SAND2015-6140C

Inexact full-space methods for simulation-based inverse problems and large-scale optimization

Denis Ridzal

Center for Computing Research

Sandia National Laboratories

USNCCM13, San Diego, 07-27-2015

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.
Sponsored by the NNSA Advanced Scientific Computing Program under *Integrated Code Research*.

Motivation: A simple inverse problem

Inexact full-space methods: TR-SQP, LS-SQP, TR-RSQP

Preconditioners for TR-SQP in PDE-constrained optimization

A hierarchy of linear systems and solvers in full-space methods

Motivation: A simple inverse problem

Inexact full-space methods: TR-SQP, LS-SQP, TR-RSQP

Preconditioners for TR-SQP in PDE-constrained optimization

A hierarchy of linear systems and solvers in full-space methods

Thermal inversion

$$\underset{\{u,z\}\in\,\mathcal{U}\times\mathcal{Z}}{\text{Minimize}} \quad \frac{1}{2}\int_{\Omega}(u-\widehat{u})^2\,dx + r(z)$$

subject to

$$-\nabla\cdot(z\nabla u) = f \text{ in } \Omega, \quad + \text{ boundary conditions.}$$

- $\mathcal{U}$ is the **state space** – simulated temperature;
- $\mathcal{Z}$ is the **parameter (control, design) space** – thermal diffusivity;
- $r:\mathcal{Z}\to\mathbb{R}$ is a regularization functional; and
- f is a Gaussian heat source at (0,0), with amplitude 5 and width 0.1.

Temperature in uniform material

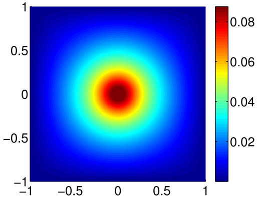

Our measured temperature $\widehat{u}$

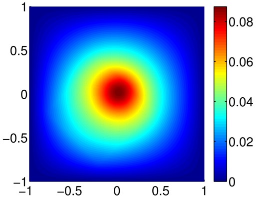

Thermal inversion

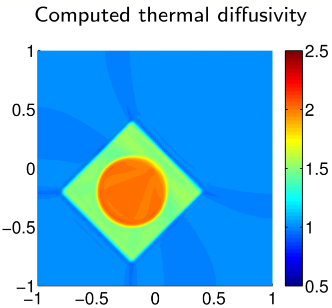

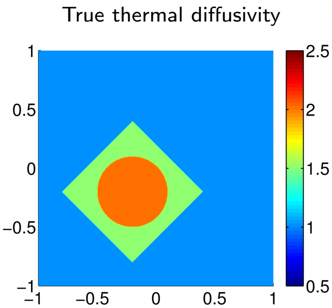

Two formulations

Full-space formulation

$$\min_{u,z} \quad \frac{1}{2}\|u - \widehat{u}\|^2 + r(z)$$

$$\text{s.t.} \quad A(z)u + Bz = f$$

- state u and control z
- the constraint is explicit in the formulation; allows us to trade feasibility for optimality
- no $A(z)^{-1}$ in the formulation

Reduced-space formulation

$$\min_{z} \quad \frac{1}{2}\|A(z)^{-1}(f - Bz) - \widehat{u}\|^2 + r(z)$$

- control z only
- the constraint is eliminated at each optimization step, by solving $A(z)u = f - Bz$
- $A(z)^{-1}$ in the objective function!

Problem classes benefiting from the full-space approach

- In full space, the solution operator $A(z)^{-1}$ is not required.

- The forward operator A(z) **is allowed to be rank deficient**.
- Examples: Acoustic inverse problems near resonance; problems without essential boundary conditions; nonlinear constraints.

- The solution operator $A(z)^{-1}$ **can be nondifferentiable**.
- Example: Multiple eigenvalues in structural optimization.

- Full-space methods can take advantage of $A(z)^{-1}$, if it is available. However, $A(z)^{-1}$ **is allowed to be very inaccurate**.
- Example: Large-scale simulations using iterative solvers.

Problem classes benefiting from the full-space approach

- In full space, the solution operator $A(z)^{-1}$ is not required.

