SIAG/OPT Views and News

A Forum for the SIAM Activity Group on Optimization

Volume 23 Number 1 April 2015

Contents

SIAM OP14

Introduction by the OP14 Co-Chairs

Miguel F. Anjos
Département de Mathématiques et de Génie Industriel
École Polytechnique de Montréal
Canada
miguel-f.anjos@polymtl.ca
http://anjos.mgi.polymtl.ca

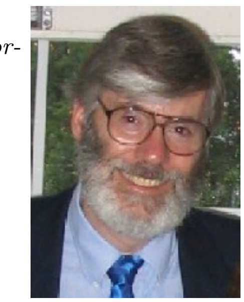

Michael J. Todd
School of Operations Research and Information Engineering
Cornell University
Ithaca, NY 14850
USA
miketodd@orie.cornell.edu
http://people.orie.cornell.edu/miketodd

The SIAM Optimization Conference in San Diego last May attracted over 650 researchers, with about half from the United States and sizable contingents from Germany (81), Canada (43), France (29), Japan (24), and the U.K. (24). In total, 31 countries were represented. Attendees were treated to invited talks from leading experts in optimization, randing from theoretical to computational to applied and from continuous to discrete. These talks, together with the SIAM Prize lecture by Yinyu Ye, the minitutorials by Monique Laurent and Didier Henrion and by Jeff Linderoth, Sven Leyffer, and James Luedtke, and the minisymposia, are available online at http://goo.gl/kwfrna.

We thought it would be useful to invite the plenary speakers to prepare short articles related to their presentations for *SIOPT Views and News*. Joaquim Martins and Franz Rendl accepted our invitation. Their articles, on wing design via numerical optimization and on relaxations for some NP-hard problems, respectively, show the breadth and depth of our field. We hope that those of you who were not able to come to San Diego, as well as those who would like a refresher after ten months, will enjoy these fascinating articles.

Two of our speakers, Andy Philpott and Rekha Thomas, were invited to write similar articles for *SIAM News*, which are excellent publicity for our activity group. Rekha's article is available at http://goo.gl/YggnKi and Andy's article is available at http://goo.gl/XabauI.

Wing Design via Numerical Optimization

Joaquim R. R. A. Martins
Department of Aerospace Engineering
University of Michigan
Ann Arbor, MI 48109
USA
jrram@umich.edu
http://mdolab.engin.umich.edu/martins

1 Introduction

With regard to flight, the wing is arguably the most crucial component. As the legendary Boeing aircraft designer Jack Steiner put it, "The wing is where you're going to fail." In a book detailing the origins of the Boeing 747 Jumbo Jet, Irving [5] writes:

> Designing the wing involved literally thousands of decisions that could add up to an invaluable asset, a proprietary store of knowledge. A competitor could look at the wing, measure it even, and make a good guess about its internal structure. But a wing has as many invisible tricks built into its shape as a Savile Row suit; you would need to tear it apart and study every strand to figure out its secrets.

These "invisible tricks" are a reflection of the complexity involved in modeling the physics governing wing performance. The function of the wing is to provide enough lift to counteract the aircraft weight, while producing the least amount of drag (which lowers the required engine thrust). Lift and drag can be predicted through aerodynamic models that vary in sophistication and computational effort. For the flight speeds of commercial airliners (78–86% the speed of sound, or 830–925 km/h), the aerodynamic flow is compressible, and the wings usually generate shock waves. This situation, together with the fact that the wing flexibility couples the aerodynamic shape to the structural layout and sizing, contributes to the "invisible tricks" mentioned above.

We must be able to model before we optimize. To model the lift and drag accurately at transonic speeds where shocks are present requires computational fluid dynamics (CFD), which solves PDEs over the three-dimensional domain. To model the flexibility of the wing, we must couple the aerodynamic model with a structural model that predicts the deflected wing shape given the aerodynamic loads. Thus the complete wing model typically involves solving a multiphysics PDE model with at least $\mathcal{O}(10^6)$ unknowns.

The "thousands of decisions" cited in the above quote can be mapped to design variables, which involve both aerodynamic shape and structural design variables. The aerodynamic flow (and hence lift and drag) is sensitive to the slightest change in aerodynamic shape, so one must parameterize the shape with a large number of local changes.

A truly practical objective function for aircraft design is difficult to define because it depends on the balance between acquisition cost and aircraft performance. This balance depends on the business model of the particular airline, as well as on the current price of fuel. Acquisition cost is notoriously difficult to model. Aircraft performance can be modeled as operating cost, which depends on two main factors: the speed and the fuel consumption. The faster the airplane can fly, the lower the costs associated with time (e.g., crew salaries) and the more productive it can be by moving more passengers. Beyond a certain point, however, speed comes at the cost of greater fuel consumption.

When optimizing both aerodynamics and structures, we need to consider the effect of the aerodynamic shape variables and structural sizing variables on the weight, which also affects the fuel burn. Thus complex multidisciplinary trade-offs are involved in such an objective function. Numerical optimization is a powerful tool that can perform these trade-offs automatically. Aerospace engineering researchers recognized this as soon as multiphysics models for wings were available, establishing the field of multidisciplinary design optimization (MDO) [4, 12]. So far, the MDO of aircraft has involved mostly low-fidelity models that are based on either simplified physics or empirical models, with few design variables and constraints.

In this article, we show a wing design example where we tackle the compounding challenges of modeling the wing with large systems of coupled PDEs while optimizing it with respect to hundreds of design variables. We are able to meet these challenges successfully through the use of high-performance parallel computing, fast coupled PDE solvers, state-of-the-art gradient-based optimization, and an efficient approach for computing the coupled derivatives for the PDEs.

2 Optimization Problem

As we mentioned, determining the real objective function in aircraft design is difficult because of the variability in the cost of time, fuel price, and airline routes. We avoid this issue by choosing the fuel burn as the objective function to be minimized. However, the methods presented here are applicable to any other objective function.

The design variables are wing shape and structural sizing parameters, as shown in Fig. 1. Two main groups of wing shape variables exist: those that define the planform shape, and those that define the airfoil sections. The planform variables determine what the wing looks like when viewed from above. We use area, sweep, span, and taper to define the planform shape. The airfoil shape requires $\mathcal{O}(10^2)$ variables so that enough freedom is provided to reduce the aerodynamic drag. Typically, $\mathcal{O}(10^1)$ airfoil sections are distributed in the spanwise direction, each of which is allowed to change its shape independently. The wing shape is then obtained by performing an interpolation in the spanwise direction.

The shape modifications due to these shape variables are applied by using free-form deformation (FFD) [1]. This approach consists in defining a volume that encloses the wing geometry and then manipulating the surface of the volume, which changes the inside of the volume continuously. The FFD variables change both the aerodynamic surface and the structure inside the wing.

The structure inside the wing, called the wing box, usually consists of a grid of spars (laid out in the spanwise direction), ribs (laid out perpendicularly to the spars), and skins that cover the wing. All these elements are thin shells, and the structural sizing variables are the thicknesses of these shells. All sizing variables are subject to constraints on the variation in thickness of adjacent elements for manufacturing reasons.

The design variables are listed in Table 1. In addition to the wing shape and structural sizing, the angle of attack is included as a design variable in order to provide the optimizer with a way to satisfy the lift constraint.

Most of the constraints in this wing design problem are there to ensure that the wing is strong enough to sustain certain maneuvers without structural failure. We consider two maneuvers: a 2.5 g pull-up maneuver and a −1 g push-over maneuver. We prevent structural failure by constraining the stress in the structure to stay below the yield stress of the material and by constraining the structure from buckling at the allowable loads. An aggregation function is used to handle these constraints [13].

The objective functions and constraints in our wing design optimization problem (Table 1) are nonlinear,

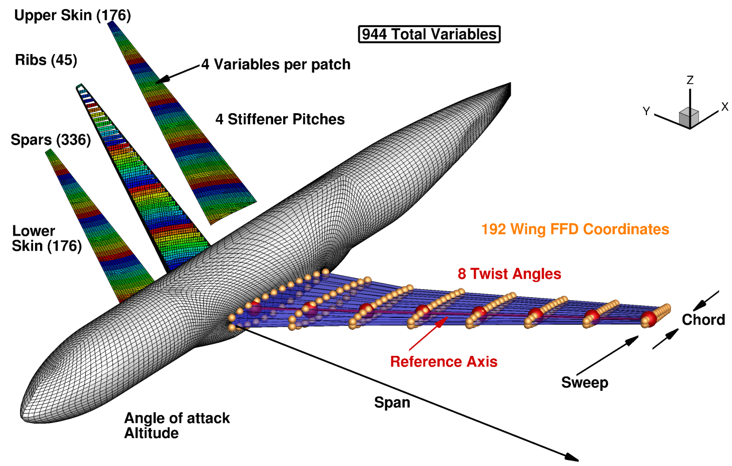

Figure 1: Wing aerostructural design variables [7].

with the exception of the adjacency and geometric constraints, which are linear. These functions are also nonconvex in general; but because of the complexity of the functions involved and the cost of the coupled PDE solutions, we currently cannot prove global optimality. However, we have studied the existence of local minima in aerodynamic shape optimization [9].

