can be truly fast [32].

Within engineering systems applying real-time MPC, there are two standard paradigms: Real Time Iteration (RTI) [34] and advanced step NMPC [60]. With RTI, a sequence of Newton steps is used directly to perform path following along the sequence of Nonlinear Programming problems defining the control for the receding horizon. The perturbation, beyond the increment in time, thus dropping one stage and adding another, includes state-model mismatch, i.e., the measured state replaces the initial state. With advanced-step MPC, 1) after a control is implemented, a new optimization problem is solved, using the previous model; 2) once the problem is solved, the real state is measured, and

a corrector Newton step updates the next control.

To account for active set changes during the procedure, predictor-corrector quadratic programs (QPs) are used to traverse these while maintaining Newton fast convergence. See, e.g., [74, 75]. Meanwhile, Filtering and Moving Horizon Estimation are optimization techniques that enable the simultaneous estimation, from real-time observations and measurements, of parameters and states [101].

8.2 Recent Advances in Mixed Integer Real Time Control

Real-time MPC has classically been limited to continuous optimization. Integer and other combinatorial decision variables are fundamentally distinct from the smooth trajectories - states and parameters changing gradually over time - amenable to Newton methods. Instead, a change in the solution involves a discrete jump, for which, generically, fast computational methods become difficult to develop with demonstrable speed and robustness [48].

Recently, the study of mixed integer optimization within the context of NMPC has begun to be addressed. Advances in the power of computational hardware have improved integer programming solvers to the point that problems of a dimension that is interesting and relevant to many industrial applications become within range for real-time numerical optimization-based control. For an early work in this space, see [67]. Extensions of RTI for mixed integer problems are explicitly developed in [30, 100].

Notably, fast computation becomes reasonable when a significant subset of the discrete decision variables is fixed, leaving a small number remaining for optimization (which can be found through a slower-time branch and bound procedure, in the references above, or through the meso-to-real-time operative integration, here). If $\mathcal{D} \subset [n_z]$ are the variables worth considering, and $\bar{e}$ is the linearized model, and F_k is the total objective (i.e., including the future time over the fixed horizon), two available approaches are available.

One can solve a Mixed Integer Linear Program, wherein the discrete variable appears explicitly in the state model, which is otherwise linearized with respect to the continuous variables:

$$\begin{aligned} \min_{d_u, d_v, z_{\mathcal{D}}} \quad & \nabla_u F_k^T d_u + \nabla_v F_k^T d_v \\ \text{s.t.} \quad & u_k(t+1) + d_u(t+1) = \bar{e}_k(z) + \nabla_u \bar{f}_k(z, t)^T d_u + \nabla_v \bar{f}_k(z, t)^T d_v \\ & -\Delta_k \leq d_v \leq \Delta_k, \quad z_{\mathcal{D}} \in \mathbf{P}(\mathcal{D}) \end{aligned}$$

where $\mathbf{P}(\mathcal{D})$ is the power set (set of all subsets) of $\mathcal{D}$.

In [100], a mixed integer quadratic optimization problem is presented, with a similar two-time-scale approach with sequential linearization of the underlying

linear model. The paper presents an approach and software for solving this set of problems, that is, of the form:

$$\begin{aligned} \min_{d_u, d_v, z_{\mathcal{D}}} \quad & \nabla_u F_k^T d_u + \nabla_v F_k^T d_v + \frac{1}{2} \begin{pmatrix} d_u \\ d_v \end{pmatrix}^T H_k \begin{pmatrix} d_u \\ d_v \end{pmatrix} \\ \text{s.t.} \quad & u_k(t+1) + d_u(t+1) = \bar{e}_k(z) + \nabla_u \bar{f}_k(z, t)^T d_u + \nabla_v \bar{f}_k(z, t)^T d_v, \end{aligned}$$

where H_k is an approximation to the Hessian of the Lagrangian with respect to the continuous variables. An efficient numerical procedure that can solve these in real time on embedded (that is, restricted computational capacity) hardware.

For a recent warm start procedure for considering stochasticity within the context of NMPC with MILPs, see [84].

8.3 Piecewise Linear Approximation of the Shallow Water Equations

A prominent approach to obtaining tractable approximations of the nonlinear Shallow Water equations relies on piecewise linear representations of the governing dynamics arising from semi-implicit space-time discretizations. In conservative form, the two-dimensional Shallow Water system reads

$$\begin{aligned} \partial_t \eta + \nabla \cdot (\eta \mathbf{u}) &= 0, \\ \partial_t (\eta \mathbf{u}) + \nabla \cdot (\eta \mathbf{u} \otimes \mathbf{u}) + g \eta \nabla H &= \mathbf{S}(\eta, \mathbf{u}), \end{aligned} \tag{18}$$

where η denotes the water depth, $\mathbf{u}$ the depth-averaged velocity, $H = h + z_b$ the free-surface elevation relative to the bathymetry z_b , g the gravitational acceleration, and $\mathbf{S}$ collects source terms such as bottom friction or Coriolis effects.

Within finite-volume or finite-difference frameworks, semi-implicit discretizations treat the gravity-driven pressure term implicitly while retaining an explicit or linearized treatment of the advective fluxes. Denoting by η_i^n and $\mathbf{u}_i^n$ the cell-averaged unknowns at time level t^n , a typical semi-implicit update can be written schematically as

$$\begin{aligned} \frac{\eta_i^{n+1} - \eta_i^n}{\Delta t} + \sum_{j \in \mathcal{N}(i)} F_{ij}^n &= 0, \\ \frac{(\eta \mathbf{u})_i^{n+1} - (\eta \mathbf{u})_i^n}{\Delta t} + \sum_{j \in \mathcal{N}(i)} \hat{F}_{ij}^n + g \eta_i^{n+1} \nabla H_i^{n+1} &= \mathbf{S}_i^{n+1}, \end{aligned} \tag{19}$$

where F_{ij}^n and $\hat{F}_{ij}^n$ denote numerical fluxes across cell interfaces and $\mathcal{N}(i)$ is the set of neighboring cells.