- The forward operator A(z) **is allowed to be rank deficient**.
- Examples: Acoustic inverse problems near resonance; problems without essential boundary conditions; nonlinear constraints.

- The solution operator $A(z)^{-1}$ **can be nondifferentiable**.
- Example: Multiple eigenvalues in structural optimization.

- Full-space methods can take advantage of $A(z)^{-1}$, if it is available. However, $A(z)^{-1}$ **is allowed to be very inaccurate**.
- Example: Large-scale simulations using iterative solvers.

Takeaway: Advances in **inexact full-space SQP methods** are enabling robust and efficient solvers for the above problem classes.

Motivation: A simple inverse problem

Inexact full-space methods: TR-SQP, LS-SQP, TR-RSQP

Preconditioners for TR-SQP in PDE-constrained optimization

A hierarchy of linear systems and solvers in full-space methods

Sequential Quadratic Programming

Solve equality-constrained optimization problem, or NLP:

$$\min_{x \in \mathcal{X}} \quad f(x)$$
$$\text{s.t.} \quad c(x) = 0$$

where $f : \mathcal{X} \to \mathbb{R}$ and $c : \mathcal{X} \to \mathcal{C}$, for some Hilbert spaces $\mathcal{X}$ and $\mathcal{C}$, and f and c are twice continuously Fréchet differentiable. We identify the spaces $\mathcal{X}$ and $\mathcal{C}$ with their duals. Note that earlier $\mathcal{X} = \mathcal{U} \times \mathcal{Z}$.

Define Lagrangian functional $\mathcal{L} : \mathcal{X} \times \mathcal{C} \to \mathbb{R}$:

$$\mathcal{L}(x, \lambda) = f(x) + \langle \lambda, c(x) \rangle_{\mathcal{C}}.$$

If *regular* point x_* is a local solution of the NLP, then there exists a $\lambda_* \in \mathcal{C}$ satisfying the *first-order necessary optimality conditions*:

$$\nabla_x f(x_*) + c_x(x_*)^* \lambda_* = 0$$
$$c(x_*) = 0.$$

Sequential Quadratic Programming

Newton's method applied to optimality conditions:

$$\begin{pmatrix} \nabla_{xx}\mathcal{L}(x_k,\lambda_k) & c_x(x_k)^* \\ c_x(x_k) & 0 \end{pmatrix} \begin{pmatrix} s \\ z \end{pmatrix} = - \begin{pmatrix} \nabla_x f(x_k) + c_x(x_k)^*\lambda_k \\ c(x_k) \end{pmatrix}.$$

If $\nabla_{xx}\mathcal{L}(x_k,\lambda_k)$ is positive definite on the null space of $c_x(x_k)$, the above KKT system is necessary and sufficient for solving the QP:

$$\begin{aligned} \min_{s\in\mathcal{X}} \quad & \frac{1}{2}\langle \nabla_{xx}\mathcal{L}(x_k,\lambda_k)s, s\rangle_{\mathcal{X}} + \langle \nabla_x\mathcal{L}(x_k,\lambda_k), s\rangle_{\mathcal{X}} + \mathcal{L}(x_k,\lambda_k) \\ \text{s.t.} \quad & c_x(x_k)s + c(x_k) = 0. \end{aligned}$$

Globalization: Trust region (TR) or line search (LS).

The choice of globalization is not arbitrary. It critically determines:

- features and limitations of quadratic subproblems, e.g., convexity;
- the type of linear systems solved at every optimization iteration;
- the preconditioner/solver options and their characteristics; and
- the mechanisms to deal with the potential rank deficiency of c_x.