The models for the coupled aerodynamic and structural PDEs that need to be solved in order to evaluate the objective and constraints are not included explicitly in the optimization constraints because they are solved with specialized algorithms. This constitutes a reduced-space approach to a PDE-constrained optimization problem. At each optimization iteration, the aerostructural solver computes the objectives and constraints for the given set of design variables.

3 Computational Models

The physics of the wing must be modeled by coupling an aerodynamics model that computes the flow field (along with the corresponding drag and lift) and a structural model that computes the wing displacement field (along with the corresponding stress field and buckling parameters).

Here we consider high-fidelity models in the form of the Reynolds-averaged Navier–Stokes (RANS) PDEs, which can model transonic flow with shocks and provides drag estimates that include both pressure drag and skin friction drag. To solve the RANS equations, we use a finite-volume, cell-centered multiblock solver [14]. The main flow is solved by using an alternating direction implicit (ADI) method method along with geometric multigrid. A segregated Spalart–Allmaras turbulence equation is iterated with the diago-

Table 1: Aerostructural wing design optimization problem (adapted from [7]).

	Function/Variable	Quantity
minimize	Fuel burn	
with respect to	Wing span	1
	Wing sweep	1
	Wing chord	1
	Wing twist	8
	FFD control point vertical position	192
	Angle of attack at each flight condition	3
	Cruise altitude	1
	Upper and lower stiffener pitch	2
	Leading and trailing edge spar stiffener pitch	2
	Rib thickness	45
	Panel thickness for skins and spars	172
	Panel stiffener thickness for skins and spars	172
	Panel stiffener height for skins and spars	172
	Panel length for skin and spars	172
	Total number of design variables	**944**
subject to	Lift=weight at each flight condition	3
	Lift coefficient ≤ 0.525 to ensure buffet margin	1
	Leading edge thickness must not decrease	20
	Trailing edge thickness must not decrease	20
	Trailing edge spar height must not be less than 80% of the initial	20
	Wing planform area must be greater than or equal to initial	1
	Wing fuel volume must be greater than or equal to initial	1
	Panel length variable must match wing geometry	172
	Aggregate stress must not exceed the yield stress at 2.5 and -1 g	4
	Aggregate buckling must not exceed the critical value at 2.5 and -1 g	4
	Thickness must not vary by more than 2.5 mm between elements	504
	Leading and trailing edge displacement constraint	16
	Total number of constraints	**938**

nally dominant alternating direction implicit (DDADI) method.

We solve the RANS equations in the three-dimensional domain surrounding the aircraft. The computation of the drag, lift, and moment coefficients consists in the numerical integration of the flow pressure and shear stress distribution on the surface of the aircraft.

The structural solver is a parallel direct solver that uses a Schur complement decomposition [6]. For the thin-shell problems typical of aircraft structures, we often have matrix condition numbers $\mathcal{O}(10^9)$, but this solver is able to handle such problems.

The coupled aerostructural system is solved by using nonlinear block Gauss–Seidel with Aitken acceleration, which has proved to be robust for the range of flight conditions considered [8].

4 Optimization Algorithm

In selecting an optimization algorithm, two fundamental choices exist: gradient-free or gradient-based methods. Our wing design optimization application faces two compounding challenges: large numbers of design variables ($\mathcal{O}(10^2)$ or more) and a high cost of evaluating the objective and constraints (which involve the solution of coupled PDEs with $\mathcal{O}(10^6)$ variables). Since the number of iterations required by gradient-free methods does not scale well with the number of optimization variables, we use a gradient-based method. In particular, we use SNOPT [3], an implementation of the sequential quadratic programming algorithm suitable for general nonlinear constrained problems. Given the efficiency of gradient-based methods, we can address the two compounding challenges mentioned above, provided we can evaluate the required gradients efficiently.

5 Computing Gradients

With a gradient-based optimizer, the efficiency of the overall optimization hinges on an efficient evaluation of the gradients of the objective and constraint functions with respect to the design variables. Several methods are available for evaluating derivatives of PDE systems: finite differences, the complex-step method, algorithmic differentiation (forward or reverse mode), and analytic methods (direct or adjoint) [11]. The computational cost of these methods is proportional either to the number of design variables, or to the number of functions being differentiated.

Since we have a large number of design variables, the best options are the reverse-mode algorithmic differentiation or the adjoint method. In our applications we tend to use a hybrid approach that combines the adjoint method with algorithmic differentiation (both reverse and forward modes).

We now derive the adjoint method for evaluating the derivatives of a function of interest, $f(x, y(x))$ (which in our case are the objective function and constraints), with respect to the design variables x. The state variable vector y is determined implicitly by the solution of the PDEs, $R(x, y(x)) = 0$, for a given x. Using the chain rule, we calculate the gradient of f with respect to x:

$$\frac{\mathrm{d}f}{\mathrm{d}x} = \frac{\partial f}{\partial x} + \frac{\partial f}{\partial y}\frac{\mathrm{d}y}{\mathrm{d}x}. \tag{1}$$

A similar expression can be written for the Jacobian of R:

$$\frac{\mathrm{d}R}{\mathrm{d}x} = \frac{\partial R}{\partial x} + \frac{\partial R}{\partial y}\frac{\mathrm{d}y}{\mathrm{d}x} = 0.$$

We can now solve this linear system to evaluate the gradients of the state variables with respect to the design variables. Substituting this solution into the evaluation of the gradient of f (1) yields

$$\frac{\mathrm{d}f}{\mathrm{d}x} = \frac{\partial f}{\partial x} - \frac{\partial f}{\partial y}\left[\frac{\partial R}{\partial y}\right]^{-1}\frac{\partial R}{\partial x}.$$

The adjoint method consists of factorizing the Jacobian $\partial R/\partial y$ with $\partial f/\partial y$. That is, we solve the adjoint equations

$$\left[\frac{\partial R}{\partial y}\right]^{T}\psi = -\frac{\partial f}{\partial y}, \tag{2}$$

where ψ is the adjoint vector. We can then substitute the result into the total gradient equation (1),

$$\frac{\mathrm{d}f}{\mathrm{d}x} = \frac{\partial f}{\partial x} + \psi^{T}\frac{\partial R}{\partial x}, \tag{3}$$

to get the required gradient. The partial derivatives in these equations are inexpensive to evaluate, since they do not require the solution of the PDEs. The computational cost of evaluating gradients with the adjoint method is independent of the number of design variables but dependent on the number of functions of interest. Thus, this method is efficient when considering the wing design problem defined in Sec. 2, which has 944 design variables and 14 nonlinear constraints (the other 924 constraints are linear, and thus their Jacobian is constant).

The discrete adjoint solver for our CFD model was developed by forming Eqs. (2) and (3), where the partial derivatives are implemented by performing algorithmic differentiation in the relevant parts of the original code [10]. A discrete adjoint method is also implemented in our structural solver [6].

The adjoint method can be extended to coupled systems, such as the aerostructural system of equations considered here [2, 8]. For the implementation of the coupled adjoint to be efficient, we ensured that the computation of each of the partial derivatives in Eqs. (2) and (3) scales well with the number of processors [8]. The coupled adjoint equations are solved by using a coupled Krylov method, which converges faster than the linear block Gauss–Seidel method [8].

6 Wing Design Optimization

We now present the solution to the design optimization problem described in Sec. 2. The initial aircraft geometry is the Common Research Model (CRM) configuration [15], which is representative of a twin-aisle long-range airliner. The CFD solver uses a structured volume grid with 745,472 cells, resulting in more than 4.47 million degrees of freedom, while the wing box structural model has 190,710 degrees of freedom.

The planform and front views of this aircraft are shown on the left side of the geometry shown in the upper left quadrant of Fig. 2. The right side of this geometry shows the optimized aircraft. The pressure coefficient contours shown on the initial wing (left) are closely spaced in the outboard area near the trailing edge, indicating a shock wave, while the optimized wing (right) shows evenly spaced contours and no shock. The optimization reduced the drag while incurring a weight penalty, resulting in a net reduction in fuel burn. The front view of the aircraft shows the deflected shapes of the wings for both the cruise and maneuver conditions.