The key observation underlying the piecewise linear formulation is that nonlinearities associated with wetting-and-drying transitions, positivity constraints on h , and friction laws can be expressed through max-min operators or complementarity conditions. For instance, enforcing non-negativity (or other, more robust bound) of the water depth leads to relations of the form

$$H_i^{n+1} = \max\{0, \eta_i^{n+1} + z_b\}. \quad (20)$$

As a result, the fully discrete problem at each time step can be cast as a piecewise linear system

$$A(\boldsymbol{\sigma}) \mathbf{x} = \mathbf{b}(\boldsymbol{\sigma}), \quad (21)$$

where $\mathbf{x}$ collects the discrete water levels and velocities, $\boldsymbol{\sigma}$ encodes the active set associated with wet and dry cells, and the matrices $A(\boldsymbol{\sigma})$ and vectors $\mathbf{b}(\boldsymbol{\sigma})$ depend linearly on this regime selection.

The combination of equations across δt -iterative time steps from some initial $t = 0$ to some final time $T = H\delta t$, together with any piece-wise kinks in the state, provide discretization knots. The system will be interpreted as exhibiting a linear interpolation between these knots, forming an optimization problem with many mixed integer linear constraints.

For such systems, specialized iterative algorithms based on nested Newton or active-set strategies admit finite termination and strong robustness properties [17, 18]. Crucially, these formulations preserve essential structural features of the continuous equations, including exact mass conservation at the discrete level, unconditional stability with respect to the time step size, and a consistent and reliable treatment of wetting-and-drying interfaces, while substantially reducing the computational burden relative to fully nonlinear implicit schemes.

High-resolution extensions of these ideas to multidimensional free-surface hydrodynamics further demonstrate that accurate advection-diffusion balances can be maintained without sacrificing stability or monotonicity [23]. In particular, the resulting schemes can be interpreted as regime-dependent linearizations embedded within a globally conservative finite-volume framework. More generally, nested Newton-type strategies for solving the associated nonsmooth algebraic problems provide a systematic mechanism for organizing such piecewise linear models within large-scale simulations [24].

From the perspective of optimization under PDE constraints, the induced piecewise linear or piecewise affine structure of the discrete dynamics is particularly appealing, as it naturally aligns with mixed-integer, complementarity-based, and hybrid optimization formulations while retaining a close connection to the underlying physical processes. At the same time, the inherent nonsmoothness and switching behavior introduced by these approximations highlight fundamental analytical and computational challenges—such as limited regularity of control-to-state mappings, sensitivity to active-set changes, and the scalability of large-scale solvers—that motivate the broader discussion developed in this work.

8.4 Open Challenges and Future Research

Even with the Model Order Reduction discussed below, the dimensionality of the optimization problems defined presents significant challenges to existing mixed-integer solvers. Fortunately, since the problem functions and solutions are not expected to significantly change with the latency under consideration, this is an exercise in the numerical algebra of warm and hot starting. This presents an impetus to develop parametric mixed-integer optimization techniques and software [8].

9 Meso Real-Time Integration

The Meso-Real-Time Integration layer provides the computational bridge between the nonlinear surrogate-augmented meso optimization and the low-latency Real-Time Iteration (RTI) controller. Its primary objective is to convert the nonlinear surrogate dynamics and the refined meso-level particle information into a model-order-reduced (MOR) and locally linearized representation. This representation is then passed to the RTI solver, which computes the actual real-time control action. The Meso-RT integration, therefore, focuses on *model preparation*, not direct real-time optimization. Overall, the proposed framework can be interpreted, in the context of dam networks, as a unified approach to bridging long-term planning, supervisory updates, and real-time hydraulic control.

Throughout this section, T denotes the finite receding prediction horizon used by the meso and real-time layers, while T_{final} denotes the (full) time horizon of the underlying time-dependent PDE-constrained optimization problem used for conceptual analysis.

Role within the Three-Tier Architecture. From the meso layer, we receive the validated surrogate $\hat{\mathcal{S}}_\theta$, the refined particle ensemble $\mathcal{P}_{t_c}$, and the best-performing local nominal trajectory. The Meso-RT interface uses this information to construct:

- a reduced-order model (ROM) via classical or data-driven MOR,
- a local linearization of the ROM around the nominal meso trajectory,
- an initialization and warm-start guess for RTI.

The real-time controller then solves a reduced-dimension quadratic or linear MPC problem in a single RTI step. Thus, information flows *from* meso-level adaptation *into* real-time execution. The real-time operation will, at the same

time, provide streaming data on the physical environment that is used for both meso and offline computational updates.

Operational Data Assimilation and Local Linearization. At each real-time instant $\hat{t}_c$, the following procedure updates the model used by RTI:

1. Assimilate the measured hydrodynamic state $u^{\hat{\xi}}(\hat{t}_c)$ and update the forecast distribution $\tilde{\rho}(\xi)$.
2. Identify the best meso particle (v^*, z^*) and its corresponding surrogate trajectory.
3. Construct a reduced-order model

$$\hat{u}_{\text{ROM}}(t) = A_{\text{ROM}} \hat{u}_{\text{ROM}}(t) + B_{\text{ROM}} v(t) + F_{\text{ROM}} + (\text{small residual}),$$

using MOR techniques (POD, operator inference, autoencoder-based ROM, etc.). Here, the terms A_{ROM} and B_{ROM} denote the reduced-order state and control matrices obtained from the MOR procedure, F_{ROM} denotes the affine forcing/bias term of the reduced model, and $\hat{u}_{\text{ROM}}(t)$ denotes the reduced-order approximation of the hydrodynamic state at time t .