Composite-step approach with trust regions

- **Composite step**:
  $s_k = n_k + t_k$

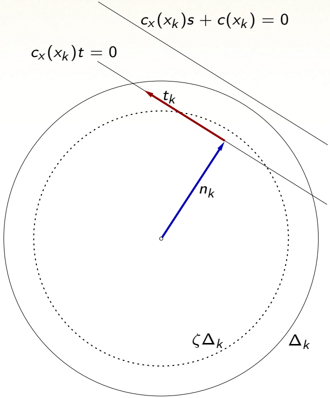

- **Quasi-normal step** n_k:
  reduces linear infeasibility

$$\min_{n \in \mathcal{X}} \quad \|c_x(x_k)n + c(x_k)\|^2_{\mathcal{C}}$$
$$\text{s.t.} \quad \|n\|_{\mathcal{X}} \leq \zeta\Delta_k$$

- **Tangential step** t_k:
  improves optimality while staying in the null space of the linearized constraints

$$\min_{t \in \mathcal{X}} \quad \frac{1}{2}\langle \nabla_{xx}\mathcal{L}(x_k, \lambda_k)(t + n_k), t + n_k\rangle_{\mathcal{X}} + \langle \nabla_x \mathcal{L}(x_k, \lambda_k), t + n_k\rangle_{\mathcal{X}} + \mathcal{L}(x_k, \lambda_k)$$
$$\text{s.t.} \quad c_x(x_k)t = 0\,, \quad \|t + n_k\|_{\mathcal{X}} \leq \Delta_k$$

Note: It is ok for the tangential step model to be nonconvex (Steihaug-Toint CG method).
Note: The quasi-normal step computation can handle rank deficiency in c_x.

Omojokun (1989), Byrd, Hribar, Nocedal (1997), Dennis, El-Alem, Maciel (1997)

Inexact TR-SQP

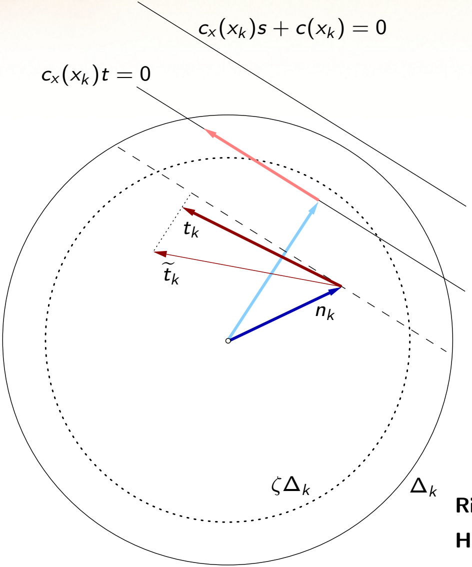

Composite step:

$s_k = n_k + t_k$

1. Compute quasi-normal step n_k using **inexact Powell dogleg**.
2. Solve tangential subproblem for $\widetilde{t}_k$ using **inexact projected ST-CG**.
3. Restore linearized feasibility, yielding tangential step t_k.
4. Update Lagrange multipliers λ_{k+1}.
5. Evaluate progress.

Ridzal, *Ph.D. Thesis, Rice University (2006)*

Heinkenschloss, Ridzal, *SIAM J. Opt. (2014)*

Linear systems

1) Given a quasi-normal Cauchy point n_k^{cp}, we solve for $\delta n_k = n_k^N - n_k^{cp}$, where n_k^N is the desired Newton step:

$$\begin{pmatrix} I & c_x(x_k)^* \\ c_x(x_k) & 0 \end{pmatrix} \begin{pmatrix} \delta n_k \\ y \end{pmatrix} = \begin{pmatrix} -n_k^{cp} + e^1 \\ -c_x(x_k) n_k^{cp} - c(x_k) + e^2 \end{pmatrix}.$$

The size of the residual $\begin{pmatrix} e^1 & e^2 \end{pmatrix} \in \mathcal{X} \times \mathcal{C}$ is restricted via

$$\|e^1\|_{\mathcal{X}}^2 + \|e^2\|_{\mathcal{C}}^2 \le \left(\xi^{qn}\right)^2 \, \|c_x(x_k) n_k^{cp} + c(x_k)\|_{\mathcal{C}}^2,$$

where $0 < \xi^{qn} \le 1$.

