Journal of Computational Physics 258 (2014) 489–508

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Adaptive change of basis in entropy-based moment closures for linear kinetic equations

Graham W. Alldredge [a,*,1,2], Cory D. Hauck [b,3], Dianne P. O'Leary [c,2], André L. Tits [a,2]

[a] Department of Electrical and Computer Engineering & Institute for Systems Research, University of Maryland, College Park, MD 20742, USA
[b] Computational Mathematics Group, Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA
[c] Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, USA

ARTICLE INFO

Article history:
Received 13 June 2013
Accepted 26 October 2013
Available online 4 November 2013

Keywords:
Convex optimization
Realizability
Kinetic theory
Transport
Entropy-based closures
Moment equations

ABSTRACT

Entropy-based (M_N) moment closures for kinetic equations are defined by a constrained optimization problem that must be solved at every point in a space–time mesh, making it important to solve these optimization problems accurately and efficiently. We present a complete and practical numerical algorithm for solving the dual problem in one-dimensional, slab geometries. The closure is only well-defined on the set of moments that are realizable from a positive underlying distribution, and as the boundary of the realizable set is approached, the dual problem becomes increasingly difficult to solve due to ill-conditioning of the Hessian matrix. To improve the condition number of the Hessian, we advocate the use of a change of polynomial basis, defined using a Cholesky factorization of the Hessian, that permits solution of problems nearer to the boundary of the realizable set. We also advocate a fixed quadrature scheme, rather than adaptive quadrature, since the latter introduces unnecessary expense and changes the computationally realizable set as the quadrature changes. For very ill-conditioned problems, we use regularization to make the optimization algorithm robust. We design a manufactured solution and demonstrate that the adaptive-basis optimization algorithm reduces the need for regularization. This is important since we also show that regularization slows, and even stalls, convergence of the numerical simulation when refining the space–time mesh. We also simulate two well-known benchmark problems. There we find that our adaptive-basis, fixed-quadrature algorithm uses less regularization than alternatives, although differences in the resulting numerical simulations are more sensitive to the regularization strategy than to the choice of basis.

1. Introduction

Moment methods are commonly used to derive reduced models of kinetic transport. Rather than fully resolve the kinetic distribution in phase space, moment models instead track the evolution of a finite number of weighted velocity averages, or

* Corresponding author.
E-mail addresses: alldredge@mathcces.rwth-aachen.de (G.W. Alldredge), hauckc@ornl.gov (C.D. Hauck), oleary@cs.umd.edu (D.P. O'Leary), andre@umd.edu (A.L. Tits).
[1] Present address: Center for Computational Engineering Science, RWTH-Aachen University, Aachen, Germany.
[2] Supported by the U.S. Department of Energy, under Grant DESC0001862.
[3] This author's research was sponsored by the Office of Advanced Scientific Computing Research and performed at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. De-AC05-00OR22725.

0021-9991/$ – see front matter
http://dx.doi.org/10.1016/j.jcp.2013.10.049

moments of the distribution. Exact equations for these moments inevitably require missing information about the unknown kinetic distribution that must be approximated via a closure. Entropy-based closures approximate the full kinetic distribution by an ansatz that solves a constrained, convex optimization problem. In the context of radiative transport [1,2], these models are commonly referred to as M_N (after G.N. Minerbo), where N is the order of the highest-order moments of the model; see [3] for additional references. These moment models preserve many fundamental properties of the kinetic description, including positivity, entropy dissipation, and hyperbolicity [4].

The primary drawback of the entropy-based approach is computational cost: at least one optimization problem must be solved at every point on a space–time grid and, except for some M_1 cases, the optimization must be done numerically. For some moments, the associated optimization problem can be particularly expensive to solve. These moments lie near the boundary of the *realizable set*, defined as the set of vectors that are moments of a positive distribution. For realizable moments near the boundary, the optimization algorithm may require a large number of iterations to converge (or may not converge at all) and the solution will be sensitive to small changes in the moments. (Indeed, in some contexts, there even are realizable moments for which the optimization problem has no solution [5–8].)