4. Linearize the ROM around the nominal meso particle:

$$\delta \hat{u}_{t+1} = A_t \delta \hat{u}_t + B_t \delta v_t + f_t,$$

producing the matrices needed for RTI. Here, the quantity $\delta \hat{u}_t$ denotes the deviation of the reduced state from the nominal meso trajectory, δv_t is the deviation of the control input from its nominal value, A_t and B_t are the Jacobians of the ROM with respect to state and control evaluated at the nominal trajectory, and f_t is the affine residual arising from the first-order linearization.

5. Produce warm-start state and control guesses for RTI (not the final real-time control) using the meso particle and the state deviation from the ROM.

Meso-Real-Time Integration Algorithm. The integration workflow is summarized below. Note that the algorithm prepares local models and initial guesses for RTI, but does not compute the control action itself; the RTI module handles that step.

Algorithm 4 Meso-Real-Time Integration for Hierarchical Control

Require: Surrogate operator $\hat{\mathcal{S}}_\theta$; refined particles $\mathcal{P}_{\hat{t}_c}$; policy π_{ϕ^*} ; tolerance ϵ

Ensure: Linearized reduced-order model and warm-start for RTI

(S0₄) Initialize time $t = \hat{t}_c$; load surrogate and meso particles.

while $t < T$ **do**

(S1₄) Acquire measured state $u^{\hat{\xi}}(t)$ and update uncertainty.

(S2₄) Select best-performing meso trajectory (v^*, z^*) .

(S3₄) Construct MOR model and obtain ROM matrices.

(S4₄) Linearize ROM around (v^*, z^*) , producing (A_t, B_t, f_t) .

(S5₄) Generate a warm-start $(v_{\text{init}}, z_{\text{init}})$ for RTI.

(S6₄) Compare ROM prediction with measurement:

 If discrepancy $> \epsilon$, register data for surrogate retraining.

(S7₄) Advance $t \leftarrow t + \Delta t$.

end while

(S8₄) Output linearized ROM and warm-start trajectory for RTI.

The output of Algorithm 4 provides the RTI layer with reduced matrices (A_t, B_t, f_t) , a warm-start control trajectory, and updated uncertainty-aware cost estimates. The RTI solver then performs a single SQP step to compute the executable real-time control. Thus, the Meso-RT integration is responsible for model preparation and adaptation, while RTI is responsible for fast optimal control computation.

Observe that the best-performing meso trajectory is expected to be drawn from memory frequently with real-time operation. This permits the consideration of asynchronous methods at the meso layer, i.e., wherein multiple parallel particle evolutions are themselves competing for the best solution or even asynchronous evolution strategies.

Connection to Parareal Real-Time PDE-Constrained Optimization.

The Meso-Real-Time integration strategy is closely related, at a structural level, to parareal time-domain decomposition methods for time-dependent PDE-constrained optimization, as developed by Ulbrich within the real-time optimization framework [120]. In that setting, parareal methods are embedded into a generalized SQP algorithm to enable parallel-in-time solution of large-scale transient PDE optimization problems, while retaining the preparation-feedback separation that is central to real-time iteration (RTI) schemes.

In Ulbrich's formulation, the time horizon $[0, T_{\text{final}}]$ (in contrast to the receding horizon T used by the meso-real-time layers), with $T_{\text{final}} > 0$, is decomposed into subintervals $[T_n, T_{n+1}]$, where $\{T_n\}_{n=0}^N$ denotes a partition of $[0, T_{\text{final}}]$ satisfying $T_0 = 0$, $T_N = T_{\text{final}}$, and $\Delta T := T_{n+1} - T_n$. The state equation is written in propagator form

$$u(T_{n+1}) = g(T_n, u(T_n)),$$

where $u(t)$ denotes the PDE state, understood as an element of an appropriate Banach or Hilbert space induced by the spatial discretization of the PDE, and g denotes the exact time- ΔT solution operator of the PDE. A coarse approximation g^Δ of g is assumed available, where g^Δ denotes a computationally inexpensive coarse-grid approximation of g , typically induced by a lower-order or more dissipative time discretization.

The parareal update for the state at iteration $k \in \mathbb{N}$ then reads

$$u^{k+1}(T_{n+1}) = g^\Delta(u^{k+1}(T_n)) + g(u^k(T_n)) - g^\Delta(u^k(T_n)), \quad (22)$$

which can be interpreted as a preconditioned fixed-point iteration for the multiple-shooting formulation of the time-dependent PDE. Ulbrich applies this correction simultaneously to the state and adjoint equations inside a generalized SQP method, allowing inexact and user-provided solvers while maintaining global convergence.

The Meso–Real-Time interface proposed here follows the same conceptual separation of responsibilities as the parareal RTI framework, but replaces adjoint-based correction by value-based and model-based coupling. Specifically, the meso layer fulfills a role analogous to the coarse global propagator in parareal–SQP: it provides a globally informed but low-latency description of future system behavior through the refined particle ensemble $\mathcal{P}_{\hat{t}_c}$, surrogate dynamics $\hat{\mathcal{S}}_\theta$, and associated ADP value estimates.

At each control time $\hat{t}_c$, the Meso–RT layer performs a preparation step in the precise sense of real-time iteration theory [15]. Here, t denotes the discrete real-time control index associated with the RTI sampling instants. The Meso–RT layer constructs a reduced-order model and a local linearization around the nominal meso trajectory (v^*, z^*) , yielding

$$\delta \hat{u}_{t+1} = A_t \delta \hat{u}_t + B_t \delta v_t + f_t,$$

which corresponds to the linearized dynamics required to assemble the quadratic program solved in the RTI feedback step. The term f_t collects affine residuals arising from the first-order linearization of the reduced-order model around the nominal trajectory. This mirrors the role of the fine propagator in parareal methods, which enforces local consistency with the high-fidelity PDE on short time intervals.

