Preface

Many engineering and scientific problems in design, control, and parameter estimation can be formulated as optimization problems that are governed by partial differential equations (PDEs). The size and complexity of the PDE simulations often present significant optimization challenges. Recent years have seen sustained progress in PDE solvers and optimization algorithms, and the rapid rise in computing capability. Accompanying these advances is a growing interest in real-time and online simulation-based optimization in such diverse areas as aerodynamics, atmospheric and geosciences, chemical process industry, environment, homeland security, infrastructure, manufacturing, and medicine.

Real-time, online optimization arises in the form of: a) *inverse problems*, in which sensor data is assimilated repeatedly into simulations of dynamic processes, b) *control problems*, in which optimal strategies are repeatedly generated based on new data and c) *design problems*, in which algorithms need to be configured to aid in the solution, as well as coordination, of components a) and b). An important underlying aspect to online optimization is the availability of powerful large-scale optimization algorithms that can target time-dependent objectives and constraints on challenging engineering systems.

The rapid solution of an optimization problem, however, is only one component needed for the successful realization of real-time optimization. The challenges for real-time and online optimization methods include the ability to:

- *Run sufficiently quickly for decision-making at relevant time scales.* Optimization algorithms must return acceptable approximations of the solution in time for decision-making. The time scales for decision-making vary by application and can range from milliseconds to days. Reliability and efficiency of the optimization algorithms are crucial. Their implementations must be capable of bootstrapping current solutions and yielding meaningful results when terminated prematurely.

- *Adjust to requirement and properties of the process.* Implementations of online optimization algorithms must adjust to different solution accuracy requirements in the process model they feed into. They must be robust and deal with ill-posedness of the problems and tolerate incomplete, uncertain, or errant data.

- *Scale to large problem sizes.* Applications for which real-time and online optimization are required or desired increase both in number and in complexity. In particular, many of these applications require modeling by PDEs instead of lumped parameter models. Consequently, online optimization algorithms must be able to handle increasingly large and complex PDE models and they have to be able to effectively utilize parallel computing environments.

This volume addresses these issues for real-time optimization of systems governed by PDEs. With fifteen contributions, it complements two recent volumes in related areas:

- *Online Optimization of Scale Large Systems*, M. Grötschel, S. Krumke, J. Rambau (eds.), Springer Verlag, Berlin (2001). Comprising the work of 25 research projects over six years, this volume deals with a broad array of real-time optimization strategies and applications. In addition to online optimization of systems of differential equations, this volume also considers concepts related to nonlinear and mixed integer optimization, along with stochastic programming and online planning.

- *Large-scale PDE-constrained Optimization*, Lecture Notes in Computational Science and Engineering, Vol. 30, L.T. Biegler, O. Ghattas, M. Heinkenschloss and B. van Bloemen Waanders (eds.), Springer Verlag, Berlin (2003). This volume deals with a variety of algorithms and strategies for optimization of systems governed by PDEs. The range of PDE applications is broad, but largely confined to design and off-line optimization.

This volume focuses on the interface of these two areas by considering inversion, control, and design problems that are governed by PDEs and have online or real-time requirements. To this end, the fifteen contributions in this volume address the following essential components for real-time optimization:

- Part I: Concepts and properties of real-time, online strategies
- Part II: Fast PDE-constrained optimization solvers
- Part III: Reduced order modeling within the context of design and control
- Part IV: Applications for real-time decision-making.

Concepts and Properties of Real-time, Online Strategies

The chapters in Part I of this volume address core problems of online optimization of systems governed by differential equations. Key issues in all of these chapters are the formulation of online problems as well as the interaction of the model, solver and the "plant" (i.e. the real-world system described by the PDEs). Important considerations are online stability properties of the closed loop system, robustness to uncertainties introduced by the plant as well as the solution algorithm, and efficient implementation of the computations. These issues are considered for differential equation models in the following chapters.