The upper right quadrant shows the wing structural box. The top two wings show a color map of the struc-

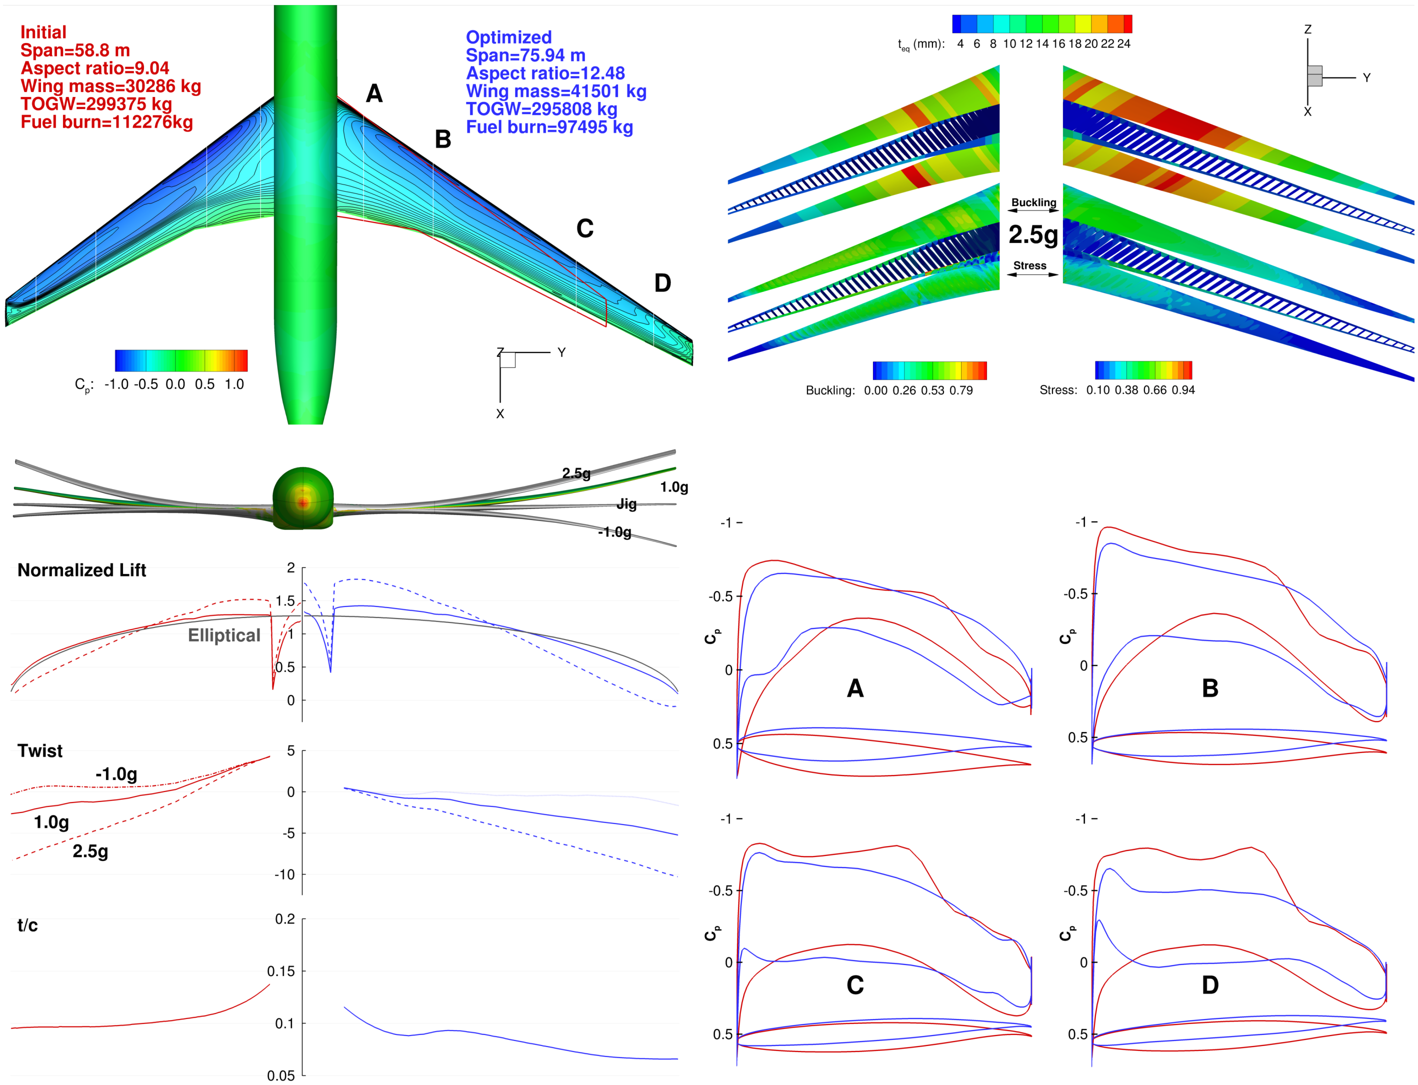

Figure 2: Initial wing design (left/red) and aerostructurally optimized wing (right/blue), showing planform view and front view (top left), wing box structure (top right), spanwise lift, twist and thickness distributions (bottom left), and airfoil sections with pressures (bottom right).

tural thickness distributions for the initial (left) and optimized (right) wing. The right wing shows the higher thicknesses that are required to strengthen the higher span wing. The bottom two wings show the values for the stress and buckling constraints, which are under the critical values (i.e., less than 1.0).

The bottom right quadrant shows four airfoil sections of the wing from the root (A) to the wingtip (D). The initial airfoils are shown in red, and the optimized airfoils are shown in blue, together with the respective pressure distributions.

7 Conclusion

In this article, we introduced a wing design problem where physics-based models of both the aerodynamics and structures were needed. Such a problem is subject to the compounding challenges of modeling the wing with large systems of coupled PDEs while optimizing the wing with respect to hundreds of design variables. We were able to tackle this problem through the use of high-performance parallel computing to solve the model PDEs, a nonlinear block Gauss–Seidel method for solving the coupled system, an SQP optimizer, and a coupled adjoint approach for computing the derivatives of the coupled PDEs. This proved to be a powerful combination that should be applicable to many other multiphysics design optimization problems.

We demonstrated these techniques in the design optimization of a large transport aircraft. The optimizer was able to tradeoff aerodynamic drag and structural weight in just the right proportions to achieve the lowest possible fuel burn. Almost one thousand geometric shape and structural sizing variables were optimized subject to a similar number of constraints. While a number of constraints still need to be considered before these results can be directly used by aircraft manufacturers, we have demonstrated the feasibility of performing wing design optimization by using high-fidelity multiphysics models.

Acknowledgments. This work was partially supported by NASA under grant number NNX11AI19A. The author thanks Gaetan Kenway and Graeme Kennedy for the results presented in this article.

REFERENCES

[1] M. Andreoli, A. Janka, and J.-A. Désidéri. Free-form-deformation parameterization for multilevel 3D shape optimization in aerodynamics. Tech. Report 5019, INRIA, Sophia Antipolis, France, November 2003.

[2] O. Ghattas and X. Li. Domain decomposition methods for sensitivity analysis of a nonlinear aeroelasticity problem. *In. J. Comput. Fluid D.*, 11:113–130, 1998.

[3] P. E. Gill, W. Murray, and M. A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Rev.*, 47(1):99–131, 2005.

[4] R. T. Haftka. Optimization of flexible wing structures subject to strength and induced drag constraints. *AIAA J.*, 14(8):1106–1977, 1977.

[5] C. Irving. *Wide-Body—The Triumph of the 747*. William Morrow and Company, New York, 1993.

[6] G. J. Kennedy and J. R. R. A. Martins. A parallel finite-element framework for large-scale gradient-based design optimization of high-performance structures. *Finite Elem. Anal. Des.*, 87:56–73, September 2014.

[7] G. K. W. Kenway, G. J. Kennedy, and J. R. R. A. Martins. Aerostructural optimization of the common research model configuration. In *15th AIAA/ISSMO Multidiscip. Anal. Optim. Conf.*, Atlanta, GA, June 2014.

[8] G. K. W. Kenway, G. J. Kennedy, and J. R. R. A. Martins. Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and derivative computations. *AIAA J.*, 52(5):935–951, May 2014.