A key difference lies in the coupling mechanism. Whereas parareal–SQP enforces continuity across time subdomains through iterative state and adjoint corrections, the present architecture enforces temporal consistency implicitly through rolling-horizon execution, warm-started RTI, and periodic surrogate validation. No adjoint sensitivities or second-order information are propagated from the real-time layer back to the meso layer, reflecting the presence of mixed-integer decisions, non-smooth dynamics, and real-time latency constraints.

From this perspective, the Meso–Real-Time Integration layer can be interpreted as a value- and model-based analogue of parareal real-time PDE optimization: global temporal structure is propagated forward via meso-level particle evolution and ADP terminal value surrogates, while local PDE fidelity is recovered through RTI-based linearized MPC at execution time. This preserves the preparation–feedback philosophy and robustness properties of classical RTI schemes, while extending parareal ideas to hybrid, stochastic, and real-time constrained PDE control settings.

Parareal–SQP and 3-Time-Scale ADP. While Ulbrich’s parareal–SQP framework provides a powerful approach for time-dependent PDE-constrained optimization, it can not be applied directly in the present setting for several structural reasons. Parareal–SQP relies on adjoint-based sensitivities, multiple-shooting continuity corrections, and iterative convergence across the full time horizon, all of which are incompatible with mixed-integer decisions, non-smooth switching dynamics, and strict real-time latency constraints. In contrast, the proposed architecture deliberately separates the global structural adaptation from local real-time execution: the meso layer propagates discrete and value information forward in time without adjoints, while the real-time layer executes a single linearized SQP step using reduced-order models. This design preserves the preparation–feedback philosophy of real-time iteration schemes while avoiding the algorithmic overhead and regularity assumptions required by parareal–SQP. At the same time, the use of coarse and fine grid methods, while they need to be adapted to heterogeneous latency, and the MGOpt framework provide a useful set of tools to, for instance, connect the finest linear discretization at the real-time layer through to the coarsest nonlinear discretization at the meso layer, for maximal information efficiency.

9.1 Open Challenges and Future Research

The Model Order Reduction from the meso to the real time layer must be quite substantial. Meso layer solutions involve highly nonlinear solutions and must incorporate a long time horizon. At the real-time layer, the problems must be of sufficiently low dimension to be tractable, while at the same time exhibiting reasonable approximation to the real flow. The propagation of approximation and uncertainty accuracy in the real-time from the meso and from the high fidelity offline layer should be investigated theoretically for deriving guarantees of some bounds, while extensive numerical tests will need to be performed validating this accuracy and precision.

10 Discussion and Conclusion

The thought exercise of considering what research and development are required to potentially address the mixed-integer optimization of hydropower cascades with river flow modeled by the shallow water equations, as well as their potential concrete implementation in water engineering, that is, with computational real-time considerations as well as uncertainty, presents a challenging problem in broad applied and computational mathematics. Any realistic implementation requires extensions along the state of the art across many relevant domains - along PDE theory, the calculus of variations, dynamic programming, approximate dynamic programming, model predictive control, uncertainty quantification, reduced order modeling, and scientific computing. Even more significantly, their *integration*, that is, traversing and systematizing across domains and levels of abstraction in a manner that is theoretically sound and computationally tractable, presents uncharted interdisciplinary investigation. As such, this highlights the broad importance, in applied and computational mathematics in general, of the necessity of a systematized understanding of multiple components of the toolkit and their synergy, to continue to leverage the field of applied and computational mathematics for solving real-world problems and continue technological innovation into the future. This problem has, in its nonlinearity and complexity, features exhibited by many contemporary systems of interest, and this paper hopes to be an illustration for and impetus to greater elucidation of methodological interdisciplinary integration towards modeling and control.

Acknowledgements V.K. and M.K. acknowledges funding support from the National Centre for Energy II (TN02000025). KK also acknowledges financial support from the Czech National Science Foundation under Project 24-11664S and M.K. acknowledges the financial support of the Austrian Science Foundation under <https://doi.org/10.55776/PAT2747625>. The authors would like to thank Sven Leyffer, Paul Manns, Wladimir Neves and Thomas Surowiec for valuable insight and discussions on the problem.

References

- [1] Remi Abgrall and Yongle Liu. A new approach for designing well-balanced schemes for the shallow water. *Appl. Math. Comput.*, 5:370–402, 2023.
- [2] A. Alla, M. Berardi, and L. Saluzzi. State Dependent Riccati for dynamic boundary control to optimize irrigation in Richards’ equation framework. *Mathematics and Computers in Simulation*, 232:261–275, 2025.
- [3] Frank Allgöwer and Alex Zheng. *Nonlinear model predictive control*, volume 26. Birkhäuser, 2012.

- [4] Hans Wilhelm Alt and Stefan Luckhaus. Quasilinear Elliptic-Parabolic Differential Equations. *Mathematische Zeitschrift*, 183(3):311–341, 1983.
- [5] Charles Audet and J. E. Dennis. Mesh adaptive direct search algorithms for constrained optimization. *SIAM Journal on Optimization*, 17(1):188–217, January 2006.
- [6] Emmanuel Audusse, François Bouchut, Marie-Odile Bristeau, Rupert Klein, and Benoît Perthame. A Fast and Stable Well-Balanced Scheme with Hydrostatic Reconstruction for Shallow Water Flows. *SIAM Journal on Scientific Computing*, 25(6):2050–2065, 2004.
- [7] Afonso S Bandeira, Katya Scheinberg, and Luis Nunes Vicente. Convergence of trust-region methods based on probabilistic models. *SIAM Journal on Optimization*, 24(3):1238–1264, 2014.
- [8] Bernd Bank and Reinhard Mandel. *Parametric integer optimization*, volume 39. Walter de Gruyter GmbH & Co KG, 2021.
- [9] Jacob Bear. *Dynamics of Fluids in Porous Media*. Elsevier, New York, 1972.
- [10] Jacob Bear and Arnold Verruijt. *Modeling Groundwater Flow and Pollution*. D. Reidel Publishing Company, 1987.
- [11] Malte Behrens, Hans Georg Bock, Sebastian Engell, Phawitphorn Khobkhun, and Andreas Potschka. Real-time pde constrained optimal control of a periodic multicomponent separation process. *Trends in PDE Constrained Optimization*, pages 521–537, 2014.
- [12] Peter Benner, Mario Ohlberger, Albert Cohen, and Karen Willcox. *Model reduction and approximation: theory and algorithms*. SIAM, 2017.
- [13] Sten Bergström and Göran Lindström. Interpretation of runoff processes in hydrological modelling—experience from the hbv approach. *Hydrological Processes*, 29(16):3535–3545, 2015.
- [14] Lorenz T Biegler, Omar Ghattas, Matthias Heinkenschloss, David Keyes, and Bart van Bloemen Waanders. *Real-time PDE-constrained Optimization*. SIAM, 2007.
- [15] Hans Georg Bock, Moritz Diehl, Ekaterina Kostina, and Johannes P. Schlöder. *1. Constrained Optimal Feedback Control of Systems Governed by Large Differential Algebraic Equations*, page 3–24. Society for Industrial and Applied Mathematics, January 2007.
- [16] A. Borzì and G. von Winckel. Multigrid methods and sparse-grid collocation techniques for parabolic optimal control problems with random coefficients. *SIAM Journal on Scientific Computing*, 31(3):2172–2192, January 2009.

