PARALLEL LAGRANGE–NEWTON–KRYLOV–SCHUR METHODS FOR PDE-CONSTRAINED OPTIMIZATION. PART I: THE KRYLOV–SCHUR SOLVER*

GEORGE BIROS† AND OMAR GHATTAS‡

Abstract. Large-scale optimization of systems governed by partial differential equations (PDEs) is a frontier problem in scientific computation. Reduced quasi-Newton sequential quadratic programming (SQP) methods are state-of-the-art approaches for such problems. These methods take full advantage of existing PDE solver technology and parallelize well. However, their algorithmic scalability is questionable; for certain problem classes they can be very slow to converge. In this two-part article we propose a new method for steady-state PDE-constrained optimization, based on the idea of using a full space Newton solver combined with an approximate reduced space quasi-Newton SQP preconditioner. The basic components of the method are Newton solution of the first-order optimality conditions that characterize stationarity of the Lagrangian function; Krylov solution of the Karush–Kuhn–Tucker (KKT) linear systems arising at each Newton iteration using a symmetric quasi-minimum residual method; preconditioning of the KKT system using an approximate state/decision variable decomposition that replaces the forward PDE Jacobians by their own preconditioners, and the decision space Schur complement (the reduced Hessian) by a BFGS approximation initialized by a two-step stationary method. Accordingly, we term the new method *Lagrange–Newton–Krylov–Schur* (LNKS). It is fully parallelizable, exploits the structure of available parallel algorithms for the PDE forward problem, and is locally quadratically convergent. In part I of this two-part article, we investigate the effectiveness of the KKT linear system solver. We test our method on two optimal control problems in which the state constraints are described by the steady-state Stokes equations. The objective is to minimize dissipation or the deviation from a given velocity field; the control variables are the boundary velocities. Numerical experiments on up to 256 Cray T3E processors and on an SGI Origin 2000 include scalability and performance assessment of the LNKS algorithm and comparisons with reduced SQP for up to 1,000,000 state and 50,000 decision variables. In part II of the article, we address globalization and inexactness issues, and apply LNKS to the optimal control of the steady incompressible Navier–Stokes equations.

Key words. sequential quadratic programming, adjoint methods, PDE-constrained optimization, optimal control, Lagrange–Newton–Krylov–Schur methods, Navier–Stokes, finite elements, preconditioners, indefinite systems, nonlinear equations, parallel algorithms

AMS subject classifications. 49K20, 65F10, 65K05, 65K10, 65J22, 65N55, 65W10, 65Y05, 65Y20, 76D05, 76D07, 76D55, 90C52, 90C55, 90C90, 93C20

DOI. 10.1137/S106482750241565X

1. Introduction. *PDE-constrained optimization* refers to the optimization of systems governed by partial differential equations (PDEs). The *forward* or *state problem* is to solve the PDEs for the *state variables*, given appropriate data, i.e., geometry, coefficients, boundary conditions, initial conditions, and source functions. The *optimization problem* seeks to determine some of these data—the *decision variables*—given an *objective function* (e.g., performance goals) and (possibly inequality) constraints

†Courant Institute of Mathematical Sciences, Department of Computer Science, New York University, New York, NY 10012 (biros@cs.nyu.edu).

‡Ultrascale Simulation Laboratory, Departments of Biomedical Engineering and Civil & Environmental Engineering, Carnegie Mellon University, Pittsburgh, PA 15213 (oghattas@cs.cmu.edu).

on the behavior of the system. Since the behavior of the system is modeled by the PDEs, they appear as constraints in the optimization problem.

PDE-constrained optimization problems arise in several contexts, including optimal design, optimal control, and parameter estimation. They share the common difficulty that PDE solution is just a subproblem associated with optimization, and that the optimization problem can be ill-posed even when the forward problem is well-posed. Thus the optimization problem is often significantly more difficult to solve than the forward problem.

To illustrate the main issues, let us consider a model problem of optimal distributed control of a steady Navier–Stokes flow:

$$\min \mathcal{F}(\boldsymbol{u}, p, \boldsymbol{b}) = \frac{1}{2} \int_{\Omega} (\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T) \cdot (\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T) d\Omega + \frac{1}{2} \int_{\Omega} \boldsymbol{b} \cdot \boldsymbol{b} \, d\Omega$$

subject to

$$-\nu \nabla \cdot (\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T) + (\nabla \boldsymbol{u})\boldsymbol{u} + \nabla p + \boldsymbol{b} = \boldsymbol{0} \quad \text{in } \Omega,$$
$$\nabla \cdot \boldsymbol{u} = \boldsymbol{0} \quad \text{in } \Omega,$$
$$\boldsymbol{u} = \boldsymbol{g} \quad \text{on } \Gamma.$$

Here $\boldsymbol{u}$ is the fluid velocity field, p the pressure field, $\boldsymbol{b}$ the body force control function, $\boldsymbol{g}$ the prescribed Dirichlet boundary conditions, and ν the inverse of the Reynolds number. The objective is to minimize the rate of dissipation of viscous energy without applying an excessively large body force. The constraints are the stationary incompressible Navier–Stokes equations with Dirichlet boundary conditions. The decision variable is the body force function $\boldsymbol{b}$.

For this problem a Lagrangian functional is given by

$$\frac{1}{2} \int_{\Omega} (\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T) \cdot (\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T) d\Omega + \frac{1}{2} \int_{\Omega} \boldsymbol{b} \cdot \boldsymbol{b} \, d\Omega$$
$$+ \int_{\Omega} \boldsymbol{\lambda} \cdot (-\nu \nabla \cdot (\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T) + (\nabla \boldsymbol{u})\boldsymbol{u} + \nabla p + \boldsymbol{b}) d\Omega + \int_{\Omega} \mu \nabla \cdot \boldsymbol{u} \, d\Omega.$$

By requiring its stationarity with respect to the state $(\boldsymbol{u}, p)$, decision $(\boldsymbol{b})$, and Lagrange multiplier (i.e., adjoint) variables $(\boldsymbol{\lambda}, \mu)$, taking the respective variations, and invoking the appropriate Green's identities, we arrive at the (strong form of) the so-called *first order necessary conditions for optimality*, or the *Karush–Kuhn–Tucker* conditions.

State equations:

$$-\nu \nabla \cdot (\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T) + (\nabla \boldsymbol{u})\boldsymbol{u} + \nabla p + \boldsymbol{b} = \boldsymbol{0} \quad \text{in } \Omega,$$
$$\nabla \cdot \boldsymbol{u} = \boldsymbol{0} \quad \text{in } \Omega,$$
$$\boldsymbol{u} = \boldsymbol{g} \quad \text{on } \Gamma.$$

Adjoint equations:

$$-\nu \nabla \cdot (\nabla \boldsymbol{\lambda} + \nabla \boldsymbol{\lambda}^T) + (\nabla \boldsymbol{u})^T \boldsymbol{\lambda} - (\nabla \boldsymbol{\lambda})\boldsymbol{u} + \nabla \mu - \nabla \cdot (\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T) = \boldsymbol{0} \quad \text{in } \Omega,$$
$$\nabla \cdot \boldsymbol{\lambda} = \boldsymbol{0} \quad \text{in } \Omega,$$
$$\boldsymbol{\lambda} = \boldsymbol{0} \quad \text{on } \Gamma.$$

Control equations:

$$\boldsymbol{b} + \boldsymbol{\lambda} = \boldsymbol{0} \quad \text{in } \Omega.$$

The first set of equations are just the original Navier–Stokes PDEs. The *adjoint* equations, which result from stationarity with respect to state variables, are themselves PDEs, and are linear in the Lagrange multipliers $\boldsymbol{\lambda}$ and μ. Finally, the *control equations* are (in this case) algebraic.

The above optimality system is a large, nonlinear, coupled, unstructured system of PDEs (or at least larger, more coupled, and less structured than usually seen by the forward PDE solver). Solving this system presents significant challenges beyond the forward PDE. One way to mitigate some of the difficulties is to eliminate state variables and Lagrange multipliers—and correspondingly, the state equations and adjoint equations—to reduce the system to a manageable one in just the decision variables $\boldsymbol{b}$. This can be done as follows: given $\boldsymbol{b}$ at some iteration, we solve the state equations for the state variables $\boldsymbol{u}, p$. Knowing the state variables then permits us to solve the adjoint equations for the Lagrange multipliers $\boldsymbol{\lambda}, \mu$. Finally, with the states and multipliers known, we can update $\boldsymbol{b}$ by taking a Newton-like iteration on the control equation. The entire process is repeated until convergence. This elimination procedure is termed a *reduced space* method, in contrast to a *full space* method, in which one solves for the states, controls, and multipliers simultaneously. The variable reduction can be done either prior to or after the Newton linearization. The former is a type of *nonlinear elimination* and is known as a *generalized reduced gradient method*, and the latter is a *linear elimination* or *reduced sequential quadratic programming method* (SQP).

Reduced space methods are attractive for several reasons. Solving the subsets of equations in sequence imparts some structure to the optimization problem. State equation solvers build on years of development of large-scale parallel PDE solvers. The strong similarities between the state and adjoint operators suggest that an existing PDE solver for the state equations can often be modified readily to handle the adjoint system. Finally, the control equations are often reasonably straightforward, at least to evaluate. Another advantage of reduction is that the full space system is often very ill-conditioned, whereas the three subsystems are typically better conditioned individually.

Given these advantages, the overwhelming popularity of reduced space methods, for optimization of systems governed by PDEs is not surprising. We will not attempt to survey such methods, but rather we give a sample of the literature on applications to various PDE problems: applications to compressible flow airfoil design [33, 40]; heat equation boundary control [28]; inverse parameter estimation [15, 27]; Navier–Stokes flow control [19]; inverse acoustic scattering [16]; chemically reacting flows [39]; and structural optimization [34, 36, 37]. In addition, parallel implementations of reduced sequential quadratic programming (RSQP) methods exhibiting high parallel efficiency and good scalability have been developed [16, 20, 33, 35, 29].

On the other hand, the major disadvantage of reduced methods is the requirement to solve the (possibly linearized) state and adjoint equations *at each iteration* of the reduced system—which is a direct consequence of the reduction onto the decision variable space. Because of this difficulty it is natural to return to the full space and ask if it is possible to solve simultaneously the entire optimality system but retain the structure-inducing, condition-improving advantages of reduced space methods—while avoiding their disadvantages.

In this two-part article we present such a method. The key idea is to solve in the full space using a Newton method, but precondition with a quasi-Newton reduced space method. We refer to the (nonlinear) Newton iterations as *outer* iterations, and use the term *inner* to describe the (linear) Krylov iterations for the Karush–Kuhn–Tucker (KKT) system that arises at each Newton iteration. This linearized KKT

system is solved using a Krylov iterative method (symmetric QMR), and it is this system to which the preconditioner is applied. The reduced space preconditioner is very effective in deflating the spectrum of the KKT matrix and in reducing the number of Krylov iterations—applying it captures the favorable structure of reduced methods. On the other hand, since the reduction is used just as a preconditioner, we can approximate the state and adjoint solves, replacing them for example by their own preconditioners. So we arrive at a method that combines rapid convergence in the outer Newton iteration (typically mesh-independent) with fast convergence of the inner Krylov iteration (which can be as good as mesh-independent).

The new method is termed *Lagrange–Newton–Krylov–Schur*. It is common in the PDE-solver literature to use the phrase *Newton–Krylov–X* to refer to Newton methods for solving PDEs that employ Krylov linear solvers, with *X* as the preconditioner for the Krylov method. Since *Lagrange–Newton* is used sometimes to describe a Newton method for solving the optimality system (i.e., an SQP method), and since a reduced space method can be viewed as a Schur complement method for the KKT system, we arrive at the concatenation *LNKS*.

The LNKS method is inspired by domain-decomposed Schur complement PDE solvers. In such techniques, reduction onto the interface space requires exact subdomain solves, so one often prefers to iterate within the full space while using a preconditioner based on approximate subdomain solution [26]. In our case, the decomposition is into states and decisions, as opposed to subdomain and interface spaces.

An early version of an LNKS method was presented in [9] and [10]. Here we extend the method in a number of directions. In part I, we analyze and compare several different variants of the KKT preconditioner. Scalability to larger problems and processor counts is studied. In part II, important globalization and inexactness issues, which were not considered in the earlier work, are introduced and analyzed. We also apply the method to some viscous flow control problems.

Haber and Asher [24] have proposed a KKT preconditioner similar to [9] and [10]. The main differences are in the reduced Hessian preconditioner and the application to inverse problems. Battermann and Heinkenschloss [7] have presented a related KKT-system preconditioner that also makes use of state and decision space decompositions. However, their preconditioner is based on congruence transformations of the original system and not on its exact factorization. The resulting preconditioned system has both positive and negative eigenvalues and its spectrum is less favorably distributed. Another important difference is that we precondition in the reduced space. We also consider parallelism and scalability issues.