[9] Z. Lyu, G. K. Kenway, and J. R. R. A. Martins. Aerodynamic shape optimization studies on the Common Research Model wing benchmark. *AIAA J.*, 2014. (In press).

[10] Z. Lyu, G. K. Kenway, C. Paige, and J. R. R. A. Martins. Automatic differentiation adjoint of the Reynolds-averaged Navier–Stokes equations with a turbulence model. In *21st AIAA Comput. Fluid Dyn. Conf.*, San Diego, July 2013.

[11] J. R. R. A. Martins and J. T. Hwang. Review and unification of methods for computing derivatives of multidisciplinary computational models. *AIAA J.*, 51(11):2582–2599, November 2013.

[12] J. R. R. A. Martins and A. B. Lambe. Multidisciplinary design optimization: A survey of architectures. *AIAA J.*, 51(9):2049–2075, September 2013.

[13] N. M. K. Poon and J. R. R. A. Martins. An adaptive approach to constraint aggregation using adjoint sensitivity analysis. *Struct. Multidiscip. Optim.*, 34(1):61–73, 2007.

[14] E. van der Weide, G. Kalitzin, J. Schluter, and J. J. Alonso. Unsteady turbomachinery computations using massively parallel platforms. In *Proc. 44th AIAA Aero. Sci. Meet. Exhibit*, Reno, NV, 2006. AIAA 2006-0421.

[15] J. C. Vassberg, M. A. DeHaan, S. M. Rivers, and R. A. Wahls. Development of a common research model for applied CFD validation studies, 2008. AIAA 2008-6919.

Relaxations for Some NP-Hard Problems Based on Exact Subgraphs

Franz Rendl
Institut für Mathematik
Alpen-Adria Universität Klagenfurt
Austria
Franz.Rendl@aau.at
https://campus.aau.at/org/visitenkarte.jsp?personalnr=2034

1 Max-Cut and Stable-Set

Many classical NP-complete graph optimization problems have relaxations based on semidefinite optimization. Two prominent examples are *Max-Cut* and *Stable-Set*.

We consider the Max-Cut problem in the following form. Given a symmetric matrix L of order n, find

$$z_{MC} := \max\{c^T Lc : \ c \in \{-1, 1\}^n\}.$$

The cut polytope CUT_n is defined as

$$\text{CUT}_n := \text{conv}\{cc^T : c \in \{-1, 1\}^n\}.$$

Clearly, $z_{MC} = \max\{\langle L, X\rangle : X \in \text{CUT}\}$. The cut polytope is contained in the spectrahedron

$$\text{CORR} := \{X : \text{diag}(X) = e, X \succeq 0\},$$

consisting of all correlation matrices, those semidefinite matrices having the all-ones vector e on the main diagonal. Optimizing over CORR yields one of the most well-studied semidefinite optimization problems,

$$z_{\text{CORR}} := \max\{\langle L, X\rangle : X \in \text{CORR}\}. \quad (1)$$

It was introduced (in dual form) by Delorme and Poljak [7]. Goemans and Williamson [8] provided a theoretical error analysis showing that $z_{MC} \geq 0.878 \cdot z_{\text{CORR}}$ for graphs with nonnegative edge weights.

For Stable-Set, we are given a graph G on n nodes. Let $\mathcal{S}$ denote the set of incidence vectors of the stable sets of G. Hence $s \in \mathcal{S}$ is a 0-1 vector, and $s_i s_j = 0$ if $[i,j]$ is an edge of G. Here the task is to find the cardinality of the largest stable set, which is denoted by $\alpha(G)$:

$$\alpha(G) := \max\{e^T s : s \in \mathcal{S}\}.$$

The stable set polytope is $\mathrm{STAB}(G) := \mathrm{conv}\{s : s \in \mathcal{S}\}$. The "squared version" of STAB is given by

$$\mathrm{STAB}^2(G) := \mathrm{conv}\{ss^T : s \in \mathcal{S}\}.$$

Any characteristic vector $s \in \mathcal{S}$ of a stable set in G yields a *stable-set matrix* $S := ss^T$ such that

$$s = \mathrm{diag}(S),$$

$$S - ss^T \succeq 0,$$

$$(S)_{ij} = s_i s_j = 0 \ \forall [i,j] \in E(G).$$

Therefore $\mathrm{STAB}^2(G)$ is contained in the spectrahedron $\mathrm{TH}^2(G)$, which is defined as

$$\begin{aligned}\mathrm{TH}^2(G) \ := \{&X : x = \mathrm{diag}(X), X - xx^T \succeq 0,\\ &x_{ij} = 0 \ \forall [i,j] \in E(G)\}.\end{aligned}$$

Lovász introduced the theta body $\mathrm{TH}(G) := \{x \in I\!R^n : \ \exists X \in \mathrm{TH}^2(G), \ x = \mathrm{diag}(X)\}$ as the projection of matrices in $\mathrm{TH}^2(G)$ to their main diagonal. From the definition it follows that $\mathrm{STAB}(G) \subseteq \mathrm{TH}(G)$. One of the most well-studied relaxations of $\mathrm{STAB}(G)$ is obtained by optimizing over $\mathrm{TH}(G)$. Its optimal value is denoted by $\theta(G)$:

$$\begin{aligned}\theta(G) :=& \ \max\{e^T x : x \in \mathrm{TH}(G)\}\\ =& \ \max\{\mathrm{tr}(X) : X \in \mathrm{TH}^2(G)\}.\end{aligned}$$

This is again a semidefinite program, and $\alpha(G) \le \theta(G)$.

2 The New Hierarchy

Several hierarchies of relaxations have been introduced recently, which result in semidefinite programs of increasing sizes but which ultimately provide an exact formulation of the problem (i.e., result in integer solutions). These hierarchies were introduced by Sherali and Adams [13] and Lovász and Schrijver [12] and have interesting theoretical properties; but they are computationally extremely challenging, since the matrix dimension grows in each level of the hierarchy.

Here we aim at hierarchies of semidefinite relaxations that are computationally accessible. Rather than using the standard hierarchies, where in each level the model increases by an order of magnitude, we maintain the original model size, and only the number of constraints will grow (exponentially). Our hierarchy is based on "exact subgraphs." We concentrate on graph problems that are "downward monotone." This means that the restriction of the problem to a vertex-induced subgraph yields a smaller problem of the same type. We note that Max-Cut and Stable-Set both have this property.

For $I \subseteq N := \{1, \ldots, n\}$ with $|I| = k$ we denote by $\pi_I(X) := X_I$ the restriction of the $n \times n$ matrix X to the $k \times k$ submatrix X_I with rows and columns indexed by I. The restriction of the cut polytope CUT_n to the subgraph induced by I is denoted by

$$\pi_I(\mathrm{CUT}_n) := \mathrm{conv}\{(cc^T)_I : c \in \{-1,1\}^n\}.$$

Downward monotonicity of Max-Cut gives $\pi_I(\mathrm{CUT}_n) = \mathrm{CUT}_k$. A matrix $X \in \mathrm{CUT}_n$ satisfies $\pi_I(X) \in \pi_I(\mathrm{CUT}_n)$, which translates into $X_I \in \mathrm{CUT}_k$, which in turn can be expressed as

$$X_I = \sum_i \lambda_i^I C_i \text{ with } \lambda_i^I \ge 0, \ \sum_i \lambda_i^I = 1.$$

The summation is over all 2^{k-1} cut matrices $C_i := c_i c_i^T$, where $c_i \in \{-1,1\}^k$. We study the following hierarchy of relaxations for Max-Cut. At level $k = 1, \ldots, n$ we consider

$$\begin{aligned}z_{MC,k} := \quad \max\{&\langle L, X\rangle : X \in \mathrm{CORR},\\ &X_I = \textstyle\sum_i \lambda_i^I C_i, \lambda_i^I \ge 0,\\ &\textstyle\sum_i \lambda_i^I = 1 \ \forall I \subseteq N, |I| = k\}.\end{aligned}$$

Clearly $k = 1$ gives the original relaxation (1), the optimal value is nonincreasing with k, and for $k = n$ we get the exact solution of the Max-Cut problem. Feasible matrices in hierarchy k have the property that all submatrices of order k are in the cut polytope and hence represent exact subgraphs.