- [17] Luigi Brugnano and Vincenzo Casulli. Iterative Solution of Piecewise Linear Systems. *SIAM J. Scientific Computing*, 30:463–472, 01 2008.
- [18] Luigi Brugnano and Vincenzo Casulli. Iterative Solution of Piecewise Linear Systems and Applications to Flows in Porous Media. *SIAM Journal on Scientific Computing*, 31(3):1858–1873, 2009.
- [19] Ana Bušić and Sean Meyn. Ordinary differential equation methods for markov decision processes and application to kullback–leibler control cost. *SIAM Journal on Control and Optimization*, 56(1):343–366, January 2018.
- [20] Yanzhao Cao, Somak Das, and Hans-Werner van Wyk. Adaptive stochastic gradient descent for optimal control of parabolic equations with random parameters. *Numerical Methods for Partial Differential Equations*, 38(6):2104–2122, February 2022.
- [21] Jesus Carrera, Alvaro Alcolea, Agustí Medina, Julian Hidalgo, and Lucas Slooten. Inverse Problem in Hydrogeology. *Water Resources Research*, 41:W09424, 2005.
- [22] Oscar Castro-Orgaz, Willi H Hager, et al. *Shallow water hydraulics*. Springer, 2019.
- [23] Vincenzo Casulli and Paola Zanolli. High resolution methods for multi-dimensional advection–diffusion problems in free-surface hydrodynamics. *Ocean Modelling*, 10(1):137–151, 2005. The Second International Workshop on Unstructured Mesh Numerical Modelling of Coastal, Shelf and Ocean Flows.
- [24] Vincenzo Casulli and Paola Zanolli. A Nested Newton-Type Algorithm for Finite Volume Methods Solving Richards’ Equation in Mixed Form. *SIAM Journal on Scientific Computing*, 32(4):2255–2273, 2010.
- [25] Michael A. Celia, Edward T. Bouloutas, and Rebecca L. Zarba. A General Mass-Conservative Numerical Solution for the Unsaturated Flow Equation. *Water Resources Research*, 26(7):1483–1496, 1990.
- [26] Ralph T Cheng, Vincenzo Casulli, and Jeffrey W Gartner. Tidal, residual, intertidal mudflat (trim) model and its applications to san francisco bay, california. *Estuarine, Coastal and Shelf Science*, 36(3):235–280, 1993.
- [27] Jean François Colombeau. *Elementary introduction to new generalized functions*, volume 113. Elsevier, 2011.
- [28] Antonio J Conejo, Miguel Carrión, Juan M Morales, et al. *Decision making under uncertainty in electricity markets*, volume 1. Springer, 2010.
- [29] Andrew R. Conn, Katya Scheinberg, and Luis N. Vicente. *Introduction to Derivative-Free Optimization*. Society for Industrial and Applied Mathematics, January 2009.

- [30] Massimo De Mauri, Wim Van Roy, Joris Gillis, Jan Swevers, and Goele Pipeleers. Real time iterations for mixed-integer model predictive control. In *2020 European Control Conference (ECC)*, pages 699–705. IEEE, 2020.
- [31] Marie Delaunay and Florin A. Radu. On the Stability and Convergence of Linearization Schemes for Richards Equation. *SIAM Journal on Numerical Analysis*, 55(2):540–565, 2017.
- [32] Peter Deufhard. *Newton methods for nonlinear problems: affine invariance and adaptive algorithms*, volume 35. Springer Science & Business Media, 2011.
- [33] María J. Castro Díaz, Alexander Kurganov, and Tomás Morales de Luna. Path-Conservative Central-Upwind Schemes for Nonconservative Hyperbolic Systems. *ESAIM: Mathematical Modelling and Numerical Analysis*, 53(3):959–985, 2019.
- [34] Moritz Diehl, Hans Georg Bock, and Johannes P Schlöder. A real-time iteration scheme for nonlinear optimization in optimal feedback control. *SIAM Journal on control and optimization*, 43(5):1714–1736, 2005.
- [35] Sakine Esmaili, M. R. Eslahchi, and Delfim F. M. Torres. Optimal control for a nonlinear stochastic pde model of cancer growth. *Optimization*, 73(9):2745–2789, July 2023.
- [36] Xinyu Fei, Lucas T Brady, Jeffrey Larson, Sven Leyffer, and Siqian Shen. Binary control pulse optimization for quantum systems. *Quantum*, 7:892, 2023.
- [37] Roger Fletcher, Sven Leyffer, Ph L Toint, et al. A brief history of filter methods. *Preprint ANL/MCS-P1372-0906, Argonne National Laboratory, Mathematics and Computer Science Division*, 36, 2006.
- [38] Gastao SF Frederico, Paolo Giordano, Alexandr A Bryzgalov, and Matheus J Lazo. Calculus of variations and optimal control for generalized functions. *Nonlinear Analysis*, 216:112718, 2022.
- [39] Claudia Garetto. Topological structures in colombeau algebras: investigation of the duals of $\mathcal{G}_c(\Omega)$, $\mathcal{G}(\Omega)$ and $\mathcal{G}_S(\mathbb{R}^n)$. *Monatshefte für Mathematik*, 146(3):203–226, 2005.
- [40] Dominik Garmatter, Margherita Porcelli, Francesco Rinaldi, and Martin Stoll. Improved penalty algorithm for mixed integer PDE constrained optimization problems. *Computers & Mathematics with Applications*, 116:2–14, 2022.
- [41] Sebastian Garreis, Thomas M. Surowiec, and Michael Ulbrich. An Interior-Point Approach for Solving Risk-Averse PDE-Constrained Optimization Problems with Coherent Risk Measures. *SIAM Journal on Optimization*, 31(1):1–29, 2021.