Part I of the article is organized as follows: in section 2 we discuss SQP methods and in particular reduced space variants; in section 3 we introduce the LNKS method and present several approaches for preconditioning the KKT matrix; in section 4 we examine the formulation of a Stokes flow problem; and in section 5 we conclude with numerical results, and study parallel and algorithmic scalability of the method.

Note on notation. We use boldface characters to denote vector-valued functions and vector-valued function spaces. We use roman characters to denote discretized quantities and italics for their continuous counterparts. For example, $\boldsymbol{u}$ will be the continuous velocity field and $\mathbf{u}$ will be its discretization. Greek letters are overloaded and whether we refer to the discretization or the continuous fields should be clear from context. We also use $+$ as a subscript or superscript to denote variable updates within an iterative algorithm.

2. Reduced space methods. We begin with a constrained optimization problem formulated as follows:

$$(2.1) \qquad \min_{\mathbf{x} \in \mathbb{R}^N} f(\mathbf{x}) \quad \text{subject to } \mathbf{c}(\mathbf{x}) = \mathbf{0},$$

where $\mathbf{x} \in \mathbb{R}^N$ are the optimization variables, $f : \mathbb{R}^N \to \mathbb{R}$ is the objective function, and $\mathbf{c} : \mathbb{R}^N \to \mathbb{R}^n$ are the constraints. We assume that the constraints consist of just the (discretized) state equations.[1] In order to exploit the structure of the PDEs, we partition $\mathbf{x}$ into the state variables $\mathbf{x}_s \in \mathbb{R}^n$ and the decision variables $\mathbf{x}_d \in \mathbb{R}^m$,

$$(2.2) \qquad \mathbf{x} = \left\{ \begin{array}{c} \mathbf{x}_s \\ \mathbf{x}_d \end{array} \right\},$$

so that $N = m + n$. By introducing the Lagrangian function $\mathcal{L}$ one can derive first- and higher-order optimality conditions. The Lagrangian is defined by

$$(2.3) \qquad \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) := f(\mathbf{x}) + \boldsymbol{\lambda}^T \mathbf{c}(\mathbf{x}),$$

where $\boldsymbol{\lambda} \in \mathbb{R}^n$ are the Lagrange multipliers (or adjoint variables). The first-order optimality conditions state that at a local minimum the Lagrangian gradient must vanish:

$$(2.4) \qquad \left\{ \begin{array}{c} \partial_x \mathcal{L} \\ \partial_\lambda \mathcal{L} \end{array} \right\} (\mathbf{x}, \boldsymbol{\lambda}) = \left\{ \begin{array}{c} \partial_x f(\mathbf{x}) + (\partial_x \mathbf{c}(\mathbf{x}))^T \boldsymbol{\lambda} \\ \mathbf{c}(\mathbf{x}) \end{array} \right\} = \mathbf{0}.$$

Points at which the gradient of the Lagrangian vanishes are often called *KKT points*. Equations (2.4) are typically known as the KKT optimality conditions. To simplify the notation further, let us define

$$\begin{array}{lll} \mathbf{A}(\mathbf{x}) := \partial_x \mathbf{c}(\mathbf{x}) & \in \mathbb{R}^{n \times N} & \text{Jacobian matrix of the constraints,} \\ \mathbf{W}(\mathbf{x}, \boldsymbol{\lambda}) := \partial_{xx} f(\mathbf{x}) + \sum_i \lambda_i \partial_{xx} \mathbf{c}_i(\mathbf{x}) & \in \mathbb{R}^{N \times N} & \text{Hessian matrix of the Lagrangian,} \\ \mathbf{g}(\mathbf{x}) := \partial_x f(\mathbf{x}) & \in \mathbb{R}^N & \text{gradient vector of the objective.} \end{array}$$

Consistent with the partitioning of the optimization variables into states and decisions, we logically partition $\mathbf{g}$, $\mathbf{A}$, $\mathbf{W}$ as follows:[2]

$$\mathbf{g} = \left\{ \begin{array}{c} \mathbf{x}_s \\ \mathbf{x}_d \end{array} \right\}, \quad \mathbf{A} = \begin{bmatrix} \mathbf{A}_s & \mathbf{A}_d \end{bmatrix}, \quad \text{and} \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}_{ss} & \mathbf{W}_{sd} \\ \mathbf{W}_{ds} & \mathbf{W}_{dd} \end{bmatrix}.$$

At each iteration of the SQP method a quadratic programming problem (QP) is solved to generate a new search direction $\mathbf{p}_x$. When SQP is derived from a Newton method for solving the optimality conditions (2.4), the resulting QP is of the form

$$(2.5) \qquad \min_{\mathbf{p}_x \in \mathbb{R}} \frac{1}{2} \mathbf{p}_x^T \mathbf{W} \mathbf{p}_x + \mathbf{g}^T \mathbf{p}_x \quad \text{subject to } \mathbf{A} \mathbf{p}_x + \mathbf{c} = \mathbf{0}.$$

---

[1] However, the methodology can be extended to problems that include additional equality or inequality constraints, e.g., through the use of a partially reduced SQP method [41] as a preconditioner.

[2] All vectors and matrices depend on the optimization variables $\mathbf{x}$; in addition, the Hessian $\mathbf{W}$ depends also on the Lagrange multipliers $\boldsymbol{\lambda}$. For clarity, we suppress identification of these dependencies.

Reduced space methods for solving (2.5) eliminate the linearized constraint by using the null and range space decomposition of the search direction,

$$(2.6) \qquad\qquad \mathbf{p}_x = \mathbf{Z}\mathbf{p}_z + \mathbf{Y}\mathbf{p}_y,$$

where the columns of $\mathbf{Z} \in \mathbb{R}^{n \times m}$ form a basis for the null space of $\mathbf{A}$ (so that $\mathbf{A}\mathbf{Z} = \mathbf{0}$). Note that $\mathbf{Y}$ is often not orthogonal to the null space of the constraints (and hence not a true range space basis for $\mathbf{A}^T$). However, $\mathbf{A}\mathbf{Y}$ should have full rank to ensure solvability for the (not strictly) range-space component $\mathbf{p}_y$.[3] The decomposition (2.6) permits the $\mathbf{p}_z$ and $\mathbf{p}_y$ components to be computed separately, as shown below. Let us define the reduced gradient of the objective function $\mathbf{g}_z$, the reduced Hessian $\mathbf{W}_z$, and an auxiliary matrix $\mathbf{W}_y$:

$$(2.7) \qquad\qquad \begin{aligned} \mathbf{g}_z &:= \mathbf{Z}^T\mathbf{g}, \\ \mathbf{W}_z &:= \mathbf{Z}^T\mathbf{W}\mathbf{Z}, \\ \mathbf{W}_y &:= \mathbf{Z}^T\mathbf{W}\mathbf{Y}. \end{aligned}$$

Then the reduced space SQP (RSQP) method (without line search) is given by Algorithm 1. The "+" subscript signifies an updated variable.

ALGORITHM 1. REDUCED SPACE SEQUENTIAL QUADRATIC PROGRAMMING (RSQP).

1: Choose $\mathbf{x}$
2: **loop**
3:      Evaluate $\mathbf{c}$, $\mathbf{g}$, $\mathbf{A}$, $\mathbf{W}$
4:      $\mathbf{g}_z = \mathbf{Z}^T\mathbf{g}$
5:      **if** $\|\mathbf{g}_z\| \leq \text{tol}$ and $\|\mathbf{c}\| \leq \text{tol}$ **then**
6:          Converged
7:      **end if**
8:      $(\mathbf{A}\mathbf{Y})\mathbf{p}_y = -\mathbf{c}$ solve for $\mathbf{p}_y$
9:      $\mathbf{W}_z\mathbf{p}_z = -(\mathbf{g}_z + \mathbf{W}_y\mathbf{p}_y)$ solve for $\mathbf{p}_z$
10:      $\mathbf{p}_x = \mathbf{Z}\mathbf{p}_z + \mathbf{Y}\mathbf{p}_y$
11:      $\mathbf{x}_+ = \mathbf{x} + \mathbf{p}_x$
12:      $(\mathbf{A}\mathbf{Y})^T\boldsymbol{\lambda}_+ = -\mathbf{Y}^T(\mathbf{g} + \mathbf{W}\mathbf{p}_x)$ solve for $\boldsymbol{\lambda}_+$
13: **end loop**

Choices on how to approximate $\mathbf{W}_z$ and on how to construct the bases $\mathbf{Z}$ and $\mathbf{Y}$ determine some of the different RSQP variants. An important step of this algorithm is "inverting" $\mathbf{W}_z$. The condition number of the reduced Hessian is affected by the choice of $\mathbf{Z}$, and ideally $\mathbf{Z}$ would come from an orthogonal factorization of $\mathbf{A}$. This approach is not possible for discretized PDE constraints, particularly in three dimensions, due to the high computational cost of dense linear algebra operations. A convenient form of $\mathbf{Z}$ that is easy to compute and takes advantage of the structure of the PDE constraints is given by

$$(2.8) \qquad\qquad \mathbf{Z} := \begin{bmatrix} -\mathbf{A}_s^{-1}\mathbf{A}_d \\ \mathbf{I} \end{bmatrix}.$$

---

[3]One obvious choice for $\mathbf{Y}$ is $\mathbf{A}^T$ but the computation of $\mathbf{p}_y$ becomes prohibitively expensive since this approach squares the condition number of the PDE Jacobian.

This form is motivated by the fact that $\mathbf{A}_s$ is the (linearized) forward PDE operator and expressing the null basis in terms of the (formal) inverse of $\mathbf{A}_s$ builds on a large body of PDE solver technology. $\mathbf{Y}$ can then be defined by

$$(2.9) \qquad \mathbf{Y} := \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \end{bmatrix}.$$

The expression for the null space basis $\mathbf{Z}$ implies that the reduced gradient is given by

$$(2.10) \qquad \mathbf{g}_z = \mathbf{g}_d - \mathbf{A}_d^T \mathbf{A}_s^{-T} \mathbf{g}_s,$$

and the reduced Hessian is given by

$$(2.11) \qquad \mathbf{W}_z = \mathbf{A}_d^T \mathbf{A}_s^{-T} \mathbf{W}_{ss} \mathbf{A}_s^{-1} \mathbf{A}_d - \mathbf{A}_d^T \mathbf{A}_s^{-T} \mathbf{W}_{sd} - \mathbf{W}_{ds} \mathbf{A}_s^{-1} \mathbf{A}_d + \mathbf{W}_{dd}.$$

Algorithm 2 states a variant of RSQP particularized to the special form of $\mathbf{Z}$ and $\mathbf{Y}$ given above, to which we refer as *Newton-RSQP* (N-RSQP). The algorithm can be decomposed into three main steps: inverting the reduced Hessian for the decision search direction (step 8), inverting the (linearized) forward operator for the state search direction (step 9), and inverting the adjoint of the (linearized) forward operator for the Lagrange multipliers (step 10).

ALGORITHM 2. NEWTON RSQP.

1: Choose $\mathbf{x}_s$, $\mathbf{x}_d$, $\boldsymbol{\lambda}$
2: **loop**
3:     Evaluate $\mathbf{c}$, $\mathbf{g}$, $\mathbf{A}$, $\mathbf{W}$
4:     $\mathbf{g}_z = \mathbf{g}_d + \mathbf{A}_d^T \boldsymbol{\lambda}$
5:     **if** $\|\mathbf{g}_z\| \le$ tol and $\|\mathbf{c}\| \le$ tol **then**
6:         Converged
7:     **end if**
8:     $\mathbf{W}_z \mathbf{p}_d = -\mathbf{g}_z + (\mathbf{W}_{ds} - \mathbf{A}_d^T \mathbf{A}_s^{-T} \mathbf{W}_{ss}) \mathbf{A}_s^{-1} \mathbf{c}$ solve for $\mathbf{p}_d$ (**Decision step**)
9:     $\mathbf{A}_s \mathbf{p}_s = -\mathbf{A}_d \mathbf{p}_d - \mathbf{c}$ solve for $\mathbf{p}_s$ (**State step**)
10:    $\mathbf{A}_s^T \boldsymbol{\lambda}_+ = -(\mathbf{g}_s + \mathbf{W}_{ss} \mathbf{p}_s + \mathbf{W}_{sd} \mathbf{p}_d)$ solve for $\boldsymbol{\lambda}_+$ (**Adjoint Step**)
11:    $\mathbf{x}_+ = \mathbf{x} + \mathbf{p}_x$
12: **end loop**

There are two ways to solve the decision equation in step 8. We can either compute and store the reduced Hessian $\mathbf{W}_z$ for factorization with a direct solver; or we can use an iterative method, which requires only matrix-vector products with $\mathbf{W}_z$. Direct solvers are very effective, but computing the reduced Hessian requires m linearized forward solves (one for each column of $\mathbf{A}_d$ in the product $\mathbf{A}_s^{-1} \mathbf{A}_d$). If we count the solves for the right-hand side of the decision equation and the adjoint and forward equations, then we have a total of $m + 4$ linearized PDE solves *at each* optimization iteration.