Turning to the Stable-Set problem, we note that $\pi_I(\mathrm{STAB}^2(G)) = \mathrm{STAB}^2(G_I)$, where G_I denotes the subgraph of G induced by the vertices in I. We investigate the hierarchical refinement of $\theta(G)$ toward $\alpha(G)$, given by

$$\begin{aligned}\theta_k(\mathrm{TH}^2(G)) := \ \max\{&\mathrm{tr}(X) : X \in \mathrm{TH}^2(G),\\ &X_I \in \mathrm{STAB}^2(G_I) \ \forall I \subseteq N,\\ &|I| = k\}.\end{aligned}$$

In this case X_I will be a convex combination of stable-set matrices from the subgraph G_I, induced by I. Again, it is clear that $\theta_1(\mathrm{TH}^2(G)) = \theta(G)$ and $\theta_n(\mathrm{TH}^2(G)) = \alpha(G)$.

3 Related Work

The idea of using smaller polytopes to tighten polyhedral relaxations has a long history in polyhedral combinatorics. Boros, Crama, and Hammer [5] introduced a hierarchy of polyhedral relaxations for Max-Cut that agrees with our hierarchy when started not from CORR but from the metric polytope relaxation. Christof and Reinelt [4] used cuts from small polytopes and applied them to the linear ordering and the betweenness problem. More recently Buchheim, Liers, and Oswald [3] introduced "target cuts" to improve polyhedral relaxations.

From a worst-case point of view, the metric polytope relaxation of Max-Cut is known to have an integrality gap of $2 - \epsilon$. Furthermore, de la Vega and Kenyon-Mathieu [6] showed that for any fixed k, the level k hierarchy of the metric polytope relaxation obtained by including all valid inequalities for the cut polytope on at most k vertices still has an integrality gap of $2 - \epsilon$.

Cutting planes generated from small polytopes have also been used for semidefinite relaxations. Helmberg and Rendl [9] considered the semidefinite relaxation (1) for the Max-Cut problem and combined it with clique inequalities and general hypermetric inequalities from small subgraphs.

4 Subgraph Separation

Including all $\binom{n}{k}$ subset constraints clearly is impractical, even for small values of k (say $k = 5$). We propose to proceed iteratively, including only a carefully selected subset of all candidate sets I. Having solved the resulting relaxation, with current solution X, we need to identify additional subsets to be included.

Thus we need to deal with the following separation problem: Given X and k, find a subset I with $|I| = k$ such that $X_I \notin \mathrm{CUT}_k$ or $X_I \notin \mathrm{STAB}^2(G_I)$, or prove that no such set exists.

Our relaxations have the property that the feasible region is contained not only in the cone of semidefinite matrices but also in the cone of completely positive matrices (consisting of all matrices that can be factorized as AA^T for some nonnegative A). For $\mathrm{STAB}^2(G)$ this is immediate from the definitions. Also Max-Cut has a relaxation based on completely positive matrices, since $\frac{1}{4}c^TLc = z^TLz$, where $c \in \{-1,1\}^n$, $z = \frac{1}{2}(c+e) \in \{0,1\}^n$, and $Le = 0$ is assumed without loss of generality.

To find a violating subset I, we can argue as follows. If X_I is not completely positive, then clearly I is violating. To show that X_I is not completely positive, we need to provide a matrix Y from the dual cone, which is the cone of copositive matrices, such that $\langle X_I, Y\rangle < 0$. The cone of completely positive matrices as well as the dual cone of copositive matrices are intractable. In our situation we are interested in these cones only for small dimensions. For instance, it is known that for $k \le 4$ the cone of copositive matrices is simply the sum of semidefinite and nonnegative matrices. For $k = 5$, Hildebrand [10] has recently given an explicit characterization of the extreme rays of the copositive cone. Moreover, Baston [2] provided a complete description of all copositive matrices with entries only in $\{-1,1\}$. Hoffman and Pereira [11] characterized those copositive matrices that have entries in $\{-1,0,1\}$.

The separation problem now amounts to the following: Given some (fixed) copositive matrix $\tilde{H}$ of order k, find I of cardinality k such that $\langle X_I, \tilde{H}\rangle < 0$. Let H denote the $n \times n$ matrix with $\tilde{H}$ as principal minor and which is zero otherwise, $H = \begin{pmatrix} \tilde{H} & 0 \\ 0 & 0 \end{pmatrix}$. Then the optimal solution ϕ of the quadratic assignment problem (QAP)

$$q^* := \min_{\phi} \sum_{ij} h_{ij} x_{\phi(i),\phi(j)}$$

over all permutations ϕ shows the following. If $q^* < 0$, then $I = \{\phi(1), \ldots, \phi(k)\}$ has $\langle \tilde{H}, X_I\rangle = q^* < 0$, and we have found a violating subset. The idea of using models based on the QAP has already been used by Christof and Reinelt [4] to find violating local cuts. The interesting point here is that using copositive matrices to identify I avoids the need to study the facial structure of small-cut or stable-set polytopes.

In Table 1 we provide some computational results with the new bound for the Stable-Set problem. Subsets I of size $k = 5$ were identified with copositive matrices of order 5 and a QAP-based heuristic as described above. In all cases the new bound (with 100 subset constraints) provides a clear improvement over the theta number $\theta(G)$. This fact is particularly impressive for the cubic graph and the grid graph. Further computational results can be found in [1].

Table 1: Results for instances of Stable-Set problems of various sizes.

Graph	n	$\theta(G)$	New Bound	$\alpha(G)$
g60-25	60	15.0058	14.71	14
cubic	74	34.8561	33.34	≥ 32
g80-25	80	17.1670	17.01	17
g100-10	100	32.1166	31.52	≥ 29
spin5	125	55.9017	51.61	≥ 50

REFERENCES

[1] E. Adams, M.F. Anjos, F. Rendl, and A. Wiegele. A hierarchy of subgraph projection-based semidefinite relaxations for some NP-hard graph optimization problems. Tech. Report, Cahier G-2013-80, GERAD, 2013.

[2] V. J. D. Baston. Extreme copositive quadratic forms. *Acta Arith.*, 15:319–327, 1968/1969.

[3] C. Buchheim, F. Liers, and M. Oswald. Local cuts revisited. *Oper. Res. Lett.*, 36(4):430–433, 2008.

[4] T. Christof and G. Reinelt. Algorithmic aspects of using small instance relaxations in parallel branch-and-cut. *Algorithmica*, 30(4):597–629, 2001.

[5] E. Boros, Y. Crama and P. L. Hammer. Upper-bounds for quadratic 0-1 maximization. *Oper. Res. Lett.*, 9(2):73–79, 1990.

[6] W. Fernandez de la Vega and C. Kenyon-Mathieu. Linear programming relaxations for maxcut. In *Proc. SODA*, pages 53–61, 2007.

[7] C. Delorme and S. Poljak. Laplacian eigenvalues and the maximum cut problem. *Math. Program. A*, 62(3):557–574, 1993.

[8] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42(6):1115–1145, 1995.

[9] C. Helmberg and F. Rendl. Solving quadratic (0,1)-problems by semidefinite programs and cutting planes. *Math. Program. A*, 82(3):291–315, 1998.

[10] R. Hildebrand. The extreme rays of the 5×5 copositive cone. *Linear Algebra Appl.*, 437(7):1538–1547, 2012.

[11] A. J. Hoffman and F. Pereira. On copositive matrices with −1,0,1 entries. *J. Combin. Theory A*, 14:302–309, 1973.

[12] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM J. Optimiz.*, 1:166–190, 1991.

[13] H. D. Sherali and W. P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM J. Discrete Math.*, 3(3):411–430, 1990.

In Memoriam

Michael J. D. Powell (1936-2015)

Mike Powell in China in 2011 *(Photo credit: Lin Wang).*

Michael James David Powell passed away on the 19th of April 2015 after a short illness. Born on the 29th of July 1936 in Kensington (London), Mike was the eldest of three children, with a sister Anne and a brother Peter. After a childhood in Sussex and Surrey, Mike studied at the Eastbourne College public school, where his interest in mathematics – especially exercises – already put him ahead of his class. Subsequently, as an undergraduate at Peterhouse College, University of Cambridge, he first passed Part 2 of the Mathematical Tripos and then took a Diploma in Numerical Analysis and Computing. After university, he started working at the Atomic Energy Research Establishment at Harwell in 1959; he also married Caroline in September of the same year. During the 17 years at Harwell, he helped physicists working on the civil side of nuclear energy improve their computational methods. He also wrote numerical codes in FORTRAN, initiating the Harwell Subroutine Library, one of the first libraries of numerical algorithms and now used worldwide. His interests in numerical analysis were broad and it is not until 1962 that he published his first paper in optimization.

