

Integration Strategies for Efficient Multizone Chemical Kinetics Models

2010-01-0576

Published
04/12/2010

Matthew J. McNenly, Mark A. Havstad, Salvador M. Aceves and William J. Pitz
Lawrence Livermore National Laboratory

Copyright © 2010 SAE International

ABSTRACT

Three integration strategies are developed and tested for the stiff, ordinary differential equation (ODE) integrators used to solve the fully coupled multizone chemical kinetics model. Two of the strategies tested are found to provide more than an order of magnitude of improvement over the original, basic level of usage for the stiff ODE solver. One of the faster strategies uses a decoupled, or segregated, multizone model to generate an approximate Jacobian. This approach yields a 35-fold reduction in the computational cost for a 20 zone model. Using the same approximate Jacobian as a preconditioner for an iterative Krylov-type linear system solver, the second improved strategy achieves a 75-fold reduction in the computational cost for a 20 zone model. The faster strategies achieve their cost savings with no significant loss of accuracy. The pressure, temperature and major species mass fractions agree with the solution from the original integration approach to within six significant digits; and the radical mass fractions agree with the original solution to within four significant digits. The faster strategies effectively change the cost scaling of the multizone model from cubic to quadratic, with respect to the number of zones. As a consequence of the improved scaling, the 40 zone model offers more than a 250-fold cost savings over the basic calculation.

INTRODUCTION

Predicting the performance of an internal combustion engine solely through simulation is one of the most challenging goals of computational science. The large span of length and time scales is simply too great at this time to completely resolve the physical phenomena in the multi-dimensional, turbulent, chemically reacting flow present in-cylinder. Ambitious efforts are underway to resolve the smallest turbulent length scales and chemical time scales in a

laboratory-scale flame using direct numerical simulation (DNS) and detailed chemical kinetics methods on petascale computing architecture [1]. Without significant algorithmic improvements, the necessary computational resources for simulating this level of detail will not appear in the desktop computer of an engine designer in industry for at least another two decades given current trends in hardware development. In the interim, the practice of simplifying the physical models for complex in-cylinder flows to the level of available computational resources will continue. Such practices commonly depend on one or more of the following: reduced dimensionality and geometric detail, reduced chemical kinetic mechanisms, simplified turbulent transport models, and simplified bulk fluid dynamics. The multizone model is an example of the last type of simplification.

The multizone model approximates the in-cylinder flow as a series of coupled, well-mixed reactors. The model does not consider the fluid velocity in the cylinder; only the conservation laws for mass, energy and species are included in the governing equations. With no bulk fluid motion to resolve, costly CFD calculations of the reacting flow may be avoided allowing most of the computational effort to be dedicated to resolving the detailed chemical kinetics. This makes the multizone model particularly well-suited to approximate reacting flows with nearly homogenous species compositions where the reactions are limited by the chemical kinetics instead of the mixing process. Consequently, the multizone model is used by many investigators [2,3,4,5,6,7] to study the homogeneous charge compression ignition (HCCI) operating regime because the engine performance is dominated by the chemical reaction rates.

Several different schemes for coupling the zones in the multizone model have been successfully applied to HCCI engine operation. Aceves *et al.* [2] obtained good agreement between the multizone model and experimental HCCI engine data. In

the multizone model in [2], the pressure is assumed to be uniform throughout the cylinder, which couples all the zones to each other at every point during the cycle. Further, the multizone model [2] assumes no transport between the zones as a result of any differences in mass, energy, or species concentration. This assumption is made because the heat release event in HCCI operation generally occurs rapidly and uniformly across the cylinder. Komninou and Hountalas [7] also showed good agreement between the multizone model and experimental HCCI engine data, albeit using a different set of modeling assumptions. In particular, the pressure is only taken to be uniform across the zones after the reactions in each zone have been updated. This approach effectively decouples the zones during the time integration of the chemical kinetics reducing the basic solution cost; however, the rate of pressure equilibration across the cylinder is now connected with the numerical time step used to march the solution. The multizone model in [7] does include transport between adjacent zones. In particular, mass transport between zones is used to achieve a uniform cylinder pressure at the end of each chemical reaction step; and a conduction type model is used to capture any energy transport between zones.

The multizone model has recently been applied by Havstad *et al.* [8] to investigate the multi-cycle instabilities arising from the chemical kinetics during the transition from HCCI to spark ignited (SI) operation. In order to approximate the flame propagation through the cylinder during SI operation, a phenomenological mixing model is used to allow transport between zones in the multizone model of Aceves *et al.* [2]. The characteristic timescale of the mixing model is adjusted to match the experimental pressure rise rate reported by Wagner *et al.* [9]. The multi-cycle multizone model is able to reproduce several of the patterns found in the experimental cycle-to-cycle variations.