If an iterative method is used to solve the decision equation in step 8, then a matrix-vector multiplication with $\mathbf{W}_z$ requires one solve with each of $\mathbf{A}_s$ and $\mathbf{A}_s$. A Krylov method such as conjugate gradients (CG) will converge in m steps (in exact arithmetic) and thus the number of forward/adjoint solves will not exceed $2m$. In practice the number of iterations depends on the spectrum of the reduced Hessian. With an optimal preconditioner the iteration count can be independent of m.

However, it is not obvious how to devise optimal preconditioners for the reduced Hessian for general problems. Furthermore, we still require four additional forward (and adjoint) solves for each SQP iteration, and two solves per CG iteration.

Even when the number of decision variables is small, it is advantageous to choose an RSQP variant that avoids computing $\mathbf{W}_z$. The main reason is that second derivatives are often difficult to compute. Moreover, Newton's method is not globally convergent and far from the solution the quality of the decision step $\mathbf{p}_z$ is questionable.

Motivated by these difficulties, quasi-Newton RSQP (QN-RSQP) methods have been developed that replace $\mathbf{W}_z$ by a quasi-Newton approximation, which drops the convergence rate from quadratic to superlinear. The QN-RSQP method is summarized in Algorithm 3. In addition, the second derivative terms are dropped from the right-hand sides of the decision and adjoint steps, at the expense of a reduction from one-step to two-step superlinear convergence [8]. An important advantage of this variant of quasi-Newton method is that only two linearized forward/adjoint problems need to be solved at each iteration, as opposed to the m forward solves required by N-RSQP for constructing $\mathbf{A}_s^{-1}\mathbf{A}_d$ in $\mathbf{W}_z$ [19]. Furthermore, no second derivatives are required. Under assumptions on the existence and uniqueness of a local minimum, one can show that the combination of a sufficiently accurate line search and an appropriate quasi-Newton update guarantees a descent search direction and thus a globally convergent algorithm [12].

ALGORITHM 3. QUASI-NEWTON RSQP.

1: Choose $\mathbf{x}_s$, $\mathbf{x}_d$, $\boldsymbol{\lambda}$, $\mathbf{B}_z$
2: **loop**
3:      Evaluate $\mathbf{c}$, $\mathbf{g}$, $\mathbf{A}$
4:      $\mathbf{g}_z = \mathbf{g}_d + \mathbf{A}_d^T \boldsymbol{\lambda}$
5:      **if** $\|\mathbf{g}_z\| \leq$ tol and $\|\mathbf{c}\| \leq$ tol **then**
6:         Converged
7:      **end if**
8:      $\mathbf{B}_z \mathbf{p}_d = -\mathbf{g}_z$ solve for $\mathbf{p}_d$ (**Decision step**)
9:      $\mathbf{A}_s \mathbf{p}_s = -\mathbf{A}_d \mathbf{p}_d - \mathbf{c}$ solve for $\mathbf{p}_s$ (**State step**)
10:    $\mathbf{A}_s^T \boldsymbol{\lambda}_+ = -\mathbf{g}_s$ solve for $\boldsymbol{\lambda}_+$ (**Adjoint Step**)
11:    $\mathbf{x}_+ = \mathbf{x} + \mathbf{p}_x$
12:    Update $\mathbf{B}_z$
13: **end loop**

QN-RSQP has been parallelized efficiently for moderate numbers of decision variables [29]; the parallelism is fine-grained across the state and adjoint fields. Unfortunately, the number of iterations taken by quasi-Newton methods often increases as the number of decision variables grows,[4] rendering large-scale problems intractable. Additional processors will not help since the bottleneck is in the iteration dimension.

On the other hand, convergence of the nonlinear iterations of N-RSQP method can be independent of the number of decision variables m, particularly when these variables are mesh related. But unless an optimal $\mathbf{W}_z$ preconditioner can be used, the large number of forward/adjoint solves per iteration preclude its use, particularly on a parallel machine, where iterative methods for the forward problem must be used.

---

[4]For example, for the limiting case of a quadratic programming problem, the popular BFGS quasi-Newton method is equivalent to conjugate gradients, which scales with the square root of the condition number of the reduced Hessian.

However, there is a way to exploit the fast convergence of the Newton method and avoid solving the PDEs exactly, and this method is proposed in the next section. The basic idea is to solve the full space system (2.4) simultaneously for states, adjoints, and decision variables but invoke QN-RSQP as a preconditioner. Since it is just a preconditioner, we can avoid exact solution of the linearized state and adjoint equations in favor of applications of state and adjoint preconditioners.

3. LNKS method. The KKT optimality conditions (2.4) define a system of nonlinear equations. The Jacobian $\mathbf{K}$ of this system is termed the *KKT matrix*. A Newton step on the optimality conditions is given by

$$(3.1) \qquad \begin{bmatrix} \mathbf{W} & \mathbf{A}^T \\ \mathbf{A} & \mathbf{0} \end{bmatrix} \left\{ \begin{array}{c} \mathbf{p}_x \\ \mathbf{p}_\lambda \end{array} \right\} = - \left\{ \begin{array}{c} \mathbf{g} + \mathbf{A}^T \boldsymbol{\lambda} \\ \mathbf{c} \end{array} \right\} \quad (\text{or } \mathbf{K}\mathbf{v} = -\mathbf{h}),$$

where $\mathbf{p}_x$ and $\mathbf{p}_\lambda$ are the updates of $\mathbf{x}$ and $\boldsymbol{\lambda}$ from current to next iterations. Assuming sufficient smoothness, and that the initial guess is sufficiently close to a solution, Newton steps obtained by the above system will converge quadratically to a solution of the KKT system [17]. Thus, the forward solves required for reduced methods can be avoided by remaining in the full space of state and decision variables, since it is the reduction onto the decision space that necessitates the solves with $\mathbf{A}_s^{-1}$. Nevertheless, the full space approach presents difficulties too: a descent direction is not guaranteed, second derivatives are required, and the KKT system itself can be difficult to solve. The size of the KKT matrix is more than twice that of the forward problem, and it is expected to be very ill-conditioned. Ill-conditioning results not only from the ill-conditioning of the forward problem, but also from the different scales between first and second derivatives submatrices. Moreover, the system is indefinite; mixing negative and positive eigenvalues is known to slow down Krylov solvers. Therefore, a good preconditioner is essential to make the method efficient, and that is the main focus of Part I of this two-part article.

The LNKS method we present here uses a proper Newton method to solve the KKT optimality conditions. To compute the Newton step we solve the KKT system using an appropriate Krylov method. At the core of the algorithm lies the preconditioner $\mathbf{P}$ for the Krylov method, which in our case is an *inexact* version of the QN-RSQP algorithm. An outline of the local LNKS algorithm is given by Algorithm 4. In part II, we discuss and analyze globalization and inexactness aspects of LNKS.

ALGORITHM 4. LAGRANGE–NEWTON–KRYLOV–SCHUR.

1: Choose $\mathbf{x}$, $\boldsymbol{\lambda}$
2: **loop**
3:     Check for convergence
4:     Compute $\mathbf{c}$, $\mathbf{g}$, $\mathbf{A}$, $\mathbf{W}$
5:     Solve $\mathbf{P}^{-1}\mathbf{K}\mathbf{v} = -\mathbf{P}^{-1}\mathbf{h}$
6:     Update $\mathbf{x}_+ = \mathbf{x} + \mathbf{p}_x$
7:     Update $\boldsymbol{\lambda}_+ = \boldsymbol{\lambda} + \mathbf{p}_\lambda$
8: **end loop**

This algorithm will produce the same steps taken by N-RSQP (in exact arithmetic, with exact solves and without globalization). It is easier to see how RSQP can be used to precondition LNKS if we rewrite the Newton step (3.1) in a block-partitioned

form:

$$(3.2) \qquad \begin{bmatrix} \mathbf{W}_{ss} & \mathbf{W}_{sd} & \mathbf{A}_s^T \\ \mathbf{W}_{ds} & \mathbf{W}_{dd} & \mathbf{A}_d^T \\ \mathbf{A}_s & \mathbf{A}_d & \mathbf{0} \end{bmatrix} \begin{Bmatrix} \mathbf{p}_s \\ \mathbf{p}_d \\ \mathbf{p}_\lambda \end{Bmatrix} = - \begin{Bmatrix} \mathbf{g}_s + \mathbf{A}_s^T \boldsymbol{\lambda} \\ \mathbf{g}_d + \mathbf{A}_d^T \boldsymbol{\lambda} \\ \mathbf{c} \end{Bmatrix}.$$

RSQP is equivalent to a block-row elimination; given $\mathbf{p}_d$, solve the last block of equations for $\mathbf{p}_s$, then solve the first to find $\mathbf{p}_\lambda$, and finally solve the middle one for $\mathbf{p}_d$, the search direction for the decision variables. Therefore RSQP can be written as a particular block factorization of the KKT matrix:

$$(3.3) \qquad \mathbf{K} = \begin{bmatrix} \mathbf{W}_{ss}\mathbf{A}_s^{-1} & \mathbf{0} & \mathbf{I} \\ \mathbf{W}_{ds}\mathbf{A}_s^{-1} & \mathbf{I} & \mathbf{A}_d^T\mathbf{A}_s^{-T} \\ \mathbf{I} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{A}_s & \mathbf{A}_d & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_z & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_{sd} - \mathbf{W}_{ss}\mathbf{A}_s^{-1}\mathbf{A}_d & \mathbf{A}_s^T \end{bmatrix}.$$

Note that these factors are permutable to block triangular form (and therefore this is equivalent to a block-LU factorization) and that $\mathbf{W}_z$ is the Schur complement for $\mathbf{p}_d$.

This block factorization suggests a preconditioner created by replacing the reduced Hessian $\mathbf{W}_z$ with an approximation $\mathbf{B}_z$, for example a quasi-Newton approximation. However, we still require four forward solves per Krylov iteration. One way to restore the two solves per iteration of QN-RSQP is to, in addition, drop second-order information from the preconditioner, just as one does when going from N-RSQP to QN-RSQP. A further simplification of the preconditioner is to replace the exact forward operator $\mathbf{A}_s$ by an approximation $\tilde{\mathbf{A}}_s$, which could be any appropriate forward problem preconditioner. With these changes, *no* forward solves need to be performed at each inner (Krylov) iteration. Thus, the work per inner iteration becomes linear in the state and adjoint variable dimension (e.g., when $\tilde{\mathbf{A}}_s$ is a constant-fill domain decomposition approximation). Furthermore, when $\mathbf{B}_z$ is based on a limited-memory quasi-Newton update, the work per inner iteration is also linear in the decision variable dimension. Since all of the steps involved in an inner iteration not only require linear work but are also readily parallelized at the fine-grained level, we conclude that each inner (Krylov) iteration will execute with high parallel efficiency and scalability.

Scalability of the entire method additionally requires mesh-independence of both inner and outer iterations. Newton methods (unlike quasi-Newton) are often characterized by a number of iterations that is independent of problem size [1]. With an "optimal" forward preconditioner and a good $\mathbf{B}_z$ approximation, we can expect that the number of inner iterations is also insensitive to the problem size. Scalability with respect to both state and decision variables would then result. Several KKT preconditioning variations on this theme are discussed in the next subsection, and some choices for the reduced Hessian preconditioner $\mathbf{B}_z$ are presented in the following subsection.

3.1. Preconditioning variants for the KKT system. Below we present several preconditioners for the KKT matrix. They are based on the block-factorization of the KKT matrix (3.3) and are therefore indefinite matrices. To examine separately the effects of discarding the Hessian terms and approximating the forward solver, we define four different variations. The subscript denotes the number of the forward PDE solves per Krylov iteration, and a tilde mark on (top of) the preconditioner symbol means that a forward or adjoint solve is replaced by a single application of its preconditioner.

