Operations Research Letters 42 (2014) 404–408

Contents lists available at ScienceDirect

Operations Research Letters

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

Solving the continuous nonlinear resource allocation problem with an interior point method

Stephen E. Wright [a,*], James J. Rohal [b]

[a] Department of Statistics, Miami University, Oxford, OH 45056, United States

[b] Department of Mathematics, North Carolina State University, Raleigh, NC 27695, United States

ARTICLE INFO

Article history:
Received 23 September 2013
Received in revised form
25 March 2014
Accepted 2 July 2014
Available online 10 July 2014

Keywords:
Convex programming
Interior point methods
Continuous knapsack

ABSTRACT

Resource allocation problems are usually solved with specialized methods exploiting their general sparsity and problem-specific algebraic structure. We show that the sparsity structure alone yields a closed-form Newton search direction for the generic primal-dual interior point method. Computational tests show that the interior point method consistently outperforms the best specialized methods when no additional algebraic structure is available.

1. Introduction

We consider the resource allocation problem in the form

$$\text{minimize } f(x) := \sum_{i=1}^{n} f_i(x_i) \text{ over all } x \tag{1}$$

$$\text{subject to } g(x) := \sum_{i=1}^{n} g_i(x_i) \le b, \tag{2}$$

$$l \le x \le u. \tag{3}$$

Here x, l, and u are n-vectors of real numbers, b is a real scalar, and the functions f_i and g_i are convex and twice differentiable on an open set containing the interval $[l_i, u_i]$. Inequalities of vectors are interpreted coordinate-wise. We shall also consider the related problem in which (2) is replaced with a linear equality constraint.

The recent survey paper of Patriksson [8] shows that such problems have a long history and diverse applications. The contexts in which the problem appears often demand that it be solved very quickly, even in high dimensions. Consequently, researchers long ago moved beyond general-purpose nonlinear programming procedures and focused on exploiting the special structure of the optimality conditions for the problem. As noted by Patriksson, two frameworks have emerged as the most competitive for solving resource allocation problems: the *pegging* or *variable-fixing* methods and the *breakpoint-search* methods. Patriksson also observes that computational studies in the literature have generally indicated that pegging is superior to breakpoint search when certain subproblems (see Section 2) common to both methods are easily solved, whereas breakpoint search is faster otherwise. Moreover, numerical comparisons of either method with general-purpose solvers are essentially nonexistent in the literature.

Here we present evidence that a primal-dual interior point method outperforms breakpoint search on problems for which the latter is traditionally considered the best possible choice, namely, when its subproblems do not admit closed-form solutions and must be solved numerically. We show that the special structure of (1)–(3) allows for a closed-form solution of the linear system defining the search directions and we present computational results showing the method's superiority. This addresses two questions posed by Patriksson [8]. First, it shows that the sparsity can be exploited within the setting of a general-purpose optimizer. Second, it provides an efficient method that also avoids the usual assumptions (see Section 2) imposed by pegging or breakpoint search methods on the domain, monotonicity or strict convexity of f_i and g_i.

In the next section, we review the optimality conditions for (1)–(3). In Section 3 we describe the breakpoint search and interior point methods, along with details of their implementation. Section 4 lays out the problem instances used for the computational tests, and the results are discussed in Section 5.

* Corresponding author.
E-mail addresses: wrightse@miamioh.edu, wrightse@muohio.edu (S.E. Wright), jjrohal@ncsu.edu (J.J. Rohal).

http://dx.doi.org/10.1016/j.orl.2014.07.001
0167-6377/

2. Optimality conditions

In this study we make the following assumptions:

(A1) The problem (1)–(3) has no optimum with $g(x) < b$.
(A2) The function f_i is decreasing on $[l_i, u_i]$ and g_i is increasing on $[l_i, u_i]$ with $g(l) < b < g(u)$.

The randomly generated test instances of Section 4 all satisfy these assumptions, which are needed for breakpoint search but not for the interior point method. In practice, we can easily determine whether either assumption holds if we know the intervals of monotonicity for each f_i and g_i. Indeed, many treatments of resource allocation problems include one or both of these assumptions because they can be inexpensively enforced through some combination of initialization, preprocessing, and data generation.

Assumptions A1–A2 imply that the problem (1)–(3) admits an optimal solution, that the Slater constraint qualification holds, and that problem (1)–(3) is equivalent to the problem in which the inequality constraint (2) is forced to hold as equality. For similar reasons, our discussion also covers the related problem where the inequality (2) is replaced by a linear equality (in which case we still assume, for simplicity, that the preprocessing has ensured that A1–A2 hold).

By Lagrangian duality, necessary and sufficient optimality conditions can be expressed as follows: $g(x) = b$ and, for some real number ρ, x is a solution to the separable optimization subproblem

$$\text{minimize } f(x) + \rho g(x) \text{ subject to } l \le x \le u. \tag{4}$$

The dual objective is

$$\rho \mapsto -b\rho + \sum_{i=1}^{n} \min_{x_i \in [l_i, u_i]} [f_i(x_i) + \rho g_i(x_i)], \tag{5}$$

which attains its maximum; moreover, any maximizer ρ is necessarily nonnegative. The subproblem (4) has coordinate-wise optimality conditions given by

$$\begin{aligned} f_i'(x_i) + \rho g_i'(x_i) &= 0, \quad \text{if } l_i < x_i < u_i, \\ f_i'(x_i) + \rho g_i'(x_i) &\ge 0, \quad \text{if } x_i = l_i, \\ f_i'(x_i) + \rho g_i'(x_i) &\le 0, \quad \text{if } x_i = u_i. \end{aligned}$$

