Journal of Computational Physics 302 (2015) 573–590

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

# Optimization and large scale computation of an entropy-based moment closure ☆

C. Kristopher Garrett [a], Cory Hauck [a,b,*], Judith Hill [c]

[a] Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN, United States
[b] Department of Mathematics, University of Tennessee, Knoxville, TN, United States
[c] Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory, Oak Ridge, TN, United States

ARTICLE INFO

*Article history:*
Received 14 August 2014
Received in revised form 13 July 2015
Accepted 5 September 2015
Available online 10 September 2015

*Keywords:*
Kinetic equations
Moment methods
GPU computing
Spherical harmonics
High performance computing

ABSTRACT

We present computational advances and results in the implementation of an entropy-based moment closure, $M_N$, in the context of linear kinetic equations, with an emphasis on heterogeneous and large-scale computing platforms. Entropy-based closures are known in several cases to yield more accurate results than closures based on standard spectral approximations, such as $P_N$, but the computational cost is generally much higher and often prohibitive. Several optimizations are introduced to improve the performance of entropy-based algorithms over previous implementations. These optimizations include the use of GPU acceleration and the exploitation of the mathematical properties of spherical harmonics, which are used as test functions in the moment formulation. To test the emerging high-performance computing paradigm of communication bound simulations, we present timing results at the largest computational scales currently available. These results show, in particular, load balancing issues in scaling the $M_N$ algorithm that do not appear for the $P_N$ algorithm. We also observe that in weak scaling tests, the ratio in time to solution of $M_N$ to $P_N$ decreases.

© 2015 Elsevier Inc. All rights reserved.

## 1. Introduction

Kinetic equations, such as the Boltzmann equation and the radiation transport equation, are integro-differential equations with up to seven independent variables: three space, three momentum, and time. Moment methods track the evolution of only a finite number of weighted momentum averages, or *moments*, of the kinetic distribution, thus reducing the dimensionality of the problem. However, this reduction requires a closure that approximates in some way the kinetic information that is lost in the averaging process. Thus, various methods will differ by the closure used in their formulation.

In the context of radiation transport, the classical moment method is the spherical harmonic expansion, colloquially termed $P_N$ in the radiation transport community [11,32,40]. This method uses a simple truncation closure that results in a linear hyperbolic balance law. However, the method may suffer from numerical artifacts, most notably large oscillations that can result in negative particle concentrations, especially in the streaming particle regime, where collisions are rare [7]. An alternative to the $P_N$ closure is a more complicated nonlinear closure based on minimizing a physically relevant, convex

☆ The submitted manuscript has been authored by a contractor of the U.S. Government under Contract No. DE-AC05-00OR22725. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

* Corresponding author.
*E-mail addresses:* garrettck@ornl.gov (C. Kristopher Garrett), hauckc@ornl.gov (C. Hauck), hilljc@ornl.gov (J. Hill).

http://dx.doi.org/10.1016/j.jcp.2015.09.008
0021-9991/© 2015 Elsevier Inc. All rights reserved.

function related to the entropy of the physical system [13,20,31]. The resulting method, colloquially termed $M_N$ in the radiation transport community, yields a nonlinear hyperbolic balance law that, unlike the $P_N$ method, formally captures the correct streaming limit [13], but at the same time, is accurate in scattering dominated regimes [9,17]. The $M_N$ method has been shown to be more accurate in several test cases [8,14,16,20] and in several different applications [6,15,18,33,36,45]. However, it requires the solution of a minimization problem at every spatial cell of the discretized domain. In general, the minimization must be solved numerically. This causes the $M_N$ method to require many more floating point operations than the $P_N$ method, even though the methods have the same data communication patterns. To avoid the computational overhead of the minimization problem, some approaches generate approximations of the entropy-based closure using look-up tables or interpolation schemes [34,35,47]. Such approximations are generally less robust than the original entropy-based closures and have, so far, been limited to low-order moment systems. Even so, in some cases, they maintain enough of the structure from the entropy-based approach to be considered as a suitable alternative.