---

2) At every CG iteration i, we compute an inexact projection $\widetilde{z}_i = \mathcal{W}_k(\widetilde{r}_i)$:

$$\begin{pmatrix} I & c_x(x_k)^* \\ c_x(x_k) & 0 \end{pmatrix} \begin{pmatrix} \widetilde{z}_i \\ y \end{pmatrix} = \begin{pmatrix} \widetilde{r}_i \\ 0 \end{pmatrix} + \begin{pmatrix} e_i^1 \\ e_i^2 \end{pmatrix},$$

where the residual $\begin{pmatrix} e_i^1 & e_i^2 \end{pmatrix} \in \mathcal{X} \times \mathcal{C}$ is controlled via

$$\|e_i^1\|_{\mathcal{X}} + \|e_i^2\|_{\mathcal{C}} \le \xi^{proj} \min\left\{ \|\widetilde{z}_i\|_{\mathcal{X}}, \|\widetilde{r}_i\|_{\mathcal{X}} \right\},$$

with $0 < \xi^{proj} \le 1$.

3) We perform another inexact null space projection,

$$\begin{pmatrix} I & c_x(x_k)^* \\ c_x(x_k) & 0 \end{pmatrix} \begin{pmatrix} t_k \\ y \end{pmatrix} = \begin{pmatrix} \widetilde{t}_k \\ 0 \end{pmatrix} + \begin{pmatrix} e^1 \\ e^2 \end{pmatrix},$$

where the residual $\begin{pmatrix} e^1 & e^2 \end{pmatrix} \in \mathcal{X} \times \mathcal{C}$ must satisfy

$$\|e^1\|_{\mathcal{X}} + \|e^2\|_{\mathcal{C}} \le \Delta_k \min\left\{ \Delta_k, \|n_k + t_k\|_{\mathcal{X}}, \xi^{tang} \|\widetilde{t}_k\|_{\mathcal{X}} / \Delta_k \right\},$$

for $0 < \xi^{tang} \le 1$.

---

4) Let $\widehat{x}_k = x_k + n_k + t_k$. We solve for $\Delta\lambda = \lambda_{k+1} - \lambda_k$:

$$\begin{pmatrix} I & c_x(\widehat{x}_k)^* \\ c_x(\widehat{x}_k) & 0 \end{pmatrix} \begin{pmatrix} z \\ \Delta\lambda \end{pmatrix} = \begin{pmatrix} -\nabla_x f(\widehat{x}_k) - c_x(\widehat{x}_k)^* \lambda_k + e^1 \\ e^2 \end{pmatrix}.$$

The residual $\begin{pmatrix} e^1 & e^2 \end{pmatrix} \in \mathcal{X} \times \mathcal{C}$ must satisfy

$$\|e^1\|_{\mathcal{X}} + \|e^2\|_{\mathcal{C}} \le \min\left\{ \xi^{lmg}, \, \xi^{lmh} \|\nabla_x f(\widehat{x}_k) + c_x(\widehat{x}_k)^* \lambda_k\|_{\mathcal{X}} \right\},$$

for $0 < \xi^{lmh} \le 1$ and a fixed $\xi^{lmg} > 0$ independent of k.

Linear systems

... are all augmented constraint systems

$$\begin{pmatrix} I & c_x(x_k)^* \\ c_x(x_k) & 0 \end{pmatrix} \begin{pmatrix} z \\ y \end{pmatrix} = \begin{pmatrix} b^1 \\ b^2 \end{pmatrix} + \begin{pmatrix} e^1 \\ e^2 \end{pmatrix}$$

- The size of $(e_1 \;\; e_2)$ is governed by various model reduction conditions, i.e., the progress of the optimization algorithm.
- These are KKT systems for the **convex quadratic programs**

$$\begin{aligned} \min \quad & \frac{1}{2}\langle z, z\rangle_{\mathcal{X}} - \langle b^1, z\rangle_{\mathcal{X}} \\ \text{s.t.} \quad & c_x(x_k)z = b^2. \end{aligned}$$

- True even if the trust-region subproblems

$$\begin{aligned} \min \quad & \frac{1}{2}\langle \nabla_{xx}\mathcal{L}(x_k,\lambda_k)s, s\rangle_{\mathcal{X}} + \langle \nabla_x\mathcal{L}(x_k,\lambda_k), s\rangle_{\mathcal{X}} + \mathcal{L}(x_k,\lambda_k) \\ \text{s.t.} \quad & c_x(x_k)s + c(x_k) = 0, \quad \|s\|_{\mathcal{X}} \le \Delta_k \end{aligned}$$

are *not convex* !

