Finite Elements in Analysis and Design 87 (2014) 56–73

Contents lists available at ScienceDirect

Finite Elements in Analysis and Design

journal homepage: www.elsevier.com/locate/finel

A parallel finite-element framework for large-scale gradient-based design optimization of high-performance structures

Graeme J. Kennedy [a,*], Joaquim R.R.A. Martins [b]

[a] Georgia Institute of Technology, School of Aerospace Engineering, Atlanta, GA 30332-0150, United States
[b] University of Michigan, Department of Aerospace Engineering, Ann Arbor, MI 48109, United States

ARTICLE INFO

Article history:
Received 6 November 2013
Received in revised form
30 March 2014
Accepted 28 April 2014
Available online 24 May 2014

Keywords:
Gradient-based optimization
Parallel computing
High-fidelity design optimization

ABSTRACT

Structural optimization using gradient-based methods is a powerful design technique that is well suited for the design of high-performance structures. However, the ever-increasing complexity of finite-element models and design formulations results in a bottleneck in the computation of the gradients required for the design optimization. Furthermore, in light of current high-performance computing trends, any methods intended to address this bottleneck must efficiently utilize parallel computing resources. Therefore, there is a need for solution and gradient evaluation methods that scale well with the number of design variables, constraints, and processors. We address this need by developing an integrated parallel finite-element analysis tool for gradient-based design optimization that is designed to use specialized parallel solution methods to solve large-scale high-fidelity structural optimization problems with thousands of design variables, millions of state variables, and hundreds of load cases. We describe the most relevant details of the parallel algorithms used within the tool. We present consistent constraint formulations and aggregation techniques for both material failure and buckling constraints. To demonstrate both the solution and functional accuracy, we compare our results to an exact solution of a pressure-loaded cylinder made with either isotropic or orthotropic material. To demonstrate the parallel solution and gradient evaluation performance, we perform a structural analysis and gradient evaluation for a large transport aircraft wing with over 5.44 million unknowns. The results show near-ideal scalability of the structural solution and gradient computation with the number of design variables, constraints, and processors, which makes this framework well suited for large-scale high-fidelity structural design optimization.

1. Introduction

Over the past few decades, increasingly powerful high-performance computational resources and the development of sophisticated numerical algorithms have enabled the solution of large-scale, high-fidelity structural design optimization problems [63]. We use the term *large-scale* to refer to design problems with a large number of design variables, structural state variables, load cases, or constraint functions, or some combination thereof, such that significant high-performance parallel computing resources are required to solve the problem within a reasonable time. This definition will change with advances in high-performance computing hardware. At present, this definition of large-scale translates to design problems with more than $\mathcal{O}(10^5)$ design variables, $\mathcal{O}(10^6)$ state variables, or $\mathcal{O}(10^2)$ load cases. Several authors have presented solution methods for large-scale problems, including structural shape and sizing problems [52,54,53], large 3D topology problems with $\mathcal{O}(10^6)$ design variables [12,64], and high-fidelity multidisciplinary design optimization problems involving structural analysis as a discipline with $\mathcal{O}(10^6)$ state variables [33].

In this paper, we present an integrated approach to parallel analysis and gradient evaluation for large-scale structural design optimization problems. We have developed this framework for the analysis and design of the thin-shell structures that are used in many high-performance aerospace applications where strength, weight, and stiffness are critical design considerations. These aerospace structures are manufactured using high-performance materials, such as laminated composites or advanced metallic alloys, that achieve high stiffness-to-weight and strength-to-weight ratios. As a result, the design problem involves the simultaneous consideration of the geometry of the structure, the

* Corresponding author.
E-mail addresses: graeme.kennedy@aerospace.gatech.edu (G.J. Kennedy), jrram@umich.edu (J.R.R.A. Martins).

http://dx.doi.org/10.1016/j.finel.2014.04.011
0168-874X/

sizing of the members and, in the case of composites, manufacturing details such as the lamination stacking sequence [1]. Therefore, the structural design problem may include stacking sequence design optimization schemes that can significantly increase the dimensionality of the design space [62,23,24,31]. In addition, slender shell structures subjected to in-plane loading are susceptible to buckling and, as a result, the structural requirements frequently includes both strength and buckling constraints. Within this framework, we impose the buckling constraints using a global–local analysis approach in which a global model determines the edge-loads for a local stiffened panel buckling problem. Other authors have reformulated these design problems using a bilevel approach where the global design problem determines the thicknesses, and the local design problem determines the lamination sequence [40,41,21,44].

Although gradient-free optimization methods have been successfully applied to many important structural design problems, including lamination stacking sequence design [35,42,2] and sizing and shape optimization problems [25], these applications involve at most $\mathcal{O}(10^2)$ design variables. Gradient-free methods are easy to use since they require only function values, but they scale very poorly with the dimensionality of the design space. Since we focus on large-scale high-fidelity applications, we use gradient-based methods and tackle the challenge of efficiently evaluating the gradient of the objective and constraint functions in the design optimization problem.

When we evaluate the gradients required for optimization there are two main concerns: computational time and accuracy. Long gradient computational times might limit the number of optimization iterations that can be performed, while low accuracy might limit the ability of the optimizer to solve the optimization problem to a tight convergence tolerance.

For large-scale structural design problems with large numbers of design variables and constraints, the computational time required to compute the gradients exceeds the computational time required for the analysis. Therefore, gradient evaluation is frequently the computational bottleneck [63]. To minimize the gradient computational time, we use the adjoint method, which requires additional computational time for each gradient. However, the overall gradient evaluation time scales very weakly with the number of design variables. We detail the costs of our adjoint implementation in Section 5.3.