The computational expense of the $M_N$ method makes it prohibitive for serial or even small-scale parallel implementations. However, for large scale computations on high performance machines, it is expected that the computing time for the $P_N$ method will eventually become dominated by communication, and in such cases, the $M_N$ method will be more competitive in time to solution. Therefore progress in this area depends on three factors: (i) algorithmic improvements in solving the minimization problem that defines the $M_N$ closure, (ii) performance improvements that leverage the available computer hardware, and (iii) scaling to extremely large problems. The work here builds on algorithmic improvements in [2,3].[1] In the current paper, we address the other two factors. First we design and test several optimizations for the $M_N$ algorithm that reduce the time to solution by as much as 10 times in some cases. We then explore scalability of the $M_N$ algorithm using an explicit time integration algorithm. Using the supercomputer Titan, which is housed at Oak Ridge National Laboratory and operated by the Oak Ridge Leadership Computing Facility, we find the $M_N$ algorithm weakly scales almost perfectly out to 17,576 compute nodes while the $P_N$ algorithm displays an increase in time per node by a factor of $1.2\times$ to $4\times$, depending on the amount of data per node. However, even with performance improvements, the time to solution of the $M_N$ algorithm is still approximately 25 times greater than the time to solution of the $P_N$ algorithm when both are run at full scale.

The layout of the paper is as follows. In Section 2, we briefly summarize the moment approach, discuss important implementation details, and introduce two test problems that will be used for numerical simulations. In Section 3, we introduce three improvements to the $M_N$ algorithm: one that leverages structure in the Hessian matrix of the $M_N$ minimization algorithm and two that use GPUs to accelerate the two most arithmetically intensive parts of the computation. In Section 4, $M_N$ statistics and timing results are presented for the two test problems. Results of weak scaling tests for $P_N$ and $M_N$ are also compared. Section 5 is for conclusions and discussion. The Appendix contains useful technical details about spherical harmonics and Gaunt coefficients. Appendix D contains a glossary for the variables used throughout the paper.

## 2. Moment equations

In this section, we briefly summarize the necessary background material on moment methods, give details on numerical implementation, and present two initial conditions used in the numerical examples.

### 2.1. Formulation

The governing equation for this study is a linear kinetic transport equation for unit speed particles in an infinite medium. This equation takes the form

$$\partial_t f + \Omega \cdot \nabla_x f = \frac{1}{4\pi}\sigma_s \langle f \rangle - \sigma_t f, \tag{1}$$

where (i) $x \in \mathbb{R}^3$ is a point in space, (ii) $\Omega \in \mathbb{S}^2$ (the unit sphere) is a velocity direction (velocity magnitude is 1), (iii) $t > 0$ is a point in time, (iv) $f(x, \Omega, t)$ is the kinetic density of particles with respect to the measure $dxd\Omega$, (v) $\sigma_s(x)$ is the scattering cross section, (vi) $\sigma_t(x)$ is the total cross section, and (vii) $\langle \cdot \rangle = \int_{\mathbb{S}^2} \cdot \, d\Omega$. In general $\sigma_t \geq \sigma_s \geq 0$; for the purposes of this paper, we set $\sigma_t = \sigma_s = 1$.

Moment equations are derived from (1). Let $\mathbf{m}(\Omega) = (m_1(\Omega), \ldots, m_M(\Omega))^T$ be a finite vector of real-valued, normalized spherical harmonics[2] (defined in Appendix A) of degree less than or equal to $N$ with length $M = (N+1)^2$. Define the finite vector of moments with respect to $\Omega$ as $\mathbf{u}_f(x,t) = \langle \mathbf{m} f \rangle$. Then according to (1), $\mathbf{u}_f$ satisfies

$$\partial_t \mathbf{u}_f + \nabla_x \cdot \langle \Omega \mathbf{m} f \rangle = -Q\,\mathbf{u}_f, \tag{2}$$

where $Q = \mathrm{diag}(0, 1, \ldots, 1)$. The system (2) is not closed because the flux $\langle \Omega \mathbf{m} f \rangle$ is a linear combination of moments up to degree $N+1$ whereas $\mathbf{u}_f$ only contains moments up to and including degree $N$.

[1] Although not discussed here, other efforts to solve the minimization problem have been documented in [3].

[2] The choice of spherical harmonics is not necessary, but it is common in radiation transport. This is because they are eigenfunctions of a more general scattering operator [32].

Download English Version:

https://daneshyari.com/en/article/518031

Download Persian Version:

https://daneshyari.com/article/518031

Daneshyari.com