• Preconditioner $\mathbf{P}_4$ replaces $\mathbf{W}_z$ by $\mathbf{B}_z$ and retains the four exact linearized PDE solves per iteration. The preconditioner is

$$(3.4) \qquad \mathbf{P}_4 = \begin{bmatrix} \mathbf{W}_{ss}\mathbf{A}_s^{-1} & \mathbf{0} & \mathbf{I} \\ \mathbf{W}_{ds}\mathbf{A}_s^{-1} & \mathbf{I} & \mathbf{A}_d^T\mathbf{A}_s^{-T} \\ \mathbf{I} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{A}_s & \mathbf{A}_d & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_z & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_y & \mathbf{A}_s^T \end{bmatrix},$$

and the preconditioned KKT matrix is

$$(3.5) \qquad \mathbf{P}_4^{-1}\mathbf{K} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_z\mathbf{B}_z^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix}.$$

• Preconditioner $\mathbf{P}_2$ replaces $\mathbf{W}_z$ by $\mathbf{B}_z$ and discards all other Hessian terms, resulting in two exact linearized PDE solves per iteration. The preconditioner is

$$(3.6) \qquad \mathbf{P}_2 = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{I} & \mathbf{A}_d^T\mathbf{A}_s^{-T} \\ \mathbf{I} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{A}_s & \mathbf{A}_d & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_z & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_s^T \end{bmatrix},$$

and the preconditioned KKT matrix is

$$(3.7) \qquad \mathbf{P}_2^{-1}\mathbf{K} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{W}_y^T\mathbf{A}_s^{-1} & \mathbf{W}_z\mathbf{B}_z^{-1} & \mathbf{0} \\ \mathbf{W}_{ss}\mathbf{A}_s^{-1} & \mathbf{W}_y\mathbf{B}_z^{-1} & \mathbf{I} \end{bmatrix}.$$

Note that the spectrum of the preconditioned KKT matrix is unaffected by dropping the second derivative terms.

• Preconditioner $\tilde{\mathbf{P}}_4$ replaces $\mathbf{W}_z$ by $\mathbf{B}_z$ and $\mathbf{A}_s$ by $\tilde{\mathbf{A}}_s$, and retains all other Hessian terms, resulting in four preconditioner applications but no forward/adjoint solves. The preconditioner is

$$(3.8) \qquad \tilde{\mathbf{P}}_4 = \begin{bmatrix} \mathbf{W}_{ss}\tilde{\mathbf{A}}_s^{-1} & \mathbf{0} & \mathbf{I} \\ \mathbf{W}_{ds}\tilde{\mathbf{A}}_s^{-1} & \mathbf{I} & \mathbf{A}_d^T\tilde{\mathbf{A}}_s^{-T} \\ \mathbf{I} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{A}}_s & \mathbf{A}_d & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_z & \mathbf{0} \\ \mathbf{0} & \tilde{\mathbf{W}}_y & \tilde{\mathbf{A}}_s^T \end{bmatrix},$$

and the preconditioned KKT matrix is

$$(3.9) \qquad \tilde{\mathbf{P}}_4^{-1}\mathbf{K} = \begin{bmatrix} \mathbf{I}_s & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_z\mathbf{B}_z^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_s^T \end{bmatrix} + \begin{bmatrix} \mathbf{0} & \mathcal{O}(\mathbf{E}_s) & \mathbf{0} \\ \mathcal{O}(\mathbf{E}_s) & \mathcal{O}(\mathbf{E}_s) & \mathcal{O}(\mathbf{E}_s) \\ \mathcal{O}(\mathbf{E}_s) & \mathcal{O}(\mathbf{E}_s) & \mathbf{0} \end{bmatrix},$$

where $\mathbf{E}_s := \mathbf{A}_s^{-1} - \tilde{\mathbf{A}}_s^{-1}$ and $\mathbf{I}_s := \mathbf{A}_s\tilde{\mathbf{A}}_s^{-1}$. The "cross-Hessian" $\tilde{\mathbf{W}}_y$ is given by (2.7) with the exception that the exact PDE solve $\mathbf{A}_s^{-1}$ is replaced by a preconditioner application $\tilde{\mathbf{A}}_s^{-1}$ in the definition of $\mathbf{Z}$. Clearly, $\mathbf{E}_s = \mathbf{0}$ and $\mathbf{I}_s = \mathbf{I}$ for exact forward PDE solves.

• Preconditioner $\tilde{\mathbf{P}}_2$ replaces $\mathbf{W}_z$ by $\mathbf{B}_z$ and $\mathbf{A}_s$ by $\tilde{\mathbf{A}}_s$, and drops all other Hessian terms, resulting in two preconditioner applications and no forward/adjoint solves. The preconditioner is

$$(3.10) \qquad \tilde{\mathbf{P}}_2 = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{I} & \mathbf{A}_d^T\tilde{\mathbf{A}}_s^{-T} \\ \mathbf{I} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{A}}_s & \mathbf{A}_d & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_z & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \tilde{\mathbf{A}}_s^T \end{bmatrix},$$

and the preconditioned KKT matrix is

$$(3.11) \qquad \tilde{\mathbf{P}}_2^{-1}\mathbf{K} = \begin{bmatrix} \mathbf{I}_s & \mathbf{0} & \mathbf{0} \\ \tilde{\mathbf{W}}_y^T & \mathbf{W}_z\mathbf{B}_z^{-1} & \mathbf{0} \\ \mathbf{W}_{ss}\tilde{\mathbf{A}}_s^{-1} & \tilde{\mathbf{W}}_y^T & \mathbf{I}_s^T \end{bmatrix} + \begin{bmatrix} \mathbf{0} & \mathcal{O}(\mathbf{E}_s) & \mathbf{0} \\ \mathbf{0} & \mathcal{O}(\mathbf{E}_s) & \mathcal{O}(\mathbf{E}_s) \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

We review some basic properties of the Krylov methods that will allow us to understand the effectiveness of the preconditioners. The performance of a Krylov method depends highly on the preconditioned operator spectrum [22, 38]. In most Krylov methods the number of iterations required to obtain the solution is at most equal to the number of distinct eigenvalues (in exact arithmetic). When solving $\mathbf{Kv} = -\mathbf{h}$ such methods satisfy the following relation for the residual at the ith iteration:

$$\mathbf{r}_i = \text{span}\left\{\mathbf{r}_0 + \mathbf{Kr}_0 + \mathbf{K}^2\mathbf{r}_0 + \cdots + \mathbf{K}^i\mathbf{r}_0\right\} = \psi(\mathbf{K})\mathbf{r}_0, \quad \psi \in \mathcal{P}_i.$$

Here $\mathbf{v}_0$ is the initial guess, $\mathbf{r}$ is defined as $\mathbf{r} = \mathbf{h} + \mathbf{Kv}$, $\mathbf{r}_0 = \mathbf{h} + \mathbf{Kv}_0$, and $\mathcal{P}_i$ is the space of monic polynomials of degree at most i. Minimum residual methods like MINRES and GMRES determine a polynomial ψ so that[5]

$$\|\mathbf{r}_i\| = \min_{\psi \in \mathcal{P}_i} \|\psi(\mathbf{K})\mathbf{r}_0\|.$$

If $\mathbf{K}$ is diagonalizable (with spectrum $\mathcal{S}(\mathbf{K})$ and $\mathbf{X}$ the matrix or its eigenvectors), then[6]

$$(3.12) \qquad \frac{\|\mathbf{r}_i\|}{\|\mathbf{r}_0\|} \leq \text{cond}(\mathbf{X}) \min_{\lambda \in \mathcal{S}(\mathbf{K})} |\psi(\lambda)|.$$

Additionally, if the spectrum lies on the positive real axis, it can be shown that

$$(3.13) \qquad \frac{\|\mathbf{r}_i\|}{\|\mathbf{r}_0\|} \leq \text{cond}(\mathbf{X})\frac{\lambda_{\max}}{\lambda_{\min}}.$$

From (3.12) we can see why clustering the eigenvalues is important. Indeed, if the eigenvalues are clustered, then the solution polynomial in the right-hand side of (3.12) can be approximated by the zeros of a low-degree polynomial (resulting in fewer Krylov iterations). When $\mathbf{K}$ is normal (unitarily diagonalizable) then $\text{cond}(\mathbf{X}) = 1$; in this case the estimates (3.12) and (3.13) are known to be sharp [22].

The preconditioned KKT matrix $\mathbf{P}_4^{-1}\mathbf{K}$ is a block-diagonal matrix with two unit eigenvalues and the m eigenvalues of the preconditioned reduced Hessian. It is immediate that GMRES takes $\mathcal{O}(\text{cond}(\mathbf{B}_z^{-1}\mathbf{W}_z))$ or at most $m + 2$ steps to converge for preconditioner $\mathbf{P}_4$. This is similar to the complexity estimates for a Krylov-based solution of the $\mathbf{B}_z$-preconditioned reduced Hessian in N-RSQP. Preconditioner $\mathbf{P}_2$ has the same effect on the spectrum of the KKT matrix, but the preconditioned KKT system is no longer a normal operator. Yet, if $\mathbf{P}_2^{-1}\mathbf{K}$ is diagonalizable and its eigenvector matrix well conditioned, relation (3.13) is still a good indicator of the effectiveness of the preconditioner.

A more efficient approach, of course, is to replace the exact forward operators $\mathbf{A}_s$ by their forward preconditioners $\tilde{\mathbf{A}}_s$, as in $\tilde{\mathbf{P}}_4$ or $\tilde{\mathbf{P}}_2$. In this case, the preconditioned KKT matrix assumes a more complicated structure. We write $\tilde{\mathbf{P}}_4^{-1}\mathbf{K}$ as the sum

---

[5]Throughout this paper we use the vector 2-norm for norms and condition numbers.

[6]Here λ is temporary overloaded to denote an eigenvalue of $\mathbf{K}$.

of two matrices; the second matrix in this sum includes terms that approach $\mathbf{0}$ as the forward preconditioner improves. The spectrum of the first matrix in (3.9) is given by $\mathcal{S}_n = \mathcal{S}(\mathbf{B}_z^{-1}\mathbf{W}_z) \cup \mathcal{S}(\tilde{\mathbf{A}}_s^{-1}\mathbf{A}_s)$. If $\tilde{\mathbf{A}}_s^{-1}\mathbf{A}_s$ is normal, then by the Bauer–Fike theorem [21] the eigenvalues are not sensitive to small KKT matrix perturbations and $\mathcal{S}(\tilde{\mathbf{P}}_4^{-1}\mathbf{K}) \cong \mathcal{S}_n$. Hence, a good preconditioner for the forward problem would bring the spectrum of $\tilde{\mathbf{P}}_4^{-1}\mathbf{K}$ close to the spectrum of $\mathbf{P}_4^{-1}\mathbf{K}$. Thus, we would expect a similar iteration count to N-RSQP, but at a much reduced cost per iteration—since the PDE solves would be replaced by preconditioner applications. For preconditioner $\tilde{\mathbf{P}}_2$, the left matrix in preconditioner (3.11) is not normal and may have ill-conditioned eigenvalues. In our numerical experiments, however, the number of Krylov iterations was unaffected by switching from $\tilde{\mathbf{P}}_4$ to $\tilde{\mathbf{P}}_2$.

Finally, we comment on our choice of a Krylov method for solving the symmetric indefinite KKT system (3.3) that has been preconditioned with any of the four preconditioners described above. The most popular method for large symmetric indefinite systems is the MINRES method. However, MINRES requires positive definite preconditioners, and the KKT preconditioners defined above are indefinite. Indefinite preconditioners are allowed in the quasi-minimum residual (QMR) method [18]. The version commonly used is the transpose-free QMR algorithm, which is designed for general nonsymmetric problems and requires two matrix-vector multiplications per Krylov iteration. In our implementation we use another variant, described in [18], that exploits symmetry and uses only one matrix-vector multiplication per iteration.

3.2. Preconditioners for the reduced Hessian. A very important component of the LNKS preconditioners is an approximation to the inverse of the reduced Hessian. In the preceding sections we suggested that a quasi-Newton approximation be used in place of the exact reduced Hessian $\mathbf{W}_z$, in the spirit of using QN-RSQP as a preconditioner. Below, we elaborate upon this choice and describe a two-step stationary iterative preconditioner that we use either stand-alone or to initialize the quasi-Newton preconditioner.