1. "Constrained Optimal Feedback Control of Systems Governed by Large Differential Algebraic Equations," H.-G. Bock, M. Diehl, E. Kostina and J.P. Schlöder

2. "A Stabilizing Real-time Implementation of Nonlinear Model Predictive Control," M. Diehl, R. Findeisen, and F. Allgöwer

3. "Numerical Feedback Controller Design for PDE Systems using Model Reduction: Techniques and Case Studies," F. Leibfritz and S. Volkwein

4. "A Least-squares Finite Element Method for Optimization and Control Problems," P.B. Bochev and M.D. Gunzburger

Chapter 1 presents an overview of recent work in nonlinear model predictive control (NMPC) by Bock and coworkers. Here the authors derive NMPC controllers based on a multiple shooting formulation. However, a number of approximations are made so that they do not require the online calculation of (expensive) sensitivity matrices. The authors develop four variants of the NMPC controller and also provide convergence proofs for their optimization strategies.

Chapter 2 by Diehl, Findeisen, and Allgöwer presents an overview of recent work in nonlinear model predictive control (NMPC) based on the dual mode control concept and stability analysis. The authors then discuss a novel real-time iteration scheme that allows incomplete solution of the dynamic optimization problem. The stability analysis is modified by coupling the dual mode approach with contractive properties of Newton's method. A experimental case study on the NMPC control of a distillation column, operating at the University of Stuttgart, is presented.

Chapter 3 by Leibfritz and Volkwein presents a set of controller designs based on static output feedback (SOF) along with linear matrix inequality (LMI) constraints to enforce stability. The controller design is then applied to reduced order models of distributed parameter systems derived using proper orthogonal decomposition (POD). The resulting controller design problems are nonconvex semi-definite programs and are solved using an interior point trust region algorithm. Three systems are considered and alternative controller designs are presented.

Chapter 4 by Bochev and Gunzburger studies problem formulation and discretization for a class of PDE constrained optimization problems. Unlike penalty methods, which enforce the PDE constraints by using well-posed least-squares penalty terms added to the original cost functional, this chapter discusses a new approach in which the cost functional is constrained by the least-squares formulation of the PDE constraints. This leads to a bilevel optimization problem. Characterization of solutions and convergence properties of finite element approximations are discussed. The approach is illustrated on a linear quadratic optimal control problem governed by the Stokes equations.

Fast PDE-constrained Optimization Solvers

Fast optimization solvers remain at the core of online optimization, and recent advances in this area are represented in the five chapters of Part 2 of this volume. In particular, optimization strategies for time-dependent PDE applications are considered. The first three chapters discuss the development of solution strategies that extend multigrid and domain decomposition methods to time-dependent PDE optimization problems, and also leverage parallel architectures.

5. "Space-time Multigrid Methods for Solving Unsteady Optimal Control Problems," A. Borzì

6. "A Time-parallel Implicit Methodology for the Near-real-time Solution of Systems of Linear Oscillators," J. Cortial and C. Farhat

7. "Generalized SQP Methods with 'Parareal' Time Domain Decomposition for Time-dependent PDE-constrained Optimization" S. Ulbrich

Chapter 5 by Borzi gives an overview of space-time multigrid methods for real-time optimal control of parabolic systems, with an application to reaction-diffusion problems. Here reaction diffusion optimal control problems are approximated by finite differences along with backward Euler schemes, and space-time multigrid schemes are formulated to solve the resulting discrete optimality systems, using a full approximation storage framework. For long time intervals this approach is enhanced with the development of a multigrid receding horizon algorithm and numerical experiments validate the numerical performance of the space-time multigrid algorithms. Finally, the optimal control of cardiac arrhythmia is presented as an example.

Chapter 6 by Cortial and Farhat treats new methods for time-parallelism in the solution of PDEs. Here, time slices are generated with a domain-decomposition in the time direction. The solution is first approximated on this coarse time-grid to provide an initial condition to each time slice. Then, the ODE solver is applied independently and therefore concurrently in each time slice to advance the solution from the starting point of this time slice to its end point. Finally, an iterative process is invoked to improve the solution and eliminate the jumps of the solution on the coarse time-grid. This approach lead to development of two related approaches, the *parareal* (parallel real-time) algorithm and the *PITA* (parallel implicit time-integrator) method. Here a new extension is presented to improve the stability properties of both methods. Favorable numerical results are demonstrated for linear hyperbolic systems.