- [42] Caroline Geiersbach and Winnifried Wollner. Optimality conditions for convex stochastic optimization problems in banach spaces with almost sure state constraints. *SIAM Journal on Optimization*, 31(4):2455–2480, 2021.
- [43] Fred Glover. Future paths for integer programming and links to artificial intelligence. *Computers & Operations Research*, 13(5):533–549, 1986.
- [44] Fred Glover. Tabu search—part i. *ORSA Journal on Computing*, 1(3):190–206, 1989.
- [45] Fred Glover. Tabu search—part ii. *ORSA Journal on Computing*, 2(1):4–32, 1990.
- [46] Fabian Gnengel, Armin Fügenschuh, Michael Hagel, Sven Leyffer, and Marcus Stiemer. A solution framework for linear pde-constrained mixed-integer problems. *Mathematical Programming*, 188(2):695–728, March 2021.
- [47] Santiago Gonzalez Zerbo, Alejandra Maestripieri, and Francisco Martínez Pería. Quadratic programming with one quadratic constraint in hilbert spaces. *Annals of Functional Analysis*, 16(2):31, 2025.
- [48] Jürgen Guddat, F Guerra Vazquez, and Hubertus Th Jongen. *Parametric optimization: singularities, pathfollowing and jumps*. Springer, 1990.
- [49] Martin Gugat and Michael Herty. Turnpike properties of optimal boundary control problems with random linear hyperbolic systems. *ESAIM: Control, Optimisation and Calculus of Variations*, 29:55, 2023.
- [50] Philipp A. Guth, Vesa Kaarnioja, Frances Y. Kuo, Claudia Schillings, and Ian H. Sloan. Parabolic PDE-constrained optimal control under uncertainty with entropic risk measure using quasi-Monte Carlo integration. *Numerische Mathematik*, 156(2):565–608, March 2024.
- [51] Mirko Hahn, Christian Kirches, Paul Manns, Sebastian Sager, and Clemens Zeile. Decomposition and approximation for PDE-constrained mixed-integer optimal control. In *Non-Smooth and Complementarity-Based Distributed Parameter Systems: Simulation and Hierarchical Optimization*, pages 283–305. Springer, 2021.
- [52] Mirko Hahn, Sven Leyffer, and Victor M Zavala. Mixed-integer PDE-constrained optimal control of gas networks. *Mathematics and Computer Science*, 113, 2017.
- [53] Andrew Hamann. Coordinated predictive control of a hydropower cascade. *Pittsburgh, Pennsylvania: Carnegie Mellon University*, 2015.
- [54] Nikolaus Hansen. *The CMA Evolution Strategy: A Comparing Review*, pages 75–102. Springer Berlin Heidelberg.

- [55] Matthias Heinkenschloss and Denis Ridzal. A matrix-free trust-region sqp method for equality constrained optimization. *SIAM Journal on Optimization*, 24(3):1507–1541, 2014.
- [56] Rainer Helmig. *Multiphase Flow and Transport Processes in the Subsurface: A Contribution to the Modeling of Hydrosystems*. Springer, Berlin, 1997.
- [57] Michael Hintermüller and Karl Kunisch. Pde-constrained optimization subject to pointwise constraints on the control, the state, and its derivative. *SIAM Journal on Optimization*, 20(3):1133–1156, 2010.
- [58] Michael Hinze, Rene Pinnau, Michael Ulbrich, Stefan Ulbrich, and Michael Ulbrich. Optimization methods in banach spaces. *Optimization with PDE Constraints*, pages 97–156, 2009.
- [59] John H. Holland. *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*. The MIT Press, April 1992.
- [60] Rui Huang, Victor M Zavala, and Lorenz T Biegler. Advanced step non-linear model predictive control for air separation units. *Journal of Process Control*, 19(4):678–685, 2009.
- [61] Oliver Jenkinson. Ergodic optimization in dynamical systems. *Ergodic Theory and Dynamical Systems*, 39(10):2593–2618, 2019.
- [62] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988.
- [63] James Kennedy. *Particle Swarm Optimization*, page 760–766. Springer US, 2011.
- [64] Morteza Kimiaei and Arnold Neumaier. Efficient unconstrained black box optimization. *Mathematical Programming Computation*, 14(2):365–414, February 2022.
- [65] Morteza Kimiaei and Arnold Neumaier. Effective matrix adaptation strategy for noisy derivative-free optimization. *Mathematical Programming Computation*, 16(3):459–501, July 2024.
- [66] Morteza Kimiaei and Arnold Neumaier. MATRS: heuristic methods for noisy derivative-free bound-constrained mixed-integer optimization. *Mathematical Programming Computation*, 17(3):505–546, May 2025.
- [67] Christian Kirches. *Fast numerical methods for mixed-integer nonlinear model-predictive control*. Springer, 2011.