Limited memory quasi-Newton updates. Quasi-Newton methods are popular methods for approximating the reduced Hessian matrix in constrained optimization. Thus, they are natural candidates to precondition the exact reduced Hessian. In particular, we use the Broyden–Fletcher–Goldberg–Shanno (BFGS) update [32, section 8.1]. At each outer iteration we compute

$$\mathbf{s} = \mathbf{p}_d^{(+)} - \mathbf{p}_d,$$
$$\mathbf{y} = \mathbf{g}_z^{(+)} - \mathbf{g}_z$$

and update the approximation to the inverse of the reduced Hessian by

$$(3.14) \qquad \mathbf{B}_+^{-1} = (I - \omega\mathbf{y}\mathbf{s}^T)^T\mathbf{B}^{-1}(I - \omega\mathbf{y}\mathbf{s}^T) + \omega\mathbf{s}\mathbf{s}^T, \qquad \omega = \frac{1}{\mathbf{y}^T\mathbf{s}}.$$

If we encounter an iterate that has $\mathbf{y}^T\mathbf{s} \le \kappa\|\mathbf{y}\|$ we skip the update. $\mathbf{B}^{-1}$ is dense; thus for large decision-space problems, storing or operating with this matrix can be expensive or even impossible. Furthermore, applying the dense $\mathbf{B}^{-1}$ to a vector in parallel involves significant all-to-all communication and may be too expensive for problems that require large numbers of processors. Thus, we have chosen a limited memory variant of BFGS (L-BFGS), which uses a limited number of vectors to approximate the inverse of the reduced Hessian. When updating we do not compute (3.14) but instead we simply store $\mathbf{s}$ and $\mathbf{y}$. Then the action of $\mathbf{B}^{-1}$ on a vector $\mathbf{x}$ is given in

ALGORITHM 5. LIMITED MEMORY BFGS.
_____
 1: $\mathbf{q} = \mathbf{x}$, and l = number of stored vectors
 2: **for** $i = k-1, k-2, \ldots, k-l$ **do**
 3:     $\alpha_i = \omega_i \mathbf{s}_i^T \mathbf{q}$
 4:     $\mathbf{q} = \mathbf{q} - \alpha_i \mathbf{y}_i$
 5: **end for**
 6: $\mathbf{z} = \mathbf{B}_{z,0}^{-1} \mathbf{q}$
 7: **for** $i = k-l, k-l+1, \ldots, k-1$ **do**
 8:     $\beta = \omega_i \mathbf{y}_i^T \mathbf{z}$
 9:     $\mathbf{z} = \mathbf{z} + \mathbf{s}_i(\alpha_i - \beta)$
10: **end for**
_____

Algorithm 5. Analysis and details of the convergence properties of this algorithm can be found in [32, sections 9.1–9.2]. We generate the initial inverse reduced Hessian approximation, $\mathbf{B}_{z,0}$, using the stationary iterative preconditioner presented below.

Two-step stationary method. As an initialization for the L-BFGS approximation above, which can also be used as a stand-alone preconditioner, we turn to stationary iterative methods. Such methods provide us with a constant preconditioner as required by most Krylov methods, while effecting an approximate solution of the reduced Hessian system. To guarantee convergence, the most popular methods (such as Jacobi and SOR) require the iteration matrix to have spectral radius less than one, which is difficult to guarantee in the general case. Furthermore, most stationary iterative methods require some kind of splitting, which is not convenient for matrix-free treatment of operators like the reduced Hessian. A method that does not have the above restrictions is a two-step stationary iterative method, first presented by Frankel in 1950 [3, section 5.2]. The method is suitable for positive-definite matrices but requires an accurate estimate of the minimum and maximum eigenvalues. The preconditioner based on this method can be arbitrarily effective, depending on the number of iterations L it is carried to. If $\lambda_1 \le \lambda_m$ are estimates for the extreme eigenvalues of $\tilde{\mathbf{W}}_z$, then the application of the preconditioner to a vector $\mathbf{d}_{\text{in}}$ is given by Algorithm 6. Step 4 requires the action of the reduced Hessian on a vector. To avoid exact forward solves, we use the approximate reduced Hessian $\tilde{\mathbf{W}}_z$ instead of $\mathbf{W}_z$. We obtain $\tilde{\mathbf{W}}_z$ from (2.11) if we replace the linearized forward PDE operators $\mathbf{A}_s^{-1}$ and $\mathbf{A}_s^{-T}$ with their preconditioners $\tilde{\mathbf{A}}_s^{-1}$ and $\tilde{\mathbf{A}}_s^{-T}$. Like CG, the method's convergence rate depends on the square root of the condition number of the (approximate) reduced Hessian. It parallelizes well because it does not require any inner products, and because in our context the matrix-vector products require just two forward preconditioner applications. To obtain estimates of the extremal eigenvalues, we use a few iterations of the Lanczos method (once per Newton iteration).

Other possibilities for reduced Hessian preconditioning. In addition to the quasi-Newton and two-step stationary iterative preconditioners we use, there are several other possibilities, which we comment on below.

• *Incomplete factorizations*. Incomplete factorizations are popular and robust preconditioners, but an assembled matrix (modulo some exceptions for structured grids) is required. Not only is the reduced Hessian usually prohibitive to assemble (for large-scale PDE optimization problems), but it is in general dense and thus impossible to store for a large number of decision variables. An incomplete factorization would be feasible only if the exact solves are replaced by the forward preconditioner and if some kind of sparsity is enforced (perhaps via element-by-element computations and

---

ALGORITHM 6. TWO-STEP STATIONARY ITERATIVE REDUCED SPACE
                    PRECONDITIONER.

1: $\rho = \frac{1-\lambda_1/\lambda_n}{1+\lambda_1/\lambda_n}$, $\alpha = \frac{2}{1+(1-\rho^2)^{1/2}}$, $\beta_0 = \frac{1}{\lambda_1+\lambda_n}$, $\beta = \frac{2\alpha}{\lambda_1+\lambda_n}$,

2: $\mathbf{r} = -\mathbf{d}_{\mathrm{in}}$, $\mathbf{d}_0 = \mathbf{0}$, $\mathbf{d}_1 = \beta_0\mathbf{r}$

3: **for** $i = 1\cdots L$ **do**

4:      $\mathbf{r} = \tilde{\mathbf{W}}_z\mathbf{d}_1 - \mathbf{d}_{\mathrm{in}}$

5:      $\mathbf{d} = \alpha\mathbf{d}_1 + (1-\alpha)\mathbf{d}_0 - \beta\mathbf{r}$

6:      $\mathbf{d}_0 = \mathbf{d}_1$, $\mathbf{d}_1 = \mathbf{d}$

7: **end for**

---

threshold-ILU methods).

• *Sparse approximate inverse-techniques.* SPAI techniques are attractive since they require just matrix-vector multiplications and a given sparsity pattern. In [14] an analysis for CFD-related Schur complements shows that this method is promising.

• *Range space preconditioners.* In [31] two different reduced Hessian preconditioners are presented, one based on a state-Schur complement factorization and the second on power series expansions. It is assumed that a (nonsingular approximation) $\tilde{\mathbf{W}}^{-1}$ is available. The basic component of these preconditioners is the application of $\mathbf{Z}^T\tilde{\mathbf{W}}^{-1}\mathbf{Z}$ on a vector. We did not test this method because in our problems $\mathbf{W}$ has thousands of zero eigenvalues and it is not clear how to construct an approximation to $\mathbf{W}^{-1}$. If $\mathbf{W}_{dd}$ is nonsingular, then a block-Jacobi-ILU technique could be used to approximate $\mathbf{W}_z^{-1}$ with $\mathbf{W}_{dd}^{-1}$ (or an approximation $\tilde{\mathbf{W}}_{dd}^{-1}$).

• *Krylov self-preconditioning.* Another option is to take a few CG iterations in the reduced space at each preconditioner application. Since we want to avoid solving the forward problem exactly, we replace $\mathbf{A}_s^{-1}$ with $\tilde{\mathbf{A}}_s^{-1}$ in (2.11). We experimented with this approach but found that full convergence of CG was needed to avoid loss of orthogonality for the Krylov vectors. This slowed down the method significantly. We have experimented also with flexible GMRES for the KKT solver, but the nice properties of a symmetric Krylov solver are lost and the algorithm was slow to converge.

• *Infinite dimension-based preconditioners.* For PDE-constrained optimization problems one may be able to derive an approximate representation of the infinite-dimensional reduced Hessian. A discretization of this expression can then be used to precondition the reduced Hessian, e.g., [2].

Notwithstanding the above alternatives, we feel the choice of a limited memory quasi-Newton approximation, initialized by Frankel's two-step stationary iterative method, strikes the best balance between effectiveness, generality, and parallel scalability. In the next section we present numerical experiments that demonstrate the effectiveness of the stationary preconditioner. To isolate the effects of the (inner) Krylov–Schur solver on LNKS, we study problems with linear constraints and quadratic objective functions. For this class of problems, the KKT system is linear, Newton's method trivially requires just one step to converge, and globalization is not an issue. In part II of this article, we discuss how the limited memory quasi-Newton scheme can be used to safeguard the Newton step.

4. Numerical results. The LNKS method has been tested on two Stokes optimal flow control problems that give rise to quadratic programming problems. Both problems include the three-dimensional (3D) Stokes equations as PDE constraints, and Dirichlet boundary conditions on some portion of the boundary as decision variables. They differ in the choice of objective function: the first is a matching velocity
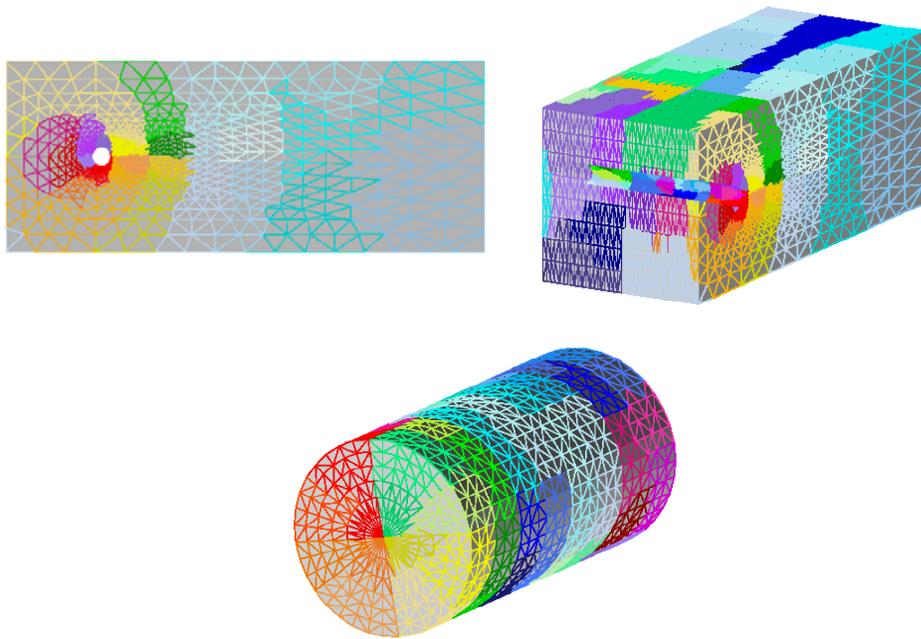
FIG. 1. *This figure depicts the (partitioned) computational meshes used in our numerical experiments. We have solved two different optimal flow control problems: that of Poiseuille flow in a pipe (bottom) and that of a flow around a cylinder embedded inside a rectangular duct (top). Both flows are interior flow problems with nonslip conditions everywhere except the outflow boundary. In the case of the Poiseuille flow the decision variables are the boundary conditions on the circumferential walls of the pipe. In the cylinder problem the decision variables are the downstream cylinder surface boundary conditions.*

problem for a Poiseuille flow, and the second is an energy dissipation minimization problem for a flow around a cylinder. Below, we define the two test problems, discuss a forward preconditioner (which as we have seen plays a key role in the effectiveness of the KKT preconditioner), and provide numerical results for both problems.

4.1. Optimal control of Stokes flows. An optimal flow control problem can be stated as follows:

$$(4.1) \qquad \min_{\boldsymbol{u}, \boldsymbol{u}_{d}, p} \mathcal{J}(\boldsymbol{u}, \boldsymbol{u}_{d}, p)$$

$$
\begin{aligned}
\text{subject to} \quad & -\nu \nabla \cdot (\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^{T}) + \nabla p = \boldsymbol{b} \text{ in } \Omega, \\
& \nabla \cdot \boldsymbol{u} = 0 \quad \text{in } \Omega, \\
(4.2) \qquad & \boldsymbol{u} = \boldsymbol{u}^{*} \quad \text{on } \Gamma_{u}, \\
& \boldsymbol{u} = \boldsymbol{u}_{d} \quad \text{on } \Gamma_{d}, \\
& -p\,\boldsymbol{n} + \nu(\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^{T})\boldsymbol{n} = \boldsymbol{0} \quad \text{on } \Gamma_{N}.
\end{aligned}
$$

Here, $\boldsymbol{u}$ is the fluid velocity, p the pressure, ν a nondimensional viscosity, $\boldsymbol{b}$ a body force, and $\boldsymbol{n}$ the outward unit normal vector on the boundary. The first problem is a velocity matching optimal control problem. A quadratic velocity profile is prescribed on the inflow boundary Γ_{u}, and we specify a traction-free outflow boundary Γ_{N}. The velocities $\boldsymbol{u}_{d}$, defined on Γ_{d}, are the decision variables. In this example, $\boldsymbol{u}^{*}$ is taken as a Poiseuille flow inside a pipe, and the decision variables correspond to boundary velocities on the circumferential surface of the pipe (Figure 1). The objective functional

is given by

$$\mathcal{J}(\boldsymbol{u}, \boldsymbol{u}_d, p) = \frac{1}{2} \int_\Omega (\boldsymbol{u}^* - \boldsymbol{u})^2 d\Omega + \frac{\rho}{2} \int_{\Gamma_d} \boldsymbol{u}_d^2 d\Omega,$$

where ρ penalizes the "cost" of the controls. Notice that the size of ρ effects the condition number of the reduced Hessian. Here we keep ρ constant and equal to 1.