It is most common to solve each optimization problem via the associated convex dual. For a smooth entropy function, a standard Newton method was proposed in [9]. In this context, the difficulty in solving optimization problems for moments near the realizable boundary is characterized by an ill-conditioned Hessian for the dual objective function. This matrix is a weighted integral in momentum space of a distribution of rank-one matrices. It becomes rank deficient (or nearly so) because the weight function is an approximation of the underlying kinetic density, and for moments near the realizable boundary, its mass will be concentrated around a small number of directions in momentum space. This is a common occurrence in radiation applications. The sensitivity is further exacerbated by the necessity of using an inexact quadrature and finite-precision arithmetic to approximate the integrals. The contributions to the quadrature may effectively be zero for most of the quadrature points, causing the computed Hessian to be singular.

The singularity in the dual Hessian has been addressed in different ways, in particular in the context of some variations of Newton's method. In [10], the author takes advantage of the structure of the Hessian, whose entries are themselves moments of a known distribution that changes at each iteration. Using orthogonal polynomials with respect to this distribution (which are found using a standard three-point recursion relation [11, Ch. 22]), the author is able to invert the ill-conditioned Hessian in a stable way. (See [12] for an efficient algorithm to evaluate the recursion coefficients.) More recently, orthogonal polynomials were used in the multi-dimensional implementation found in [13,14]. There the author applies a BFGS quasi-Newton method (see, e.g., [15, Ch. 6]) and, when the approximate Hessian becomes ill-conditioned, a Gram–Schmidt procedure is applied to change into a polynomial basis for which the Hessian is the identity.[4] In [16], a penalized version of the primal problem was introduced in order to handle non-realizable moments. This modification amounts to Tikhonov regularization of the dual problem, which also reduces ill-conditioning of the Hessian for realizable moments that are near the realizable boundary. In [17], two modifications to the Newton method for the solution of the dual problem were introduced. The first, which is only practical for a relatively small number of moments, is to generate an initial guess for the Newton solver by interpolating values from look-up tables. The second, which is only practical in one-dimension, is a root finding algorithm to guide the placement of nodes in the adaptive quadrature used to evaluate the objective function and its derivatives. Finally, in [3], where a damped Newton method is used, ill-conditioning of the Hessian near the realizable boundary is avoided in two ways. First, adaptive quadrature is used to better capture the support of the Hessian weight function and thereby increase the number of significant rank-one contributions. Second, a regularization method is introduced to move the moments away from the realizable boundary. These two remedies are automatically invoked as needed, since manual intervention is impractical.

Several specialized techniques for solving the dual problem have also been suggested. In [18,19], the authors show that, in the case of the Maxwell–Boltzmann entropy, the solution of the dual problem is the unique solution of a certain finite set of linear equations. Unfortunately, setting up this linear system requires the knowledge of additional moments that are not available to the closure. In [20], the authors solve the dual problem by means of a coordinate descent method, also known as Bregman's balancing method [21,22], where each sub-problem is solved using a multiple algebraic reconstruction technique [23,21]. Coordinate descent, however, is known to converge rather slowly (see, e.g., [15, p. 230]).

In this paper, along the lines of [3], we employ a damped Newton method and investigate ways to better handle hard optimization problems near the realizable boundary. We note that adaptive quadrature complicates realizability, so we do not use it. Further, we show that regularization affects accuracy and therefore should be avoided whenever possible. For this, we adopt the change-of-basis procedure introduced in [13,14], albeit with a different implementation. We focus on the Maxwell–Boltzmann entropy in the one-dimensional setting on a bounded domain, although our methodology is applicable to general smooth entropies and moments defined over bounded domains of arbitrary dimension.

As in [13,14], we observe that the change of basis makes the optimization more stable and effectively removes the need for an adaptive quadrature. Regularization is still required for robustness but is invoked far less frequently, resulting in noticeable improvements in accuracy in manufactured solution simulations we performed. We perform a series of numerical tests to quantify the overall performance of the algorithm, to assess the interplay between the change of basis and the

[4] In contrast to our problem, ill-conditioning in [13,14] occurs because the moments are integrals over an unbounded domain, so the Hessian is dominated by moments corresponding to the highest-order polynomials [5].