Chapter 7 by S. Ulbrich extends the recent *parareal* time-domain decomposition concept to PDE-constrained optimization problems. The underlying optimization algorithm is a generalized sequential quadratic programming (SQP) method. Unlike traditional SQP methods that require the solution of the linearized state equation, the generalized SQP method studied in this chapter allows the incorporation of approximate state equation solvers. Here parareal time-domain decomposition is applied to approximately solve the state as well as the adjoint equation system. Global convergence of the generalized SQP method is shown, and the method is applied to the optimal control of a semilinear parabolic equation. Numerical tests demonstrate the efficiency of the approach.

The following two chapters deal with essential improvements to accelerate the optimization of PDE-constrained problems.

- "Simultaneous Pseudo-Timestepping for State Constrained Optimization Problems in Aerodynamics," S.B. Hazra, V. Schulz

- Digital Filter Stepsize Control in DASPK and its Effect on Control Optimization Performance," K. Meeker, C. Homescu, L. Petzold, H. El-Samad, M. Khammash and G. Söderlind

Chapter 8 by Hazra and Schulz presents a method for solving aerodynamic shape optimization problems involving state constraints, based on simultaneous pseudo-timestepping. This approach solves the KKT system by evolving an embedded non-stationary system of equations in pseudo-time. Advantages of this method include no additional globalization techniques in the design space, and direct extensions to a previously existing pseudo-timestepping method for the states only. Moreover, a block preconditioner that stems from reduced SQP methods can be used for convergence acceleration. The approach is demonstrated on two airfoil application examples to reduced drag with constant lift. Results show a significant reduction of the computational cost compared to usual gradient methods that solve the state and adjoint PDEs exactly at each optimization iteration.

Finally, Chapter 9 by Meeker et al. deals with the issue of nonsmoothness of adaptive ODE and DAE solvers, with respect to perturbations in initial conditions or other problem parameters. Here a recently developed digital filter stepsize controller leads to a much smoother dependence of the solution on problem parameters along with better convergence characteristics for the control and optimization of dynamical systems. This chapter discusses the implementation of the digital filter stepsize controller in the DAE solver DASPK3.1. The improved performance of the optimizer, due to the new stepsize controller, is demonstrated on a biological problem regarding the heat shock response of *Escherichia coli*.

Reduced Order Modeling

An important consideration for online optimization is an accurate representation of the real-world system. Often, the online solution of a detailed plant model, even with efficient state-of-the-art optimization algorithms, is precluded by the computational effort engendered by the governing PDEs. This barrier can often be overcome by the application of reduced order models. Here a key factor is the development of reduced order models that capture the accuracy needed for online purposes over a range of input parameters. Clearly, a trade-off occurs between the accuracy achieved by the reduced order model and the effort needed to solve it. The following chapters deal with the generation of accurate reduced order models and demonstrate their potential for online optimization.

10. "Certified Rapid Solution of Partial Differential Equations for Real-Time Parameter Estimation and Optimization," M.A. Grepl, N.C. Nguyen, K. Veroy, A.T. Patera, and G.R. Liu

11. "Model Reduction for Large-scale Applications in Computational Fluid Dynamics," K. Willcox

12. "Suboptimal Feedback Control of Flow Separation by POD Model Reduction," K. Kunisch and L. Xie

Chapter 10 by Grepl et al. derives reduced-order approximations to input-output functions (representing PDE-based systems) with certified accuracy. Here the focus is on processes operating with associated *Assess-Act* scenarios. In the *Assess* stage, robust parameter estimation (inverse) procedures are pursued that map measured-observable outputs to (all) possible system-characteristic and environment-state inputs. In the subsequent *Act* stage, adaptive design and optimization procedures are considered that map mission-objective outputs to best control-variable inputs. The chapter describes a method for real-time certified evaluation of PDE input-output relations. The two ingredients are reduced-basis approximation as well as *a posteriori* error estimation. Both linear elliptic second-order PDEs as well as extensions to certain nonlinear and parabolic equations are considered.