In the second problem we control the flow around a cylinder that is anchored inside a rectangular duct. A quadratic velocity profile is used as an inflow Dirichlet condition and we prescribe a traction-free outflow. The decision variables are defined to be the velocities on the downstream portion of the cylinder surface. In this problem the objective functional is given by

$$\mathcal{J}(\boldsymbol{u}, \boldsymbol{u}_d, p) = \frac{1}{2} \int_\Omega |\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^T|^2 d\Omega + \frac{\rho}{2} \int_{\Gamma_d} \boldsymbol{u}_d^2 d\Omega.$$

In this problem we also choose $\rho = 1$.

The Stokes equations are discretized by the Galerkin finite element method, using tetrahedral Taylor–Hood elements [23]. We implemented the reduced-space algorithm as well as the LNKS method with the four different preconditioners. Our code is based on the PETSc library [4, 5, 6] and makes use of PETSc parallel domain-decomposition preconditioners for the approximate forward solves. We use the symmetric QMR for both the KKT and Stokes linear systems. The two flow control problems have quadratic objective functions and linear constraints and thus Newton's method takes only one iteration to converge. Hence, it is not possible to build a quasi-Newton approximation (as described above) to use within the KKT preconditioner (but see [30] for other options). Instead, we use the two-step stationary algorithm to precondition the reduced Hessian.

4.2. The Stokes forward preconditioner. It is evident that one of the two major components of the LNKS method is the forward solver preconditioner (the choice of which is, of course, problem-dependent). The Stokes problem, in algebraic form, is given by

$$(4.3) \qquad \begin{bmatrix} \mathbf{V} & \mathbf{P}^T \\ \mathbf{P} & \mathbf{0} \end{bmatrix} \left\{ \begin{array}{c} \mathbf{u} \\ \mathbf{p} \end{array} \right\} = \left\{ \begin{array}{c} \mathbf{b}_u \\ \mathbf{b}_p \end{array} \right\}.$$

In our earlier numerical experiments (results are reported in [9, 10]) we used the following preconditioner:

$$(4.4) \qquad \begin{bmatrix} \mathbf{V}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^{-1} \end{bmatrix}.$$

PETSc's block-Jacobi preconditioners with local ILU(0) were used to provide a domain decomposition approximation of $\mathbf{V}^{-1}$ and $\mathbf{M}^{-1}$, which are the discrete Laplacian and discrete pressure "mass" matrices, respectively. We found the performance of this preconditioner somewhat unsatisfactory, especially when it was part of the KKT preconditioner. For this reason we have switched to a preconditioner that is based on the following exact factorization of the inverse of (4.3),

$$(4.5) \qquad \begin{bmatrix} \mathbf{I} & -\mathbf{V}^{-1}\mathbf{P}^T \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{V}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{P}\mathbf{V}^{-1} & \mathbf{I} \end{bmatrix},$$

where $\mathbf{S} := -\mathbf{P}\mathbf{V}^{-1}\mathbf{P}^T$ is the Schur complement for the pressure. Based on this factorization, the Stokes preconditioner is defined by replacing the exact solves $\mathbf{V}^{-1}$ with domain decomposition approximations $\tilde{\mathbf{V}}^{-1}$. For the pressure Schur complement block, we use the two-step stationary method described in Algorithm 6. Performance statistics are presented in Table 4.1. We can see that the factorization-based preconditioner reduces significantly solution times compared to the earlier block-Jacobi variant. Nevertheless, the efficiency of the new preconditioner is still mesh-dependent. It is known that block-Jacobi-ILU preconditioners require work that is superlinear in problem size. This can be overcome with an additive Schwarz domain-decomposition preconditioner, with generous overlap and a coarse grid component [13].

TABLE 4.1

Effectiveness of various preconditioners for Stokes forward solver with increasing number of processors. The first two columns list the number of processors (PEs) *and the problem size* (n). *The remaining columns give the number of Krylov iterations required to satisfy* $\|\mathbf{r}\|/\|\mathbf{r}_0\| \leq 1 \times 10^{-8}$; none *signifies no preconditioner;* bj *is a block-Jacobi preconditioner, with each diagonal block mapped to a processor and approximated by ILU*(0); psc3 *is the factorization-based preconditioner, with three (two-step) stationary iterations for the Schur complement block* $\mathbf{S}$; psc7 *employs seven stationary iterations. In parentheses is wall-clock time in seconds on the Pittsburgh Supercomputing Center's Cray T3E-900.*

		Poiseuille flow			
PEs	n	None	bj	psc3	psc7
16	64,491	18,324 (250)	1,668 (86)	140 (30)	128 (39)
32	114,487	24,334 (620)	2,358 (221)	197 (39)	164 (45)
64	280,161	27,763 (1410)	3,847 (515)	294 (71)	249 (84)
128	557,693	32,764 (1297)	5,010 (620)	446 (109)	357 (123)
256	960,512	49,178 (2272)	6,531 (780)	548 (122)	389 (128)
		Flow around a cylinder			
PEs	n	None	bj	psc3	psc7
16	50,020	24,190 (720)	2,820 (206)	251 (72)	234 (103)
32	117,048	35,689 (1284)	4,512 (405)	363 (120)	327 (176)
64	389,440	41,293 (1720)	7,219 (1143)	581 (332)	497 (456)
128	615,981	52,764 (2345)	10,678 (1421)	882 (421)	632 (512)
256	941,685	71,128 (3578)	13,986 (1578)	1,289 (501)	702 (547)

4.3. Numerical results for Poiseuille optimal flow control problem. Results for the Poiseuille flow control problem are presented in Table 4.2. We have solved for five different problem sizes on up to 256 processors in order to assess the performance and scalability of the LNKS algorithm. Our first and most important finding is that the fastest variant is LNKS-II (which uses $\tilde{\mathbf{P}}_2$ as a preconditioner; i.e., it performs two approximate forward solves per KKT iteration), which is approximately 30 *times* faster than QN-RSQP. Another observation is that LNKS-I (which uses $\mathbf{P}_2$, i.e., two exact forward solves), despite its discarding second-order terms and approximating the reduced Hessian, is very effective in reducing the number of iterations—note the difference in KKT iterations between LNK (unpreconditioned) and LNKS-I for each problem instance. Indeed, the mesh-independence of LNKS-I iterations suggests that the two-step stationary preconditioner is very effective for the reduced Hessian,[7] and that if a forward preconditioner that is spectrally equivalent to the forward operator is available, one can expect mesh-independence for the entire KKT preconditioner. Even if a suboptimal forward preconditioner is used (witnessed

---

[7]Also, the weak increase in quasi-Newton iterations suggests that the reduced Hessian for this problem is well-conditioned.

TABLE 4.2

Performance of LNKS and comparisons with QN-RSQP for the Poiseuille flow problem as a function of increasing number of state and decision variables and number of processors. Here, QN-RSQP *is quasi-Newton reduced-space SQP;* LNK *is the full-space Lagrange–Newton–Krylov method with no preconditioning for the KKT system;* LNKS-I *uses* $\mathbf{P}_2$ *KKT preconditioner—which requires two Stokes solves—combined with the two-step stationary preconditioner for the reduced Hessian* (3) *iterations);* LNKS-II *uses the* $\tilde{\mathbf{P}}_2$ *preconditioner with the same reduced Hessian preconditioner;* LNKS-III *is the same as* LNKS-II *except that no reduced Hessian preconditioning is used, i.e.,* $\mathbf{B}_z = \mathbf{I}$*;* QN iter *reports the number of quasi-Newton iterations (the number of outer iterations for the LNKS variants is just one, since this is a quadratic optimization problem);* KKT iter *is the number of Krylov (symmetric QMR) iterations for the KKT system to converge; finally* time *is wall-clock time in seconds on a Cray T3E-900. Both QN-RSQP and LNKS methods are converged so that* $\|\mathbf{c}\|/\|\mathbf{c}_0\| \leq 1 \times 10^{-6}$ *and* $\|\mathbf{g}\|/\|\mathbf{g}_0\| \leq 1 \times 10^{-6}$ *in all cases but the unpreconditioned KKT case, in which the Krylov method was terminated when the number of iterations exceeded* 500,000*.*

States Decisions	Method	QN iter	KKT iter	Time
64,491	QN-RSQP	221	—	15,365
7,020	LNK	—	274,101	19,228
(16 PEs)	LNKS-I	—	27	1,765
	LNKS-II	—	170	482
	LNKS-III	—	1,102	969
114,487	QN-RSQP	224	—	19,493
10,152	LNK	—	499,971	39,077
(32 PEs)	LNKS-I	—	26	2,089
	LNKS-II	—	258	519
	LNKS-III	—	1,525	1225
280,161	QN-RSQP	228	—	33,167
18,144	LNK	—	>500,000	—
(64 PEs)	LNKS-I	—	29	4,327
	LNKS-II	—	364	934
	LNKS-III	—	1,913	1,938
557,693	QN-RSQP	232	—	53,592
28,440	LNK	—	>500,000	—
(128 PEs)	LNKS-I	—	29	6,815
	LNKS-II	—	603	1,623
	LNKS-III	—	2,501	3,334
960,512	QN-RSQP	241	—	63,865
40,252	LNK	—	>500,000	—
(256 PEs)	LNKS-I	—	32	8,857
	LNKS-II	—	763	1,987
	LNKS-III	—	2,830	3,735

by the growth in KKT iterations with increasing problem size for LNKS-II), one can trade off increasing KKT iterations for reduced work per iteration to realize a significant reduction in overall work. For example, in the largest problem size case, we give up a factor of 24 in KKT iterations, but we gain a factor of 4.5 in execution time.

The number of unpreconditioned KKT iterations illustrates the severe ill-conditioning of the KKT matrix (even for a problem as simple as controlling a Poiseuille flow). The comparison between LNKS-I and QN-RSQP simply illustrates the better convergence properties of a true Newton, as opposed to quasi-Newton, method. The comparison between LNKS-II and LNKS-III reveals the significance of a good preconditioner for the reduced space. When the two-step preconditioner is used, LNKS runs approximately twice as quickly as with no preconditioning (and requires four times fewer iterations).

The scalability of LNKS is studied by tracking the execution time for LNKS as

TABLE 4.3

Isogranular scalability results for LNKS-I. PEs is the number of processors; max Mflop/s/PE is the per-processor maximum (across PEs) sustained Mflop/s; Mflop/s/PE is the average sustained Mflop/s per processor. A comparison between the second and third columns is an indication of load imbalance. Implementation efficiency (η_i) is based on Mflop rate; optimization work efficiency (η_a) is based on number of optimization iterations; forward solver work efficiency (η_f) is based on the number of forward solver Krylov iterations psc3 (Table 4.1); overall efficiency (η) is based on execution time; (η') is an estimate of the overall efficiency given by $\eta' = \eta_f \times \eta_i \times \eta_a$.

PEs	Max Mflop/s/PE	Mflop/s/PE	η_i	η_a	η_f	η	η'
16	76.5	52.0	1.00	1.00	1.00	1.00	1.00
32	74.8	51.1	0.98	1.04	0.71	0.84	0.72
64	74.9	49.2	0.96	0.93	0.48	0.41	0.43
128	71.2	47.8	0.91	0.93	0.31	0.26	0.26
256	69.8	45.1	0.87	0.84	0.26	0.18	0.19

the problem size and number of processors increase proportionately. The problem size per processor is held (relatively) constant, and execution time increases from about 8 minutes for 16 processors (65K states, 7K decisions) to about 30 minutes for 256 processors (960K states, 40K decisions). One may conclude that the method is not scalable. However, a glance at the KKT iterations column reveals that the number of optimization iterations when using the LNKS-I variant (so that the effect of inexact forward problem solves is factored out) is largely independent of problem size, and thus the work efficiency of LNKS should be very high. Furthermore, the Mflop rate drops (probably due to a poorer-quality mesh partition) only slowly as the problem size increases, suggesting good implementation efficiency. What, then, accounts for the increase in execution time?

Table 4.3 provides the answer. Following [25], the overall parallel efficiency η (based on execution time) has been decomposed into an implementation efficiency η_i (based on Mflop/s and which represents costs like latency, load imbalance, and interprocessor communication), and a work efficiency η_a (based on the number of optimization iterations), and the forward solver work efficiency η_f (based on Table 4.1). Whereas the optimization algorithm and the implementation are reasonably scalable (84% and 87% efficiency over a 16-fold increase in number of processors), the forward solver's parallel efficiency drops to near 26% for the largest problem size. The last column, η', gives an estimate of what the overall efficiency would be if we did not use the time measurement but instead set $\eta' = \eta_i * \eta_a * \eta_f$. The agreement between the last two columns makes apparent that the overall efficiency of the algorithm depends greatly upon the forward solver, which is essentially a question of forward preconditioning. The parallel inefficiency of the forward preconditioner can be addressed by switching to a more scalable approximation than the one we are currently using (nonoverlapping local ILU(0) block-Jacobi). With a better forward preconditioner, we anticipate good overall scalability.