In this productive period at Harwell, one cannot but mention some notable highlights. The first is the well-known publication, together with Roger Fletcher, of the pivotal paper [3] presenting the Davidon-Fletcher-Powell (DFP) variable-metric method for unconstrained optimization. This technique was originally described by Bill Davidon, but its truly international status was boosted by Mike and Roger's work that provided clear numerical evidence of computational efficiency. The DFP method, which uses only function values and gradients, was such an improvement on other first-order techniques that it generated huge interest, both theoretical and practical, and spawned a complete subfield of optimization (*quasi-Newton methods*) with hundreds of papers in the subject over more than twenty years. The state-of-the-art BFGS algorithm is a prominent member of this vast family of methods. Mike was proud that NASA used his DFP method in the planning of the Apollo lunar missions. A second highlight is the publication in 1969 of the original idea for the *augmented Lagrangian* [5], which was the seed for many later contributions. A third highlight is the paper on the Powell-Symmetric-Broyden (PSB) quasi-Newton method [6], published in 1970, where Mike gave the first convergence proof for *trust-region methods*, a topic that still today is generating much research worldwide. A fourth contribution is an economical method to estimate sparse Jacobian matrices [1], whose brilliant idea was allegedly discussed with his colleagues and friends Alan Curtis and John Reid while queuing in the famous Harwell canteen. Once more, their simple and elegant technique led to several generations of papers devoted to its extensions and adaptations in other contexts.

In 1976, Mike was offered the John Humphrey Plummer Professorial Chair of Numerical Analysis at Cambridge University and moved to Milton Road with his wife and three children. Shortly after, he was elected fellow of Pembroke College. He started his teaching by giving an excellent course on approximation theory (the other mathematical love of his life besides optimization), which the second author had the privilege to attend in 1977 as his doctoral student. Life at the Department of Applied Mathematics and Theoretical Physics was active and provided the opportunity to meet most interesting people. The second author has, in particular, vivid memories of meeting Stephen Hawking and Sir James Lighthill, whose office was next to Mike's. At that time Mike wrote his first paper on *sequential quadratic programming* [7], a technique still in widespread use today for the solution of constrained optimization problems. He was elected to the Royal Society in 1983. Since it was not in Mike Powell's character to be distracted from research by university committees and other administrative duties, he kept producing top-level research in optimization up to (and well into) his retirement at the age of 65. In fact Mike never stopped doing research and saw retirement mostly as an opportunity for doing more of it! Recently, he returned with renewed focus to one of his earliest interests, *derivative-free optimization* [4], a class of minimization methods that use only function values. Motivated by its wide applicability in practice, Mike developed some of the best codes (BOBYQA, LINCOA) in this area and kept working on this topic until his last days. An initiative is under way to preserve Mike's codes for the future (see http://bit.ly/1yIaTea).

Mike made several other fundamental contributions to our understanding of optimization methods. In the realm of superbly elegant proofs, he showed convergence of the DFP method with exact linesearches on strictly convex functions [8] and of the BFGS method with inexact linesearches on convex functions [10]. Furthermore, Mike is well known not only as a constructor of proofs but also as a deconstructor, when he thought a proof could not be given. His keen counterexample-building skills have greatly improved our understanding – such as his worst-case analysis of Karmarkar's interior-point algorithm on the discretization of the unit circle [11], showing increasingly poor performance with an increase in the number of constraints. Similarly, his efforts sealed the fate of some methods – such as his example of failure of coordinate search methods to converge to stationary points [9]. An excellent survey of his major contributions is given in [2].

Mike received numerous accolades such as the Dantzig prize of the joint MOS-SIAM Societies and the IMA Gold Medal. In addition to the fellowship of the Royal Society, he was elected Foreign Member of the U.S. National Academy of Sciences (2001) and Corresponding Member of the Australian Academy of Science (2007).

Known for his acute understanding, Mike was notoriously difficult to convince, but, once convinced, he was your best advocate. He applied the highest standards to all aspects of his life, including research. A keen competitor, he had endless interests: walking up mountains, golf (which inspired his well-known "dog-leg" technique [6]), wine, food, traveling the canals of England and Wales on a narrowboat, hockey, bridge, Pembroke bowls, squash, road-rally racing (as the navigator), and many more.

To say that optimization has lost a founding father is an understatement. The entire optimization community recognizes his fundamental contributions and grieves over the loss of one of the giants whose shoulders help us all to see further toward the horizon.

Memories, comments, and farewells to Mike can be left on the website kindly set up by Dominique Orban at http://michaeljdpowell.blogspot.ca.

— *Coralia Cartis and Philippe Toint,*
Mike's last and first doctoral students at Cambridge

Coralia Cartis
Mathematical Institute, University of Oxford, UK,
coralia.cartis@maths.ox.ac.uk
Philippe L. Toint
Namur Center for Complex Systems (naXys) and Dept. Mathematics, University of Namur, Belgium,
philippe.toint@unamur.be

REFERENCES

[1] A. Curtis, M. J. D. Powell, and J. Reid. On the estimation of sparse Jacobian matrices. *J. Inst. Math. Appl.*, 13:117–119, 1974.

[2] R. Fletcher. The contributions of Mike Powell to optimization. In *Approximation Theory and Optimization: Tributes to M.J.D. Powell*, pp 13–19. Cambridge Univ. Press, 1997.

[3] R. Fletcher and M. J. D. Powell. A rapidly convergent descent method for minimization. *Comput. J.*, 6:163–168, 1963.

[4] M. J. D. Powell. An efficient method for finding the minimum of a function of several variables without calculating derivatives. *Comput. J.*, 17:155–162, 1964.

[5] M. J. D. Powell. A method for nonlinear constraints in minimization problems. In R. Fletcher, editor, *Optimization*, pp 283–298, London, 1969. Academic Press.

[6] M. J. D. Powell. A new algorithm for unconstrained optimization. In J. B. Rosen, O. L. Mangasarian, and K. Ritter, editors, *Nonlinear Programming*, pp 31–65, London, 1970. Academic Press.

[7] M. J. D. Powell. A fast algorithm for nonlinearly constrained optimization calculations. In G. A. Watson, editor, *Numerical Analysis, Dundee 1977*, no. 630 in Lecture Notes in Mathematics, pp 144–157, Berlin, 1978. Springer Verlag.

[8] M.J.D. Powell. On the convergence of the variable metric algorithm. *J. Inst. Math. Appl.*, 7:21–36, 1971.

[9] M.J.D. Powell. On search directions for minimization algorithms. *Math. Program.*, 4:193–201, 1973.

[10] M.J.D. Powell. Some global convergence properties of a variable metric algorithm for minimization without exact line searches. In *Nonlinear Programming SIAM-AMS Proceedings, Vol. IX*, pp 53–72. AMS, 1976.

[11] M.J.D. Powell. Karmarkar's algorithm: a view from nonlinear programming. *IMA Bull.*, 26:165–181, 1990.

Bulletin

Email items to siagoptnews@lists.mcs.anl.gov for consideration in the bulletin of forthcoming issues.

1 Event Announcements

1.1 ISMP 2015 in Pittsburgh

The 22nd International Symposium on Mathematical Programming (ISMP 2015) will take place in Pittsburgh, PA USA, July 12-17, 2015. ISMP is the world congress of mathematical optimization where scientists as well as industrial users of mathematical optimization meet in order to present the most recent developments and results and to discuss new challenges from theory and practice.

Plenary Speakers: Laurent El Ghaoui, Univ. California, Berkeley; Jim Geelen, Univ. Waterloo, Canada; Daniel Kuhn, EPFL, Switzerland; Dan Spielman, Yale Univ.; Steve Wright, Univ. Wisconsin

Semi-Plenary Speakers: Sam Burer, Univ. Iowa; Roberto Cominetti, Univ. Chile, Chile; Michele Conforti, Univ. Padova, Italy; Tammy Kolda, Sandia Labs; Andrea Lodi, Univ. Bologna, Italy; Asu Ozdaglar, MIT; Werner Roemisch, Humboldt Univ. Berlin, Germany; Frank Vallentin, Univ. Koeln, Germany; Pascal van Hentenryck, NICTA, Australia; Yaxian Yuan, Chinese Academy of Sciences, China

The symposium will take place at the Wyndham Grand Pittsburgh Downtown Hotel located at the confluence of Pittsburgh's famed three rivers. The conference registration is now open; more details are available on the conference website http://www.ismp2015.org.