To address gradient accuracy, we focus on minimizing the amount of computational time required to evaluate the gradient to the maximum accuracy possible. To achieve this goal, we do not use finite-difference methods to evaluate the derivatives. Instead, we exclusively use hand-coded derivative routines that achieve good computational performance while avoiding the subtractive cancelation issues suffered by finite-difference methods. We note that other authors have used automatic differentiation methods, rather than hand-coded routines, to obtain accurate derivatives [47]. To verify the accuracy of our derivative implementation, we use the complex-step derivative evaluation technique. The complex-step method uses a complex perturbation of the variables to determine the derivative and does not suffer from subtractive cancellation. As a result, a very small step size may be used, yielding derivatives with the same number of significant digits as the functional estimate [61,49].

In this paper we present a fully verified, integrated framework for the parallel analysis and gradient-evaluation of shell structures. We verify that our methods achieve the optimal solution and functional accuracy. In our opinion, within the context of design optimization, functional accuracy and solution accuracy are of equal importance, yet functional accuracy is often overlooked in structural optimization applications. In addition, we verify our gradient evaluation methods using a complex-step derivative approximation technique that enables accurate verification of the derivatives without loss of accuracy due to subtractive cancellation. We have integrated the developments presented in this paper into a sophisticated parallel finite-element code that we call the Toolkit for the Analysis of Composite Structures (TACS). We have used TACS for large-scale structural analysis of composite beams [29], structural topology optimization [36,37], lamination sequence design [31], and both static and dynamic aeroelastic design optimization [28,33,30,39,32].

1.1. The model optimization problem

Since structural weight reduction is critical in many aerospace applications, the most common structural design problem is to minimize the structural mass subject to stress and possibly buckling constraints. These structural constraints are imposed at a series of design load cases to ensure the safety of the aerospace vehicle within a prescribed operational envelope. With this standard structural design optimization problem in mind, we pose the following generic structural design optimization problem:

$$\begin{aligned}
&\text{minimize} && f(\mathbf{x}, \mathbf{u}_1, \ldots, \mathbf{u}_{n_\ell}) \\
&\text{with respect to} && \mathbf{x},\ \mathbf{u}_1, \ldots,\ \mathbf{u}_{n_\ell} \\
&\text{governed by} && \mathbf{R}_i(\mathbf{X}^N(\mathbf{x}_G), \mathbf{x}_M, \mathbf{u}_i) = 0 \\
&\text{for } 1 \le i \le n_\ell \text{ such that} && \mathbf{f}_i(\mathbf{x}, \mathbf{u}_i) \le 1 \qquad \mathbf{x}_l \le \mathbf{x} \le \mathbf{x}_u
\end{aligned} \tag{1}$$

where $f(\mathbf{x}, \mathbf{u}_1, \ldots, \mathbf{u}_{n_\ell})$ is the objective function and $\mathbf{f}_i(\mathbf{x}, \mathbf{u}_i) \in \mathbb{R}^{n_f}$ represents a vector of constraints for the ith load case. Note that there are a total of n_ℓ load cases. The design variables $\mathbf{x} = (\mathbf{x}_G, \mathbf{x}_M) \in \mathbb{R}^{n_x}$ are partitioned into either geometric or material design variables that we denote $\mathbf{x}_G \in \mathbb{R}^{n_{xg}}$ and $\mathbf{x}_M \in \mathbb{R}^{n_{xm}}$, respectively. The distinction between geometric and material design variables arises at the element level: geometric design variables modify the element nodes, and material design variables modify the element constitutive behavior. This distinction between design variables is critical for high-performance derivative evaluation methods. Treating the geometric and material variables uniformly would lead to computationally inefficient derivative evaluation methods. The finite-element residuals $\mathbf{R}_i \in \mathbb{R}^{6n}$ depend on the finite-element nodal locations $\mathbf{X}^N(\mathbf{x}_G) \in \mathbb{R}^{3n}$, the material design variables $\mathbf{x}_M$, and the state variables $\mathbf{u}_i \in \mathbb{R}^{6n}$, for the i^{th} load case. Note that we write the nodal locations as functions of the geometric design variables to ensure that the partial derivative $\partial \mathbf{X}^N / \partial \mathbf{x}_G$ is computed only once in each gradient evaluation.

While there are numerous techniques available for solving the optimization problem (1), we employ a reduced-space approach where the governing equations for each load case, $\mathbf{R}_i(\mathbf{X}^N(\mathbf{x}_G), \mathbf{x}_M, \mathbf{u}_i) = 0$, are solved at each optimization iteration, and the optimization problem is recast solely in terms of design variables. In the reduced-space approach, which is also referred to as the nested analysis and design (NAND) architecture [20], the state variables are implicit functions of the design variables, and the adjoint or direct method must be used to determine the objective and constraint gradients. The reduced-space method is an alternative to full-space approaches that solve the design and analysis problems simultaneously [20,9,10]. We do not solve the optimization problem (1) directly using our framework, but we instead provide the objective and constraint values and gradients to a gradient-based optimizer. Typically, we solve the optimization problem with the Python-based optimization interface pyOpt [57], which provides access to several optimization packages. The optimization package we use here is SNOPT [18]. While there is ongoing work to parallelize optimization algorithms, especially in the context of PDE-constraint optimization [9], for many