Timings and flop measurements were performed by using PETSc logging routines, which were validated with native performance analyzers on the T3E and Origin platforms. At first glance CPU performance appears to be mediocre—less than 10% of the peak machine performance. Recall, however, that unstructured grid computations are not cache coherent and therefore the bottleneck is in memory bandwidth and not in the CPU flop rate.

Remark 1. We do not advocate the use of exact forward solves within the LNKS context. If one can afford exact solves, then one should iterate in the reduced space.

If both exact forward solves and Hessian terms are retained, then 4 solves per KKT–Krylov iteration are required; iterating with N-RSQP requires only 2 solves per $\mathbf{W}_z$-Krylov iteration. Even when the Hessian terms are dropped (which is equivalent to retaining only $\mathbf{g}_z$ on the right-hand side of the decision equation (in Algorithm 2)), the positive definiteness of the reduced Hessian permits the use of CG in the reduced space, which seems more natural (in conjunction with a good reduced Hessian preconditioner) than iterating in the full space.

On the other hand, an order-of-magnitude argument can be used (assuming $\mathbf{A}_s$ is symmetric) to illustrate why it should be advantageous to stay in the full space when approximate solves are used. If the condition number of the preconditioned reduced Hessian is given by μ_m/μ_1 and the condition number of the preconditioned forward operator is given by λ_n/λ_1, the complexity for inverting the exact reduced Hessian is $\mathcal{O}(\sqrt{\mu_m/\mu_1} \times \sqrt{\lambda_n/\lambda_1} \times N)$, since the number of expected iterations is $\mathcal{O}(\sqrt{\mu_m/\mu_1})$ and for each iteration we need to solve the forward problem. (Recall that N is the number of states plus the number of decision variables.) Using the fact that (3.9) is block-diagonal we see that the complexity for inverting the preconditioned KKT system is given by $\mathcal{O}(\sqrt{\max(\mu_m, \lambda_n)/\min(\mu_1, \lambda_1)} \times N)$. If, in addition, the spectra of the preconditioned forward and reduced Hessian operators overlap, then a reduced method has an effective condition number that is approximately the square of the KKT system's condition number. (It is exactly equal to the square of the condition number of the reduced Hessian, if $\lambda_n = \mu_m$ and $\lambda_1 = \mu_1$.)

Remark 2. We have tested the four different preconditioners for the LNKS algorithm. Here we report results only for preconditioners $\mathbf{P}_2$ and $\tilde{\mathbf{P}}_2$. In our numerical experiments preconditioners $\mathbf{P}_4$ and $\mathbf{P}_2$ took approximately the same number of KKT iterations to converge. Since $\mathbf{P}_2$ requires two solves fewer, it is twice as fast as $\mathbf{P}_4$ and for this reason we report results only for $\mathbf{P}_2$. The same trend is evident when comparing the preconditioners $\tilde{\mathbf{P}}_4$ and $\tilde{\mathbf{P}}_2$ although the difference is less pronounced.

Numerical results for cylinder optimal flow control problem. In order to further test the KKT solver we performed additional numerical experiments for the flow around a cylinder. Figure 2 compares the controlled and the uncontrolled flow. Optimization drives the downstream surface of the cylinder to become a sink for the flow. As we can see in Figure 2, the locally controlled flow (around the cylinder) appears to be more irregular than the uncontrolled one. Yet the overall dissipation is reduced. The results can be explained if we consider the dissipation function as the sum along both the upstream and the downstream part of the duct. Flow suction results in reduced mass flow (i.e., velocity) on the downstream portion of the duct and thus minimizes dissipation losses.

Table 4.4 presents scalability results for this problem. We observe similar performance to the Poiseuille flow. Using exact second derivatives accelerates the algorithm by a factor of two. Switching to the inexact preconditioners makes the method another 6 (or more) times faster. For the 256 processor problem (941,685 states and 11,817 controls) quasi-Newton RSQP required 38 hours, whereas the LNKS algorithm with the $\tilde{\mathbf{P}}_2$ preconditioner required only 3 hours, almost 13 times faster. Notice that the fastest LNKS variant for this example (LNKS-III) does not precondition in the reduced space. Although LNKS-II takes fewer iterations to converge, the cost per KKT Krylov iteration is increased due to load imbalance—the Poiseuille flow problem has more uniform distribution of decision variables across processors.

The CPU performance has dropped compared to the Poiseuille flow results, but not significantly, as we can see in Table 4.5. The overall efficiency appears to be gov-

(a)   (b)

(c)   (d)

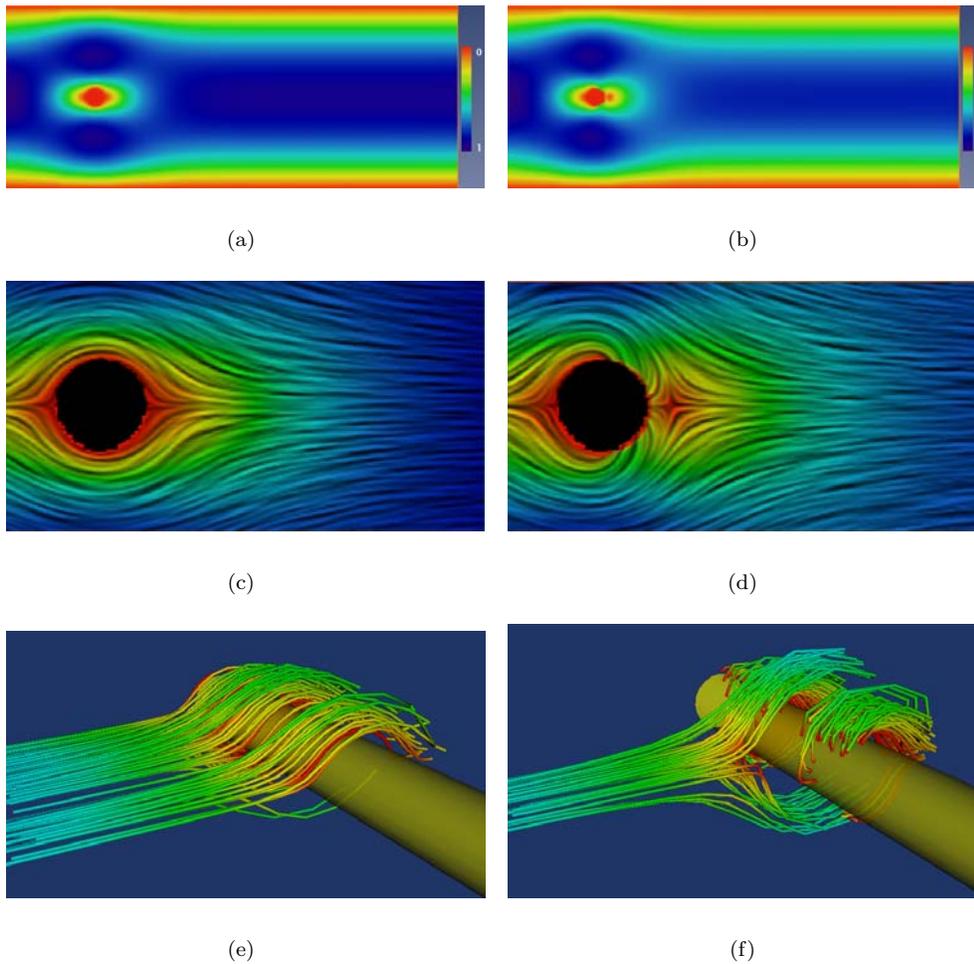(e)   (f)

Fig. 2. *Optimal boundary control problem of a steady incompressible Stokes flow around a cylinder. The objective is to minimize the energy dissipation. The controls are injection points (velocity Dirichlet boundary conditions) on the downstream portion of the cylinder surface. The left column is the uncontrolled flow and the right column is the controlled one. Images (a) and (b) depict a color map of the velocity magnitude on a cross section at the middle of the duct. Comparing image (a) with (b) notice the lighter gray throughout the downstream region. This indicates smaller velocities and thus reduced dissipation. Images (c) and (d) show a closeup snapshot of the flow around the cylinder. The suction Dirichlet conditions (set by the optimizer) are clearly visible. Finally, the last two images show stream tubes of the flow around the cylinder. Although the flow appears to be more irregular (locally) for the controlled flow, overall the dissipation is reduced.*

erned by the forward problem preconditioner; however, there is a slight disagreement between the last two columns. An explanation can be found in η_a, which drops faster than in the Poiseuille flow case. The reason is that the faster LNKS variant for the cylinder example does not precondition in the reduced space.

Our performance analyses are based on isogranular scaling. More common are fixed-problem size scalability analyses. Although these kinds of tests are very good indicators of the performance of an algorithm, they do not capture the important issue

TABLE 4.4

Performance and scalability of the LNKS algorithm for the control of a 3D flow around a cylinder. Here the objective function is the energy dissipation and the constraints are the Stokes equations. QN-RSQP is quasi-Newton reduced-space SQP; LNK *is the full-space Lagrange–Newton–Krylov method with no preconditioning;* LNKS-I *uses the* $\mathbf{P}_2$ *preconditioner combined with the two-step stationary preconditioner for the reduced Hessian;* LNKS-II *uses the* $\tilde{\mathbf{P}}_2$ *preconditioner with the same reduced Hessian preconditioning;* LNKS-III *is the same as LNKS-II but the reduced Hessian is not preconditioned;* time *is wall-clock time in seconds on a T3E-900. The reduced gradient and the constraints were converged to a relative norm of* 1×10^{-6}.

States Controls	Method	N or QN iter	KKT iter	Time
50,020	QN-RSQP	116	—	9,320
1,653	LNK	1	390,456	37,670
(16 PEs)	LNKS-I	1	41	6.833
	LNKS-II	1	1,101	3,100
	LNKS-III	1	1,312	1,812
117,048	QN-RSQP	120	—	27,669
2,925	LNK	1	>500,000	—
(32 PEs)	LNKS-I	1	39	10,780
	LNKS-II	1	1,522	5,180
	LNKS-III	1	1,731	2,364
389,440	QN-RSQP	128	—	92,874
6,549	LNK	1	> 500,000	—
(64 PEs)	LNKS-I	1	50	34,281
	LNKS-II	1	2,987	18,451
	LNKS-III	1	3,637	15,132
615,981	QN-RSQP	132	—	113,676
8,901	LNK	1	> 500,000	—
(128 PEs)	LNKS-I	1	53	54,678
	LNKS-II	1	3,150	17,144
	LNKS-III	1	4,235	9,325
941,685	QN-RSQP	138	—	140,085
11,817	LNK	1	> 500,000	—
(256 PEs)	LNKS-I	1	52	59,912
	LNKS-II	1	4,585	20,384
	LNKS-III	1	5,687	11,028

of the dependence of number of Krylov iterations on the number of unknowns (mesh size). For completeness, we present a standard fixed-problem-size scalability analysis for the (117,048 states, 2,925 controls) cylinder problem for 4 different partitions and across two different platforms: a T3E-900 and an Origin 2000. Our experiments showed Origin to have superior performance, which is surprising since the T3E has a much faster interconnect. The fact that the Origin has bigger cache size is a possible explanation. Another observation is that the effectiveness of the LNKS algorithm degrades with the number of processors. The overall efficiency drops to 71% for the T3E and to 83% for the Origin. The work efficiency remains constant but part of it is hidden in η because the time per preconditioner application increases. Recall that the condition number of a linear system preconditioned with block-Jacobi is approximately $\mathcal{O}((pn)^{1/6})$ and therefore the flop count increases with the number of processors p. Our forward solver implementation uses such a preconditioner and this is why we observe the decrease in overall efficiency. We have not conducted any measurements on latency and bandwidth dependencies on the number of processors. The inferior performance of the T3E is somewhat surprising but this would probably

TABLE 4.5

Isogranular scalability results for LNKS-III ($\bar{\mathbf{P}}_2$ KKT preconditioner with no preconditioning for the reduced Hessian). The Mflop rates given are sustained Mflop per processor; the first column is maximum across processors and the second column lists the average rate. The difference is an indication of imbalance. Efficiency (η_i) is based on Mflop rate; work efficiency (η_a) is based on number of optimization iterations; forward solver work efficiency (η_f) is based on Table 4.1; overall efficiency (η) is based on execution time; η' is an estimate of the overall efficiency given by $\eta' = \eta_f \times \eta_i \times \eta_a$.