2 Book Announcements

2.1 Variational Analysis in Sobolev and BV Spaces: Applications to PDEs and Optimization, Second Edition

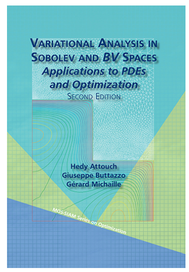

By Hedy Attouch, Giuseppe Buttazzo, and Gérard Michaille

Publisher: SIAM

Series: MOS-SIAM Series on Optimization, Vol. 17

ISBN: 978-1-611973-47-1, xii + 793 pages

Published: October 2014

http://bookstore.siam.org/mo17/

About the book: This volume is an excellent guide for anyone interested in variational analysis, optimization, and PDEs. It offers a detailed presentation of the most important tools in variational analysis as well as applications to problems in geometry, mechanics, elasticity, and computer vision. New elements in this second edition include: coverage of quasi-open sets and quasicontinuity; an increased number of examples in the areas of linearized elasticity system, obstacle problems, convection-diffusion, and semilinear equations; and a new subsection on stochastic homogenization.

2.2 Introduction to Nonlinear Optimization Theory, Algorithms, and Applications with MATLAB

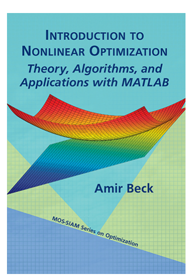

By Amir Beck

Publisher: SIAM

Series: MOS-SIAM Series on Optimization, Vol. 19

ISBN: 978-1-611973-64-8, x + 282 pages

Published: October 2014

http://bookstore.siam.org/mo19/

About the book: Readers will find more than 170 theoretical, algorithmic, and numerical exercises that deepen and enhance their understanding of nonlinear optimization. The author includes several subjects not typically found in optimization books—for example, optimality conditions in sparsity-constrained optimization, hidden convexity, and total least squares. The book also discusses a large number of applications theoretically and algorithmically.

2.3 Vehicle Routing: Problems, Methods, and Applications, Second Edition

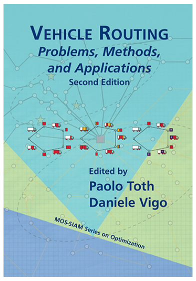

By Paolo Toth and Daniele Vigo

Publisher: SIAM

Series: MOS-SIAM Series on Optimization, Vol. 18

ISBN: 978-1-611973-58-7, xviii + 463 pages

Published: December 2014

http://bookstore.siam.org/mo18/

ABOUT THE BOOK: This second edition replaces The Vehicle Routing Problem (DT09, ISBN: 978-0-898715-79-8), which is no longer available.

Vehicle routing problems, among the most studied in combinatorial optimization, arise in many practical contexts (freight distribution and collection, transportation, garbage collection, newspaper delivery, etc.). Operations researchers have made significant developments in the algorithms for their solution, and Vehicle Routing: Problems, Methods, and Applications, Second Edition reflects these advances. The text of the new edition is either completely new or significantly revised and provides extensive and complete state-of-the-art coverage of vehicle routing by those who have done most of the innovative research in the area; emphasizes methodology related to specific classes of vehicle routing problems and, since vehicle routing is used as a benchmark for all new solution techniques, it contains a complete overview of current solutions to combinatorial optimization problems; includes several chapters on important and emerging applications, such as disaster relief and green vehicle routing.

2.4 Arc Routing: Problems, Methods, and Applications

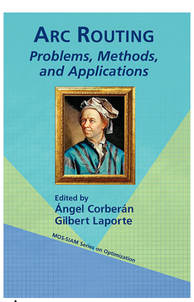

By Ángel Corberán and Gilbert Laporte

Publisher: SIAM

Series: MOS-SIAM Series on Optimization, Vol. 20

ISBN: 978-1-611973-66-2, xviii + 392 pages

Published: January 2015

http://bookstore.siam.org/mo20/

ABOUT THE BOOK: This book provides a thorough and up-to-date discussion of arc routing by world-renowned researchers. Organized by problem type, the book offers a rigorous treatment of complexity issues, models, algorithms, and applications. The book opens with a historical perspective of the field and is followed by three sections that cover complexity and the Chinese Postman and the Rural Postman problems; the Capacitated Arc Routing Problem and routing problems with min-max and profit maximization objectives; and important applications, including meter reading, snow removal, and waste collection.

2.5 Evaluating Gas Network Capacities

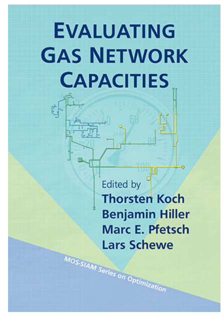

Edited by Thorsten Koch, Benjamin Hiller, Marc E. Pfetsch, and Lars Schewe

Publisher: SIAM

Series: MOS-SIAM Series on Optimization, Vol. 21

ISBN: 978-1-611973-68-6, xvi + 376 pages

Published: March 2015

http://bookstore.siam.org/mo21/

ABOUT THE BOOK: This book addresses a seemingly simple question: Can a certain amount of gas be transported through a pipeline network? The question is difficult, however, when asked in relation to a meshed nationwide gas transportation network and when taking into account the technical details and discrete decisions, as well as regulations, contracts, and varying demands, involved. This book provides an introduction to the field of gas transportation planning, discusses the advantages and disadvantages of several mathematical models that address gas transport within the context of its technical and regulatory framework, and shows how to solve the models by using sophisticated mathematical optimization algorithms. The book also includes examples of large-scale applications of mathematical optimization to this real-world industrial problem. Readers will also find a glossary of gas transport terms, tables listing the physical and technical quantities and constants used throughout the book, and a reference list of regulation and gas business literature.

3 Other Announcements

3.1 SIAG on Optimization Prize

At the 2014 SIAM Conference on Optimization, Stanford University's Yinyu Ye was awarded the SIAM Activity Group on Optimization Prize for his paper "The Simplex and Policy-Iteration Methods Are Strongly Polynomial for the Markov Decision Problem with a Fixed Discount Rate," *Mathematics of Operations Research*, 36(4):593-603, 2011.

This paper proved that the classic policy-iteration method (Howard 1960), including the simplex method (Dantzig 1947) with the most-negative-reduced-cost pivoting rule, is also a strongly polynomial-time exact algorithm for solving the Markov decision problem (MDP) exactly for any fixed discount factor. Furthermore, the computational complexity of the policy-iteration method (including the simplex method) is better than that of a suitably designed interior-point algorithm (that Ye also showed is strongly polynomial-time on MDP), which matches its superior practical performance. The result is surprising because the simplex method with the same pivoting rule was shown to be exponential-time in the worst case for solving a general linear programming problem (Klee and Minty, 1972), and the simplex method with the smallest index pivoting rule was shown to be exponential-time in the worst case for solving an MDP regardless of the discount rate. Subsequent to this work, Ian Post and Yinyu Ye proved that the simplex method is strongly polynomial-time for deterministic MDP regardless of the discount factor ("The simplex method is strongly polynomial for deterministic Markov decision processes," *SODA 2013*).

The simplex method remains one of most effective methods for solving linear programs, especially for Markov decision problems, and the method is arguably one of the most widely used decision models/methodologies in practice. Prior to Ye's result, however, the simplex method was known to be strongly polynomial only for linear programs with totally unimodular matrices (such as those arising from networks). Ye's result represents a first step that goes beyond the class of totally unimodular linear programs. As such, it provides a valuable contribution toward settling the question of strongly polynomial time solvability of all linear programs. Moreover, it succeeds in showing the strongly polynomial property for a classical and widely used pivoting strategy as opposed to an artificial construct built for theoretical purposes only. In addition, the mathematical proofs are elegant and readable by a wide optimization audience.

The SIAM Activity Group on Optimization (SIAG/OPT) Prize, established in 1992, is awarded to the author(s) of the most outstanding paper, as determined by the prize committee, on a topic in optimization published in English in a peer-reviewed journal. This year's prize committee consisted of Mihai Anitescu (chair, Argonne National Lab.), Robert Freund (MIT), Dorit Hochbaum (Univ. California Berkeley), Tom Luo (Univ. Minnesota), and Luís Nunes Vicente (Univ. Coimbra).

3.2 2015 SIAM Fellows Announced

Each year, SIAM designates as Fellows of the society those who have made outstanding contributions to the fields of applied mathematics and computational science. This year, 31 members of the community were selected for this distinction.

These new Fellows include five members of the SIAG, whose citations are included below. Full details on the SIAM Fellow program can be found at http://www.siam.org/prizes/fellows/index.php. Congratulations to all the new Fellows!

Aharon Ben-Tal
Technion - Israel Institute of Technology
For contributions to continuous optimization, both theory and applications, including the field of robust optimization.

Stephen P. Boyd
Stanford University
For fundamental contributions to the development, teaching, and practice of optimization in engineering.