Inexact LS-SQP

- **[1] Byrd, Curtis, Nocedal,** *SIAM J. Optim., 2008;* **[2] Byrd, Curtis, Nocedal,** *Math. Prog., 2008;* **[3] Curtis, Nocedal, Wächter,** *SIAM J. Optim., 2009.*
- Geared at inexactly solving the full KKT system:

$$\begin{pmatrix} H(x_k,\lambda_k) & c_x(x_k)^* \\ c_x(x_k) & 0 \end{pmatrix} \begin{pmatrix} s \\ z \end{pmatrix} = - \begin{pmatrix} \nabla_x f(x_k) + c_x(x_k)^* \lambda_k \\ c(x_k) \end{pmatrix} + \begin{pmatrix} e_1 \\ e_2 \end{pmatrix}.$$

- The size of the residual $\begin{pmatrix} e^1 & e^2 \end{pmatrix} \in \mathcal{X} \times \mathcal{C}$ is governed by a model-reduction condition inspired by trust-region literature.
- A backtracking line search is used to compute a steplength satisfying Armijo conditions for the merit function $\phi(x,\pi) = f(x) + \pi \|c(x)\|$.
- The operator $H(x_k,\lambda_k)$ must be positive definite on the null space of c_x; in **[2]** an iterative **inertia correction** procedure is suggested.
- To handle potential rank deficiency in c_x, in **[3]** a **composite-step strategy** is borrowed from the trust-region literature.

Inexact TR-RSQP

- **Heinkenschloss, Vicente,** *SIAM J. Optim., 2001.*
- A "reduced" SQP method, where a decomposition of the optimization variables x into basic and nonbasic variables is assumed, e.g., state variables u and control variables z.
- Very similar to inexact TR-SQP, with some simplifications. In particular, the approach only uses inexact applications of
  - the state Jacobian inverse, c_u^{-1}; and
  - its adjoint, c_u^{-*}.
- The latter is also a limitation for rank-deficient problems.
- Precursor to both TR-SQP and LS-SQP.

Reduced-space result for thermal inversion

Study inaccurate solution operator. Apply Newton-CG with trust regions. Use ML to compute $A(z)^{-1}(f - Bz)$ to tolerance tol.

tol$=10^{-1}$; ~~convergence~~

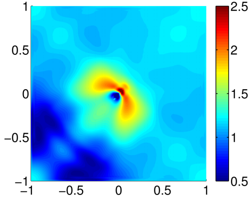

tol$=10^{-2}$; ~~convergence~~

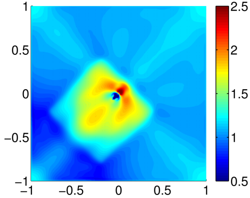

tol$=10^{-4}$; convergence

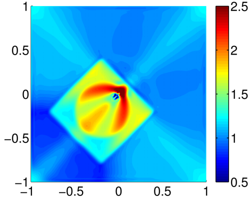

tol$=10^{-8}$; convergence

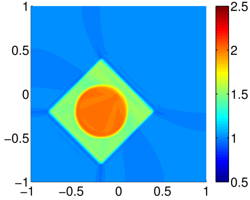

Full-space result for thermal inversion

Study inaccurate solution operator, as preconditioner$^{(*)}$. Apply the inexact full-space TR-SQP algorithm. Use ML to apply $A(z)^{-1}$.