- [68] Sanja Konjik, Michael Kunzinger, and Michael Oberguggenberger. Foundations of the calculus of variations in generalized function algebras. *Acta Applicandae Mathematicae*, 103(2):169–199, 2008.
- [69] Drew P Kouri, Matthias Heinkenschloss, Denis Ridzal, and Bart G van Bloemen Waanders. A trust-region algorithm with adaptive stochastic collocation for pde optimization under uncertainty. *SIAM Journal on Scientific Computing*, 35(4):A1847–A1879, 2013.
- [70] Drew P Kouri, Matthias Heinkenschloss, Denis Ridzal, and Bart G van Bloemen Waanders. Inexact objective function evaluations in a trust-region algorithm for PDE-constrained optimization under uncertainty. *SIAM Journal on Scientific Computing*, 36(6):A3011–A3029, 2014.
- [71] Drew P. Kouri, Mathias Staudigl, and Thomas M. Surowiec. A relaxation-based probabilistic approach for PDE-constrained optimization under uncertainty with pointwise state constraints. *Computational Optimization and Applications*, 85(2):441–478, February 2023.
- [72] Drew P Kouri and Thomas M Surowiec. Risk-averse PDE-constrained optimization using the conditional value-at-risk. *SIAM Journal on Optimization*, 26(1):365–396, 2016.
- [73] Vyacheslav Kungurtsev. Quasi-ergodic control of multi-periodic autoregressive processes: Formulation and examples. *arXiv preprint arXiv:2509.24729*, 2025.
- [74] Vyacheslav Kungurtsev and Moritz Diehl. Sequential quadratic programming methods for parametric nonlinear optimization. *Computational Optimization and Applications*, 59(3):475–509, 2014.
- [75] Vyacheslav Kungurtsev and Johannes Jaschke. A predictor-corrector path-following algorithm for dual-degenerate parametric optimization problems. *SIAM Journal on Optimization*, 27(1):538–564, 2017.
- [76] Alexander Kurganov and Guergana Petrova. A Second-Order Well-Balanced Positivity-Preserving Central-Upwind Scheme for the Saint-Venant System. *Communications in Mathematical Sciences*, 5(1):133–160, 2007.
- [77] Ailsa H. Land and Alison G. Doig. An automatic method of solving discrete programming problems. *Econometrica*, 28(3):497–520, 1960.
- [78] Randall J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge Texts in Applied Mathematics. Cambridge University Press, 2002.
- [79] Sven Leyffer and Paul Manns. Sequential linear integer programming for integer optimal control with total variation regularization. *ESAIM: Control, Optimisation and Calculus of Variations*, 28:66, 2022.

- [80] Sven Leyffer and Paul Manns. McCormick envelopes in mixed-integer pde-constrained optimization. *Mathematical Programming*, January 2025.
- [81] Wei Liu. Existence and uniqueness of solutions to nonlinear evolution equations with locally monotone operators. *Nonlinear Analysis: Theory, Methods & Applications*, 74(18):7543–7561, 2011.
- [82] Stefano Lucidi and Marco Sciandrone. A derivative-free algorithm for bound constrained optimization. *Computational Optimization and Applications*, 21(2):119–142, February 2002.
- [83] Berardi M., Difonzo F. V., and Guglielmi R. A preliminary model for optimal control of moisture content in unsaturated soils. *Computational Geosciences*, 27:1133–1144, 2023.
- [84] Berend Markhorst, Markus Leitner, Joost Berkhout, Alessandro Zocca, and Rob van der Mei. A two-step warm start method used for solving large-scale stochastic mixed-integer problems. *arXiv preprint arXiv:2412.10098*, 2024.
- [85] Matthieu Martin, Sebastian Krumscheid, and Fabio Nobile. Complexity analysis of stochastic gradient methods for pde-constrained optimal control problems with uncertain parameters. *ESAIM: Mathematical Modelling and Numerical Analysis*, 55(4):1599–1633, 2021.
- [86] Ali Mesbah. Stochastic model predictive control: An overview and perspectives for future research. *IEEE Control Systems Magazine*, 36(6):30–44, 2016.
- [87] Johannes Milz. Consistency of monte carlo estimators for risk-neutral PDE-constrained optimization. *Applied Mathematics Optimization*, 87(3), April 2023.
- [88] LG Mitten. Branch-and-bound methods: General formulation and properties. *Operations research*, 18(1):24–34, 1970.
- [89] Michael Oberguggenberger. Generalized solutions to semilinear hyperbolic systems. *Monatshefte für Mathematik*, 103(2):133–144, Jun 1987.
- [90] Claudio Paniconi and Mario Putti. Physically based modeling in catchment hydrology at 50: Survey and outlook. *Water Resources Research*, 51(9):7090–7129, 2015.
- [91] Carlos Parés. Numerical Methods for Nonconservative Hyperbolic Systems: A Theoretical Framework. *SIAM Journal on Numerical Analysis*, 44(1):300–321, 2006.
- [92] Shige Peng. Stochastic hamilton–jacobi–bellman equations. *SIAM Journal on Control and Optimization*, 30(2):284–304, 1992.