Chapter 11 by Willcox reviews model reduction of computational fluid dynamics systems and compares two reduction methods, proper orthogonal decomposition (POD) and

Fourier model reduction (FMR) in the context of an aerodynamic application, i.e., active flow control for a supersonic diffuser. In particular for linear systems, FMR has advantages in the computation of rigorous, though noncomputable error bounds, and the development of an efficient reduction process when FMR is combined with balanced truncation. Numerical results show significant improvements in approximation over the POD approach. Open questions related to applications of FMR for nonlinear systems are also discussed.

Finally, Chapter 12 by Kunisch and Xie addresses optimal control of fluid dynamical systems using POD applied to the solution of the Hamilton-Jacobi-Bellman (HJB) equations. Both optimal control and boundary feedback controls are developed for control problems of Navier-Stokes flows. The approach is implemented in parallel on a cluster, by using a message passing interface to exchange data between MATLAB sessions. This parallel implementation reduces the computing time to an acceptable level for online use.

Applications

The chapters in this final section focus on applications related to online optimization and decision-making. While all chapters in this volume incorporate applications to demonstrate the effectiveness of their approaches, the following chapters emphasize problem formulations and solution strategies that are tailored to a particular application. All three chapters deal with inverse problems, which is a core element of real-time optimization and control. The first two chapters are concerned with imaging for medical applications, while the third deals with source detection in water networks.

13. "A Combined Shape-Newton and Topology Optimization Technique in Real-time Image Segmentation," M. Hintermüller

14. "Cofir: Coarse and Fine Image Registration," J. Modersitzki and E. Haber

15. "Real-time, Large Scale Optimization of Water Network Systems using a Subdomain Approach," C. D. Laird, L. T. Biegler, B. G. van Bloemen Waanders

Chapter 13 by Hintermüller addresses the determination of tumor boundaries from high resolution magnetic resonance images (MRIs). For this purpose, an edge-detector based image segmentation algorithm is presented that utilizes both shape as well as topological sensitivity information. Here shape sensitivity and topological derivatives are employed in a serial fashion and no synchronous blending of information takes place. Typically, the algorithm is initiated by the topological phase for computing initial guesses and, if necessary, shape sensitivity is used to drive the contour. A benefit of this approach is that the segmentation becomes fully automatic, without manual selection of an initial contour. This reduction of required user interaction is an important feature in real-time applications. The approach is demonstrated on a number of MRI examples.

Chapter 14 by Modersitzki and Haber presents a general framework for image registration based on a variational principle. Image registration seeks to find a transformation between two related images. Compared to other variational registration schemes, the approach has two new features. First, the transformation space is separated into coarse and fine spaces, orthogonal complements, and only the fine part of the transformation is regularized. Second, a simultaneous solution scheme is developed for both parts. Compared to a multi-level sequential scheme, this can lead to the avoidance of local minima without added computational cost. These features are significant for real-time clinical applications. The approach is applied to registration of a magnetic resonance image for a human knee.

Chapter 15 by Laird et al. describes a subdomain approach to the problem of source detection to detect contaminants in municipal water networks. Previous work led to efficient and reliable optimization formulations for this problem. The chapter adapts this approach to online implementations and considers a, possibly moving, window of sensor measurements along with problem formulations that can be modified over time. Because the solution of these online optimization problems is time-critical, the moving window approach allows the consideration of very large networks while still maintaining a manageable problem size. The approach is demonstrated on a water distribution network characteristic of a large city.

Conclusions and Acknowledgments

This volume summarizes recent advances toward the solution of time-critical PDE-constrained optimization problems that arise in online implementations. Addressing key issues of online formulations, fast algorithms, reduced order models, and applications, these chapters deal with essential ingredients for realization of real-time optimization for challenging applications. They also pose a number of open questions that will spark future work in this exciting area.

Finally, we wish to acknowledge the Computer Science Research Institute at Sandia National Laboratories; National Science Foundation grants 0121360, 0121667, 0314647, 0326449, 0352334, 0427464, 0427985, 0438279, 0457379, 0511624, 0520280, and 0540372; and Department of Energy grant DE-FC02-01ER25477, for financial support of this project.

L.T. Biegler
O. Ghattas
M. Heinkenschloss
D.E. Keyes
B. van Bloemen Waanders