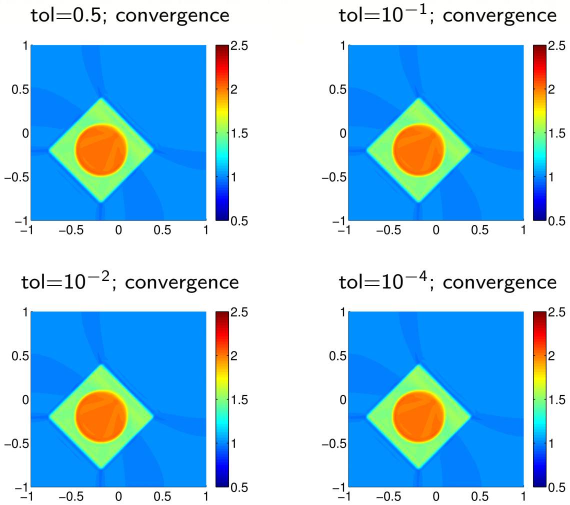

Full-space inversion in 3D

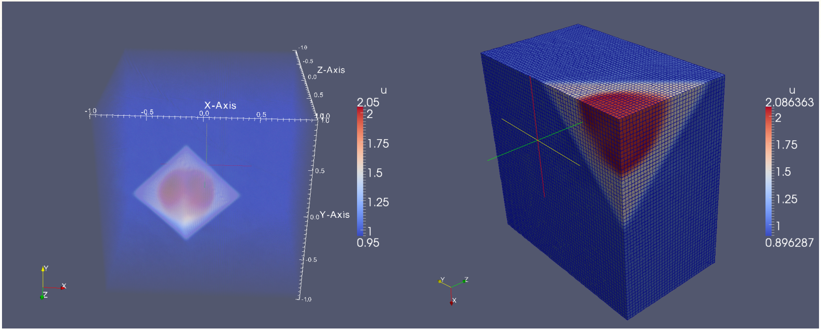

- Setup similar to the 2D example.
- One million elements, runs on my workstation.
- Converges to 10^{-16} in 22 SQP iterations and $\approx$ 1300 CG iterations.
- A single V-cycle of multigrid used to apply $A(z)^{-1}$.
- Parallelizes as well as ML does.

Calibration of electrical circuit models

Nonlinear constraints, with ill-conditioned constraint Jacobians.

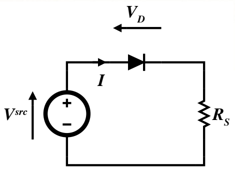

A simple diode.

Shockley diode equation:

$$I = I_S \left(\exp \left(\frac{V^{src} - IR_S}{\eta V^{th}} \right) - 1 \right).$$

Estimate parameters I_s and R_s in a large number of experiments where V^{src} is varied.

— Initial condition 1: $I_S = $ 1e-10, $R_S = 1.0$

Method	#iterations	time (sec)
Reduced space, LS	> 1000	—
Reduced space, TR	204	3.34
Full space, TR-SQP	46	0.10

— Initial condition 2: $I_S = $ 1e-13, $R_S = 0.5$

Method	#iterations	time (sec)
Reduced space, LS	13	0.04
Reduced space, TR	97	1.51
Full space, TR-SQP	23	0.08

Motivation: A simple inverse problem

Inexact full-space methods: TR-SQP, LS-SQP, TR-RSQP

Preconditioners for TR-SQP in PDE-constrained optimization

A hierarchy of linear systems and solvers in full-space methods

The augmented system in PDE optimization

- Reintroduce **state variables** u and **control variables** z:

$$\begin{aligned} \min_{u,z} \quad & f(u,z) \\ \text{s.t.} \quad & c(u,z) = 0 \end{aligned}$$

- Write augmented system matrices as 3×3 block matrices

$$\begin{pmatrix} I & 0 & c_u(u_k, z_k)^* \\ 0 & I & c_z(u_k, z_k)^* \\ c_u(u_k, z_k) & c_z(u_k, z_k) & 0 \end{pmatrix}$$

- Compress notation:

$$\begin{pmatrix} I & 0 & C_u^T \\ 0 & I & C_z^T \\ C_u & C_z & 0 \end{pmatrix}$$

Schur preconditioners

Consider the **exact** and the **approximate** preconditioners, resp.:

$$\mathcal{P}^* = \begin{pmatrix} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & (C_u C_u^T + C_z C_z^T)^{-1} \end{pmatrix} \quad \text{and} \quad \mathcal{P} = \begin{pmatrix} I & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & (C_u C_u^T)^{-1} \end{pmatrix}$$

- $\mathcal{P}^*$-preconditioned GMRES converges in **three iterations**.
- $\mathcal{P}$ amounts to applying C_u^{-1} and C_u^{-T}, i.e., **forward/adjoint solves**.
- These forward/adjoint solves can be **very coarse**!
- Documented **physics-independent** performance!