- [93] Sebastian Pfaff and Stefan Ulbrich. Optimal boundary control of nonlinear hyperbolic conservation laws with switched boundary data. *SIAM Journal on Control and Optimization*, 53(3):1250–1277, 2015.
- [94] Sebastian Pfaff and Stefan Ulbrich. Optimal control of nonlinear hyperbolic conservation laws by on/off switching with an application to traffic flow modeling, 2015.
- [95] Sebastian Pfaff and Stefan Ulbrich. Optimal control of nonlinear hyperbolic conservation laws by on/off-switching. Technical report, Department of Mathematics, TU Darmstadt, 2016.
- [96] Alois Pichler, Rui Peng Liu, and Alexander Shapiro. Risk-averse stochastic programming: Time consistency and optimal stopping. *Operations Research*, 70(4):2439–2455, 2022.
- [97] Warren B Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.
- [98] Alexander Poznyak, Sebastian Noriega-Marquez, Alejandra Hernandez-Sanchez, Mariana Ballesteros-Escamilla, and Isaac Chairez. Min-max dynamic programming control for systems with uncertain mathematical models via differential neural network bellman’s function approximation. *Mathematics*, 11(5):1211, March 2023.
- [99] Sri Redjeki Pudjaprasetya and Ikha Magdalena. Momentum conservative schemes for shallow water flows. *East Asian Journal on Applied Mathematics*, 4(2):152–165, 2014.
- [100] Rien Quirynen and Stefano Di Cairano. Sequential quadratic programming algorithm for real-time mixed-integer nonlinear mpc. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 993–999. IEEE, 2021.
- [101] James B Rawlings and Bhavik R Bakshi. Particle filtering and moving horizon estimation. *Computers & chemical engineering*, 30(10-12):1529–1541, 2006.
- [102] Rommel G. Regis and Christine A. Shoemaker. A stochastic radial basis function method for the global optimization of expensive functions. *INFORMS Journal on Computing*, 19(4):497–509, November 2007.
- [103] Eveline Rosseel and Garth N Wells. Optimal control with stochastic pde constraints and uncertain controls. *Computer Methods in Applied Mechanics and Engineering*, 213:152–167, 2012.
- [104] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 3rd edition, 2010.

- [105] Sebastian Sager. *Numerical methods for mixed-integer optimal control problems*. Der Andere Verlag Tönning, 2005.
- [106] Luca Santosuosso and Sonja Wogrin. Distributed stochastic model predictive control with temporal aggregation for the joint dispatch of cascaded hydropower and renewables. *arXiv preprint arXiv:2510.11998*, 2025.
- [107] Anton Schiela. Barrier methods for optimal control problems with state constraints. *SIAM Journal on Optimization*, 20(2):1002–1031, 2009.
- [108] Tobias Schwedes, David A Ham, Simon W Funke, and Matthew D Piggott. Mesh dependence in PDE-constrained optimisation. In *Mesh Dependence in PDE-Constrained Optimisation: An Application in Tidal Turbine Array Layouts*, pages 53–78. Springer, 2017.
- [109] Alexander Shapiro and Lingquan Ding. Periodical multistage stochastic programs. *SIAM Journal on Optimization*, 30(3):2083–2102, 2020.
- [110] Alexander Shapiro and Alois Pichler. Time and dynamic consistency of risk averse stochastic programs. *Optimization Online*, 2016.
- [111] M Soleimani-Damaneh. Optimality and invexity in optimization problems in banach algebras (spaces). *Nonlinear Analysis: Theory, Methods & Applications*, 71(11):5522–5530, 2009.
- [112] B Talaei, S Jagannathan, and J Singler. Adaptive dynamic programming boundary control of uncertain coupled semi-linear parabolic pde. In *2015 IEEE International Symposium on Intelligent Control (ISIC)*, pages 918–923. IEEE, 2015.
- [113] Behzad Talaei, Sarangapani Jagannathan, and J Singler. Boundary control of two dimensional burgers pde using approximate dynamic programming. In *2016 American Control Conference (ACC)*, pages 5243–5248. IEEE, 2016.
- [114] Behzad Talaei, Sarangapani Jagannathan, and John Singler. Boundary control of linear uncertain 1-d parabolic pde using approximate dynamic programming. *IEEE Transactions on Neural Networks and Learning Systems*, 29(4):1213–1225, 2018.
- [115] El-Ghazali Talbi. *Metaheuristics: From Design to Implementation*. Wiley, 2009.
- [116] Anna Thünen, Sven Leyffer, and Sebastian Sager. State elimination for mixed-integer optimal control of partial differential equations by semi-group theory. *Optimal Control Applications and Methods*, 43(4):867–883, 2022.
- [117] Virginia Torczon. On the convergence of pattern search algorithms. *SIAM Journal on Optimization*, 7(1):1–25, February 1997.

- [118] Eleuterio F. Toro. *Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction*. Springer-Verlag, 3rd edition, 2009.
- [119] Stefan Ulbrich. A sensitivity and adjoint calculus for discontinuous solutions of hyperbolic conservation laws with source terms. *SIAM journal on control and optimization*, 41(3):740–797, 2002.
- [120] Stefan Ulbrich. 7. *Generalized SQP Methods with “Parareal” Time-Domain Decomposition for Time-Dependent PDE-Constrained Optimization*, page 145–168. Society for Industrial and Applied Mathematics, January 2007.
- [121] Stefan Ulbrich. Generalized SQP methods with “parareal” time-domain decomposition for time-dependent PDE-constrained optimization. In *Real-time PDE-constrained optimization*, pages 145–168. SIAM, 2007.
- [122] Sifan Wang, Mohamed Aziz Bhouiri, and Paris Perdikaris. Fast PDE-constrained optimization via self-supervised operator learning. *arXiv preprint arXiv:2110.13297*, 2021.
- [123] Zhangping Wei and Yafei Jia. A depth-integrated non-hydrostatic finite element model for wave propagation. *International Journal for Numerical Methods in Fluids*, 73(11):976–1000, 2013.
- [124] Maximilian Willmann, Jesus Carrera, and Xavier Sánchez-Vila. Model Reduction for Groundwater Flow and Transport. *Advances in Water Resources*, 33(11):135–144, 2010.
- [125] Eberhard Zeidler. *Nonlinear Functional Analysis and It’s Applications: Fixed-point theorems*. Springer-Verlag, 1986.
- [126] Clemens Ulrich Zeile. *Combinatorial integral decompositions for mixed-integer optimal control*. PhD thesis, 2021.
- [127] Ying Zheng, Xudong Fu, and Jiahua Wei. Evaluation of power generation efficiency of cascade hydropower plants: a case study. *Energies*, 6(2):1165–1177, 2013.