PEs	Max Mflop/s/PE	Mflop/s/PE	η_i	η_a	η_f	η	η'
16	69.1	50.9	1.00	1.00	1.00	1.00	1.00
32	69.2	48.5	0.94	1.05	0.69	0.64	0.67
64	73.9	45.7	0.89	0.82	0.62	0.20	0.45
128	65.6	42.9	0.84	0.77	0.29	0.13	0.18
256	66.4	39.3	0.77	0.78	0.19	0.12	0.11

TABLE 4.6

Fixed-size scalability analysis for the 117,028 state variables problem. The experiments were performed on a SGI Origin 2000 and on a T3E-900. The 450 MHz Compaq/Alpha 21164 processor on the T3E is equipped with 8 KB L1 cache and 96 KB L2 cache. The Origin uses the 250 MHz MIPS R10000 processor with 32 KB L1 and 4MB L2 cache. Bisection bandwidth is 128GB/s for the T3E and 20.5 GB/s for the Origin (128 PEs).

CRAY T3E-900

Procs	Agr Gflop/s	Its	Time	Speedup	η_i	η_a	η
16	0.81	38	18,713	1.00	1.00	1.00	1.00
32	1.55	39	10,170	1.84	0.95	0.97	0.92
64	3 14	40	5,985	3.13	0.97	0.95	0.92
128	4.86	40	3,294	5.68	0.75	0.95	0.71

SGI ORIGIN 2000

Procs	Agr Gflop/s	Its	Time	Speedup	η_i	η_a	η
16	1.09	37	13,512	1.00	1.00	1.00	1.00
32	2.13	40	6,188	1.84	0.96	0.93	0.89
64	6.08	38	3,141	3.13	0.96	0.97	0.93
128	7.90	39	1,402	5.68	0.87	0.95	0.83

change had we increased the number of processors further.[8]

4.4. LNKS parallel scalability. In this section, we provide arguments for verifying the scalability of the LNKS method, with respect to increasing problem size and number of processors. For scalability, we require that the work increases near-linearly with problem size (work scalability) and that it parallelizes well. Let us examine the major components.

Formation of the KKT matrix-vector product. For PDE-constrained optimization, the Hessian of the Lagrangian function and the Jacobian of the constraints are usually sparse with structure dictated by the mesh (particularly when the decision variables are mesh-related). Thus, formation of the matrix-vector product at each QMR iteration is linear in both state and decision variables, and parallelizes well due to a high computation-to-communication ratio and minimal sequential bottlenecks characteristic of parallel matvecs for PDE-type problems.

---

[8] At the time these experiments were performed, the 256 partition on the Origin was not available.

Application of the KKT preconditioner. The main work involved is application of the PDE preconditioner $\tilde{\mathbf{A}}_s$ and its transpose, and "inversion" of an approximation to the reduced Hessian, $\mathbf{B}_z$. We can often make use of scalable, parallel PDE preconditioners that require $\mathcal{O}(n)$ work to apply (as in various domain decomposition preconditioners or multigrid for elliptic problems). The stationary preconditioner for the reduced Hessian involves only matrix-vector multiplications and Jacobian preconditioner applications, and therefore requires linear work and parallelizes well. The same can be said about the limited-memory quasi-Newton update. Therefore, we conclude that application of the KKT preconditioner requires linear work and parallelizes well.

The Krylov (inner) iteration. As argued above, with an "optimal" state preconditioner and a good $\mathbf{B}_z$ approximation, we can expect that the number of inner (Krylov) iterations will be relatively insensitive to the problem size.

The Lagrange–Newton (outer) iteration. The number of outer (Newton) iterations is often independent of problem size for PDE-type problems, and the PDE-constrained optimization problems we have solved exhibit this type of behavior as well. For an extended discussion and results from numerical experiments on the outer iteration we refer the reader to part II of this article [11].

This combination of linear work per Krylov iteration, weak dependence of Krylov iterations on problem size, and independence of Lagrange–Newton iterations on problem size suggests a method that scales well with increasing problem size and number of processors.

5. Conclusions. The main new idea offered by the LNKS algorithm for PDE-constrained optimization problems is the use of a reduced space quasi-Newton method, not as a driver, but rather as a preconditioner for the KKT system. The preconditioning exploits the PDE structure of the constraints by permitting approximate PDE solves at each KKT iteration along with preconditioning of the reduced Hessian. We argue that in order to achieve algorithmic scalability, a bona fide Newton method is necessary. LNKS requires second derivatives (but only in the form of matrix-vector multiplications, which allows for a directionally differenced approximation) and the adjoint operator of the forward problem, beyond the information required by a (forward sensitivity-based) quasi-Newton RSQP method. The most important result is that the problem of devising a good preconditioner for the KKT system is reduced to that of finding a good preconditioner for the forward PDE operator and the reduced Hessian operator.

The problems we have chosen to investigate are relatively simple, yet provide a reasonable testbed for algorithmic tuning and experimentation. The results obtained thus far are very encouraging: the full space LNKS method with reduced space preconditioning is by a factor of 10–30 faster than a popular reduced space method. We expect that other problems will exhibit similar behavior. The method can be parallelized efficiently, and work efficiency can be achieved provided a good forward preconditioner is available. Scalability then results from the combination of these two.

In part II, we extend our discussion to nonlinear PDE constraints and examine issues of robustness and globalization of the LNKS method.

REFERENCES

[1] E. L. Allgower, K. Böhmer, F. A. Potra, and W. C. Rheinboldt, *A mesh-independence principle for operator equations and their discretizations*, SIAM J. Numer. Anal., 23 (1986), pp. 160–169.

[2] E. Arian and S. Ta'asan, *Analysis of the Hessian for Aerodynamic Optimization*, Technical report 96-28, Institute for Computer Applications in Science and Engineering, Hampton, VA, 1996.

[3] O. Axelsson, *Iterative Solution Methods*, Cambridge University Press, Cambridge, UK, 1994.

[4] S. Balay, K. Buschelman, W. D. Gropp, D. Kaushik, L. C. McInnes, and B. F. Smith, *PETSc home page*, http://www.mcs.anl.gov/petsc (2001).

[5] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, *Efficient management of parallelism in object oriented numerical software libraries*, in Modern Software Tools in Scientific Computing, E. Arge, A. M. Bruaset, and H. P. Langtangen, eds., Birkhäuser, Boston, 1997, pp. 163–202.

[6] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, *PETSc Users Manual*, Technical report ANL-95/11, Revision 2.1.1, Argonne National Laboratory, Argonne, IL, 2001.

[7] A. Battermann and M. Heinkenschloss, *Preconditioners for Karush-Kuhn-Tucker matrices arising in the optimal control of distributed systems*, in Optimal Control of Partial Differential Equations. Internat. Ser. Numer. Math. 126, W. Desch, F. Kappel, and K. Kunisch, eds., Birkhäuser-Verlag, 1998, pp. 15–32.

[8] L. T. Biegler, J. Nocedal, and C. Schmid, *A reduced Hessian method for large-scale constrained optimization*, SIAM J. Optim., 5 (1995), pp. 314–347.

[9] G. Biros and O. Ghattas, *Parallel Newton-Krylov algorithms for PDE-constrained optimization*, in Proceedings of SC99, The SCxy Conference Series, Portland, OR, 1999, ACM/IEEE.

[10] G. Biros and O. Ghattas, *Parallel preconditioners for KKT systems arising in optimal control of viscous incompressible flows*, in Parallel Computational Fluid Dynamics 1999, D. E. Keyes, A. Ecer, J. Periaux, and N. Satofuka, eds., North–Holland, Amsterdam, 1999.

[11] G. Biros and O. Ghattas, *Parallel Langrange–Newton–Krylov–Schur methods for PDE-constrained optimization. Part II: The Lagrange–Newton solver and its application to optimal control of steady viscous flows*, SIAM J. Sci. Comput., 27 (2005), pp. 714–739.

[12] R. Byrd and J. Nocedal, *An analysis of reduced Hessian methods for constrained optimization*, Math. Program., 49 (1991), pp. 285–323.

[13] X.-C. Cai and O. B. Widlund, *Domain decomposition algorithms for indefinite elliptic problems*, SIAM J. Sci. Statist. Comput., 13 (1992), pp. 243–258.

[14] E. Chow and Y. Saad, *Approximate inverse techniques for block-partitioned matrices*, SIAM J. Sci. Comput., 18 (1997), pp. 1657–1675.

[15] J. E. Dennis and R. M. Lewis, *A Comparison of Nonlinear Programming Approaches to an Elliptic Inverse Problem and a New Domain Decomposition Approach*, Technical report TR-9433, Department of Computational and Applied Mathematics, Rice University, Houston, TX, 1994.

[16] C. Farhat, R. Tezaur, and R. Djellouli, *On the Solution of Three Dimensional Inverse Obstacle Scattering Problems by a Regularized Newton Method*, Technical report CU-CAS-01-11, Center for Aerospace Structures, University of Colorado, Boulder, CO, 2001.

[17] R. Fletcher, *Practical Methods of Optimization*, 2nd ed. John Wiley and Sons, Chichester, UK, 1987.

[18] R. W. Freund and N. M. Nachtigal, *An implementation of the QMR method based on coupled two-term recurrences*, SIAM J. Sci. Comput., 15 (1994), pp. 313–337.

[19] O. Ghattas and J.-H. Bark, *Optimal control of two- and three-dimensional incompressible Navier-Stokes flows*, J. Comput. Phys., 136 (1997), pp. 231–244.

[20] O. Ghattas and C. E. Orozco, *A parallel reduced Hessian SQP method for shape optimization*, in Multidisciplinary Design Optimization: State of the Art, N. M. Alexandrov and M. Y. Hussaini, eds., SIAM, Philadelphia, 1997, pp. 133–152.

[21] G. H. Golub and C. H. V. Loan, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, MD, 1996.

[22] A. Greenbaum, *Iterative Methods for Solving Linear Systems*, SIAM, Philadelphia, 1997.

[23] M. D. Gunzburger, *Finite Element Methods for Viscous Incompressible Flows*, Academic Press, Boston, 1989.

[24] E. Haber and U. Ascher, *Preconditioned all-at-once methods for large sparse parameter estimation problems*, Inverse Problems, 17 (2000), pp. 1847–1864.

[25] D. E. Keyes, *How scalable is domain decomposition in practice?*, in Proceedings of the 11th International Conference on Domain Decomposition Methods, C.-H. Lai, M. Cross, and O. Widlund, eds., 1998.

[26] D. E. Keyes and W. D. Gropp, *A comparison of domain decomposition techniques for elliptic partial differential equations and their parallel implementation*, SIAM J. Sci. Statist. Comput., 8 (1987), pp. S166–S202.

[27] K. Kunisch and E. W. Sachs, *Reduced SQP methods for parameter identification problems*, SIAM J. Numer. Anal., 29 (1992), pp. 1793–1820.

[28] F.-S. Kupfer and E. W. Sachs, *Numerical solution of a nonlinear parabolic control problem by a reduced SQP method*, Comput. Optim. Appl., 1 (1992), pp. 113–135.

[29] I. Malčević, *Large-Scale Unstructured Mesh Shape Optimization on Parallel Computers*, Master's thesis, Carnegie Mellon University, Pittsburgh, PA, 1997.

[30] J. L. Morales and J. Nocedal, *Automatic preconditioning by limited memory quasi-Newton updating*, SIAM J. Optim., 10 (2000), pp. 1079–1096.

[31] S. G. Nash and A. Sofer, *Preconditioning reduced matrices*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 47–68.

[32] J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer, New York, 1999.

[33] C. E. Orozco and O. Ghattas, *Massively parallel aerodynamic shape optimization*, Comput. Systems Engrg., 1–4 (1992), pp. 311–320.

[34] C. E. Orozco and O. Ghattas, *A reduced SAND method for optimal design of nonlinear structures*, Internat. J. Numer. Methods Engrg., 40 (1997), pp. 2759–2774.

[35] J. Reuther, J. Alonso, M. J. Rimlinger, and A. J. Jameson, *Aerodynamic shape optimization of supersonic aircraft configurations via an adjoint formulation on parallel computers*, Comput. & Fluids, 29 (1999), pp. 675–700.

[36] U. Ringertz, *Optimal Design of Nonlinear Shell Structures*, Technical report FFA TN 91-18, The Aeronautical Research Institute of Sweden, 1991.

[37] U. Ringertz, *An algorithm for optimization of nonlinear shell structures*, Internat. J. Numer. Methods Engrg., 38 (1995), pp. 299–314.

[38] Y. Saad, *Iterative Methods for Sparse Linear Systems*, PWS, Boston, 1996.

[39] A. Salinger, R. Pawlowski, J. Shadid, B. van Bloemen Waanders, R. Barlett, G. Itle, and L. Biegler, *RSQP optimization of large-scale reacting flow applications with MP-Salsa*, in Proceedings of First Sandia Workshop on PDE-Constrained Optimization, Santa Fe, NM, 2001.

[40] V. Schulz, *Solving discretized optimization problems by partially reduced SQP methods*, Comput. Vis. Sci., 1 (1997), pp. 83–96.

[41] V. H. Schulz and H. G. Bock, *Partially reduced SQP methods for large-scale nonlinear optimization problems*, in Proceedings of the Second World Congress of Nonlinear Analysis, Elsevier, 1997.