William W. Hager
University of Florida
For contributions to optimal control, optimization theory, and numerical optimization algorithms.

Tamara G. Kolda
Sandia National Laboratories
For contributions to numerical algorithms and software in multi-linear algebra, optimization, and graph analysis.

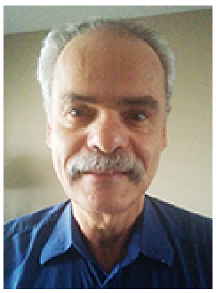

Henry Wolkowicz
University of Waterloo
For contributions to convex optimization and matrix theory.

3.3 2015 John von Neumann Theory Prize – Call for Nominations

The John von Neumann Theory Prize is awarded annually to a scholar who has made fundamental, sustained contributions to theory in operations research and the management sciences. The award is given each year at the INFORMS Annual Meeting if there is a suitable recipient. Although the prize is normally given to a single individual, in the case of accumulated joint work, the recipients can be multiple individuals.

The prize is awarded for a body of work, typically published over a period of several years. Although recent work should not be excluded, the prize typically reflects contributions that have stood the test of time. The criteria for the prize are broad and include significance, innovation, depth, and scientific excellence.

The award is $5,000, a medallion, and a citation.

Application Process: The Prize Committee is currently seeking nominations, which should be in the form of a letter (preferably email) addressed to the prize committee chair (below), highlighting the nominee's accomplishments. Although the letter need not contain a detailed account of the nominee's research, it should document the overall nature of the nominee's contributions and their impact on the profession, with particular emphasis on the prize's criteria. The nominee's curriculum vitae, while not mandatory, would be helpful. Please compress electronic files if larger than 10 MB. Please see http://goo.gl/uWT0pn for complete details.

Nominations should be submitted to the committee chair as soon as possible, but no later than June 1, 2015: George Nemhauser, Professor, Georgia Institute of Technology, Atlanta, GA 30332, USA, george.nemhauser@isye.gatech.edu.

Chair's Column

Welcome to the latest edition of the SIAG Optimization *Views and News*. I am deeply honored to have been elected chair for this SIAG, and I'm looking forward to the next three years. In addition to myself, the three new officers are Martine Labbe (vice chair), Michael Friedlander (program director), and Kim-Chuan Toh (secretary/treasurer).

I would like to start off by recognizing the great work of the previous officers. The team of Michael Todd, Mihai Anitescu, Miguel Anjos, and Marina Epelman have done a fantastic job of guiding the SIAG these past three years, as was evident by last year's outstanding conference.

You may have also noticed that we have two new editors for the *Views and News*. Stefan Wild (Argonne National Laboratory) and Jennifer Erway (Wake Forest) have kindly agreed to take on this responsibility. Our thanks go out to the former editors, Sven Leyffer and Franz Rendl, both of whom have done an outstanding job with the *Views and News* for the past six years.

As you will see, this edition is devoted to a recap of our highly successful Optimization Conference in San Diego this past May. If you were there, you know that it was one of the best conferences we've ever had with an attendance of 677, making it the largest conference in our history. It was exciting to see old friends again and meet new ones as well.

I would also like to recognize Yinyu Ye from Stanford, who was this year's recipient of the SIAG Optimization Prize. His lecture on the Efficiency of the Simplex and Policy Iteration Methods for Markov Decision Processes, which you can find through our web site (http://www.siam.org/activity/optimization/) is worth a look. In addition, 2014 saw six of our members become SIAM Fellows: Omar Ghattas, Philip Gill, Dorit Hochbaum, Masakazu Kojima, Jean Lasserre, and Christine Shoemaker. Recently, the class of 2015 Fellows were announced; see the announcement in this issue. Congratulations to all of you!

Our membership numbers are healthy, with over 1,000 total members, although the number of student members has declined over the past two years. I hope all our members will encourage students at their institutions to sign up for the Optimization SIAG. As a reminder, students at many of our universities receive free membership, and it's a great way for them to get to know SIAM and to start their professional networks.

You may be wondering where the next optimization conference will be. We had a short discussion on this topic during the business meeting in San Diego, and several ideas were presented, but as of now we have no formal proposals. If you have any ideas or would like to submit a formal proposal to us, please feel free to contact any one of the officers directly. For that matter, if you have any ideas for activities that the SIAG should be involved in, please let us know.

Juan Meza, SIAG/OPT Chair
School of Natural Sciences, University of California, Merced, Merced, CA 95343, USA, jcmeza@ucmerced.edu, http://bit.ly/1G8HUxO

Comments from the Editors

To Print or Not to Print?

We are excited to take over the *Views and News* reins from the great precedence set by Sven Leyffer (who served 2008–2014) and Franz Rendl (2011–2014). Thanks to Miguel Anjos and Mike Todd's look back on the 2014 Conference on Optimization, it appeared as if we were on track for an easy first issue...

...and then it came time to print.

It turns out that *Views and News* is currently the only SIAG newsletter that is physically printed and mailed to members. Over the past two years, SIAM has been encouraging SIAGs to make greater use of SIAM-hosted wikis and websites and to produce electronic-only newsletters.

Indeed, *Views and News* is archived in electronic (pdf) form back to the inaugural 1992 issue via SIAG/OPT's wiki at http://wiki.siam.org/siag-op.

As reported at last year's business meeting, the semi-annual expenses associated with *Views and News* printing and mailing do not place a significant financial burden on the SIAG. Furthermore, SIAM maintains a list of SIAG/OPT members who have opted out of physical mailings (historically, the opt-out list has consisted of less than 1% of members).

However, we would like to use our first (and maybe last?!) issue to survey SIAG/OPT's members:

- Should physical copies of *Views and News* no longer be printed and mailed to members?
- Should there be an "opt-in" list for physical copies (i.e., membership default will be "opt-out")?
- What format/frequency/alert mechanism should electronic issues of *Views and News* take? Would you contribute to a SIAG/OPT blog?

What do other optimization societies do? The Mathematical Optimization Society has a quarterly newsletter, *Optima*, which is currently printed and mailed to all MOS members.

The INFORMS Optimization Society has an annual newsletter, *INFORMS OS Today*, which is currently emailed to all OS members.

What do other SIAGs do? Seven SIAGs have maintained an electronic newsletter in recent years.[1] Most other SIAGs have a web presence (edited annually or sporadically) but do not maintain a newsletter.[2]

Optimization has the third largest SIAG membership (just behind Dynamical Systems), and the largest SIAG (Computational Science & Engineering) arguably spans a greater part of industrial and applied mathematics than does SIAG/OPT. Are there additional factors that support the printing of our SIAG's newsletter?

Send us your feedback! Jennifer and I are among those who "grew up" (academically) flipping through *Views and News* and *Optima* back issues left in student lounges, and I am more likely to encounter a copy of these newsletters than a hardcopy journal at an optimization coffee break.[3]

I am sure Mike Powell would have loved to send feedback on the subject. It was not only an honor to arrange for an overhead projector at optimization meetings for his foils (transparencies), it was also kind of fun, and definitely memorable.

We welcome your feedback, (e-)mailed directly to us or to siagoptnews@lists.mcs.anl.gov. Suggestions for new issues, comments, and papers are always welcome! See you in six months (in your mailbox and/or inbox),

Stefan Wild, Editor
Mathematics and Computer Science Division, Argonne National Laboratory, USA, wild@anl.gov, http://www.mcs.anl.gov/~wild
Jennifer Erway, Editor
Department of Mathematics, Wake Forest University, USA, erwayjb@wfu.edu, http://www.wfu.edu/~erwayjb

[1] Analysis of Partial Differential Equations (last issue 02/2011) http://siags.siam.org/siagapde, Discrete Mathematics (last issue 02/2012) http://bit.ly/1bjLEUQ, Dynamical Systems (quarterly) http://www.dynamicalsystems.org, Financial Mathematics and Financial Engineering (annually) http://wiki.siam.org/siag-fm, Imaging Science (last issue 06/2010) http://siags.siam.org/siagis, Linear Algebra (blog-like) http://siags.siam.org/siagla, and Orthogonal Polynomials and Special Functions (bimonthly) http://math.nist.gov/opsf.

[2] These include Algebraic Geometry, http://wiki.siam.org/siag-ag; Computational Science and Engineering: http://wiki.siam.org/siag-cse; Geometric Design, http://siags.siam.org/siaggd; Geosciences, http://wiki.siam.org/siag-gs; and Mathematical Aspects of Materials Science, http://wiki.siam.org/siag-ms.

[3] Of course, smartphones are even more prevalent; I'm sure a reader will use one to document a journal-dominated coffee break.