A recent result for the Helmholtz equation

Mesh \ ω	112.5	225	450	900	1800	3600
50×50	**7.8**					
100×100	7.4	**6.3**				
200×200	6.7	6.0	**5.3**			
400×400	5.5	5.1	4.6	**4.4**		
800×800	4.7	4.4	4.5	4.4	**3.9**	
1600×1600	4.5	4.4	4.5	3.5	3.1	**2.7**

Theorem (Tsuji/Kouri/Ridzal/Tuminaro)

Under suitable assumptions, the eigenvalues μ of the preconditioned system $\mathcal{PA}$ satisfy:

$$\begin{aligned} &\textit{either} && \mu = 1, \\ &\textit{or} && \frac{1}{2}\left(1+\sqrt{5}\right) \le \mu \le \frac{1}{2}\left(1+\sqrt{5+a_1 c^2(\omega)}\right), \\ &\textit{or} && \frac{1}{2}\left(1-\sqrt{5+a_1 c^2(\omega)}\right) \le \mu \le \frac{1}{2}\left(1-\sqrt{5}\right). \end{aligned}$$

where a_1 is a positive constant independent of the discretization parameters, ω is the system frequency, and $c(\omega) \sim 1/\omega$.

Motivation: A simple inverse problem

Inexact full-space methods: TR-SQP, LS-SQP, TR-RSQP

Preconditioners for TR-SQP in PDE-constrained optimization

A hierarchy of linear systems and solvers in full-space methods

Linear systems in inexact SQP methods for PDE-constrained optimization

Inexact line-search SQP:

$$\begin{pmatrix} H_{11} & H_{12} & C_u^T \\ H_{21} & H_{22} & C_z^T \\ C_u & C_z & 0 \end{pmatrix} \begin{pmatrix} u \\ z \\ \lambda \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} + \begin{pmatrix} e_1 \\ e_2 \\ e_3 \end{pmatrix}$$

Linear systems in inexact SQP methods for PDE-constrained optimization

Inexact trust-region SQP:

$$\begin{pmatrix} I & & C_u^T \\ & I & C_z^T \\ C_u & C_z & 0 \end{pmatrix} \begin{pmatrix} u \\ z \\ \lambda \end{pmatrix} = \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} + \begin{pmatrix} e_1 \\ e_2 \\ e_3 \end{pmatrix}$$

Linear systems in inexact SQP methods for PDE-constrained optimization

Inexact trust-region "reduced" SQP:

$$\begin{pmatrix} & C_u^T \\ C_u & \end{pmatrix} \begin{pmatrix} u \\ \lambda \end{pmatrix} = \begin{pmatrix} b_1 \\ b_3 \end{pmatrix} + \begin{pmatrix} e_1 \\ e_3 \end{pmatrix}$$

Summary of inexact SQP methods

Method	Linear systems	Linear solves	Indefinite Hessian	Rank deficiency
LS-SQP	KKT systems	Specialized KKT solvers; can combine constraint preconditioning with certain objective functions	Inertia correction	Hybrid methods, using a general composite step strategy
TR-SQP	Augmented constraint systems	Constraint preconditioning through linearized state and adjoint solves; can use specialized KKT solvers	Conjugate gradients with Steihaug-Toint stopping conditions	Built-in, through general composite steps
TR-RSQP	Constraint systems	Linearized state and adjoint solves	Conjugate gradients with Steihaug-Toint stopping conditions	N/A

- TR-SQP is implemented in the Rapid Optimization Library (ROL).
- The constraint (Schur) preconditioner is also available.
- Currently implementing LS-SQP and TR-RSQP.