The left-hand sides give the Karush–Kuhn–Tucker multipliers for the bounds $l_i \le x_i$ and $x_i \le u_i$, respectively, as

$$\begin{aligned} \lambda_i &:= \max\{0, -[f_i'(x_i) + \rho g_i'(x_i)]\}, \\ \mu_i &:= \max\{0, f_i'(x_i) + \rho g_i'(x_i)\}. \end{aligned}$$

Letting $s := u - x$ denote the vector of slack variables for the upper bounds on x, we express the Karush–Kuhn–Tucker (KKT) conditions for (1)–(3) as

$$\nabla f(x) + \rho \nabla g(x) - \lambda + \mu = 0, \tag{6}$$
$$x + s = u, \tag{7}$$
$$x \ge l, \quad \lambda \ge 0, \quad s \ge 0, \quad \mu \ge 0, \tag{8}$$
$$\operatorname{diag}(x - l)\lambda = 0, \qquad \operatorname{diag}(s)\mu = 0, \tag{9}$$
$$g(x) = b. \tag{10}$$

Here $\operatorname{diag}(z)$ denotes the diagonal matrix whose diagonal entries are the entries of the vector z.

The three solution frameworks discussed in Section 1 utilize the optimality conditions in different ways:

- Pegging methods solve subproblems of the form (1)–(2), but for which some variables are held fixed while the bounds (3) for all remaining variables are omitted.
- Breakpoint search methods maximize the dual objective (5) by solving a sequence of subproblems of the form (4) at various values of ρ.
- Primal-dual interior point methods apply Newton's method to perturbations of the KKT system (6)–(10).

The pegging and breakpoint search methods both benefit considerably when minimization of $x_i \mapsto f_i(x_i) + \rho g_i(x_i)$ can be handled efficiently. Because we focus on problems for which breakpoint search dominates pegging, we do not include pegging methods in this study. In fact, the pegging approach is not even well-defined for some of the problems we consider, because the pegging subproblems do not admit optimal solutions.

3. Methods and implementation

In this section, we describe the two main approaches considered in our computational study.

3.1. Breakpoint search

Breakpoint search is based on the observation that the dual objective (5) is concave and defined piecewise with a finite number of easily calculated breakpoints. The derivative, or subdifferential, of this objective is nonincreasing. A binary search of the breakpoints therefore identifies either one that is a root or a pair that most closely brackets a root.

There are at most $2n$ breakpoints, occurring at ρ-values where some $x_i \mapsto f_i(x_i) + \rho g_i(x_i)$ attains its minimum over $[l_i, u_i]$ at an endpoint l_i or u_i. Equivalently, a breakpoint makes the derivative $x_i \mapsto f_i'(x_i) + \rho g_i'(x_i)$ nonnegative at l_i or nonpositive at u_i. Consequently, all breakpoints have the form $\rho_i^+ := -f_i'(l_i)/g_i'(l_i)$ or $\rho_i^- := -f_i'(u_i)/g_i'(u_i)$. The monotonicity of f_i and g_i allow us to define $\rho_i^+ = \infty$ when $g_i'(l_i) = 0$ and to guarantee that $g_i'(u_i) > 0$ in the definition of ρ_i^-. The convexity and monotonicity of f_i and g_i also guarantee that $0 \le \rho_i^- \le \rho_i^+$.

The binary search sequentially refines a bracketing $\rho^- < \rho^* < \rho^+$ until the true root ρ^* lies between two consecutive breakpoints. The bracket is adjusted inward by finding a breakpoint ρ within it and testing the sign of the derivative of the dual objective (5). To evaluate that derivative at ρ, we first fix

$$x_i := \begin{cases} l_i, & \text{if } \rho \ge \rho_i^+, \\ u_i, & \text{if } \rho \le \rho_i^-. \end{cases} \tag{11}$$

The remaining minimizers are critical points: $f_i'(x_i) + \rho g_i'(x_i) = 0$ and $l_i < x_i < u_i$. Depending on the problem data, these critical points might be found (a) in closed form, (b) by using a problem-specific implementation of Newton's method, or (c) by means of a general-purpose Newton's method with Armijo linesearch for sufficient decrease and damping (as needed) to maintain $l_i < x_i < u_i$. The derivative value at ρ is then given by $-b + \sum_i g_i(x_i)$, the sign of which determines whether ρ becomes the new ρ^- or ρ^+. This in turn determines, through (11), that some values of x_i shall remain fixed and can therefore be removed from further consideration.

The final bracket, if nontrivial, consists of two closest breakpoints with the optimal value of ρ lying somewhere between them. To interpolate between them, our implementation finds ρ and the unfixed x_i-coordinates (denoted by $i \in I$) simultaneously by applying a multi-dimensional Newton's method with Armijo linesearch to the corresponding Lagrange multiplier conditions $\sum_{i \in I} g_i(x_i) = \hat{b}$ and $f_i'(x_i) + \rho g_i'(x_i) = 0$ for $i \in I$.

Throughout the procedure, the subproblem optimizations are initialized using the corresponding solutions from prior iterations. Also, we extract the required median values without sorting the list of breakpoints in advance, which can yield significant computational savings if each subproblem solution requires only a few operations per index i [1,4,7,9].