The focus of this investigation is to develop computationally efficient integration strategies for the type of multizone model used in [8]. The engine performance is obtained with the multizone model by integrating a set of first-order, ordinary differential equations (ODEs) representing the conservation of mass, energy and species in each zone. A detailed chemical kinetics mechanism generally involves a wide range of characteristic timescales, the fastest of which are several orders of magnitude smaller than the period of the engine cycle. It is therefore impractical to use an explicit ODE integrator (*e.g.* Runge-Kutta) because the time step is limited for numerical stability by the smallest timescale in the system. When a large range of timescales exists in the system of ODEs, the system is said to be “stiff.” Stiff ODE integrators are designed to take larger implicit time steps. However, updating the system state after an implicit time step typically requires a solution to a non-linear system of equations arising from the discretization of the time derivative. This adds an appreciable amount of complexity to the design of the stiff ODE integrator.

The purpose of this investigation is not to create a new stiff ODE integrator, but rather find out which integration strategies, using existing stiff ODE integrator packages yield the greatest performance gains for the fully-coupled multizone model. The most expensive part involved in the integration of the multizone model is the construction and linear system solution associated with the Jacobian matrix. The Jacobian matrix is used in stiff ODE integrators to iterate toward the non-linear solution that updates the system state after an implicit time step. Based on the structure of the Jacobian matrix for the multizone model, three integration strategies are proposed here to accelerate the computation. These include: (i) approximating the Jacobian as a banded matrix; (ii) using a set of quasi-independent well-mixed reactors to estimate the block diagonal terms of the Jacobian matrix; and (iii) using the quasi-independent well-mixed reactor estimate of the Jacobian as a preconditioner in an iterative Krylov-type linear system solver.

The three integration strategies are compared to the original approach in [2] using the basic features of the DVODE integrator developed as part of the EPSIODE/ODEPACK/VODE family of integrators at Lawrence Livermore National Laboratory [10,11,12,13]. The first strategy, based on the banded Jacobian, is considerably slower than the original. The second strategy, using a decoupled multizone model to generate an approximate Jacobian, yielded a 35-fold reduction in the computational cost for a 20 zone model. By using the same approximate Jacobian in the second strategy as a preconditioner for an iterative Krylov-type linear system solver, the third strategy achieves a 75-fold reduction in the computational cost for a 20 zone model. The faster strategies achieve their cost savings with no significant loss of accuracy. The pressure, temperature and major species mass fractions agree with the solution from the original integration approach to within six significant digits; and the radical mass fractions agree with the original solution to within four significant digits. The faster strategies effectively change the cost scaling of the multizone model from cubic to quadratic, with respect to the number of zones. Even greater performance gains can therefore be achieved for multizone models with more than 20 zones. As tested in this investigation, the 40 zone model offers more than a 250-fold cost savings over the original calculation. This integration strategy for the coupled multizone model now offers the opportunity to calculate hundreds of cycles in the time it used to take the original method, which allows the multi-cycle stability simulations conducted in [8] to be performed in a few hours instead of several days.

It is worth noting that the reductions in computational cost cited above are based on comparisons between the different integration schemes with no modifications to the ODE integrator source code. As distributed, the source code for the EPSIODE/ODEPACK/VODE family of integrators includes linear system solvers based on the other LINPACK linear

algebra library. The newer LAPACK linear algebra library [14] has largely superseded LINPACK offering marked improvement on modern computer architectures especially for large matrices. For the largest matrices involved in the ODE integration of the coupled multizone model, the LAPACK linear system solvers are approximately 15 times faster than the original LINPACK library. Modifying the ODE integrator source code to use LAPACK for the original integration strategy in [2] results in a 2.4-fold reduction in cost for a 20 zone model, and a 4.3-fold reduction in cost for a 40 zone model. The proposed integration strategies achieve a significant cost reduction by reducing the size of the linear systems solved in the implicit ODE integration. The matrices in the proposed strategies are sufficiently small that there is no appreciable performance difference between LINPACK and LAPACK. Therefore, the performance improvements offered by the proposed strategies are now somewhat lower using LAPACK: a 28-fold reduction in the computational cost for a 20 zone model; and a 60-fold reduction in the computational cost for a 40 zone model. The LAPACK library tested is tuned specifically for the architecture of a 2.4 GHz Intel Quad Core processor, which is used in all the LAPACK-based calculations presented here. The choice of linear algebra library does not impact the asymptotic trends presented. The original strategy for the coupled multizone model in [2] still scales as $O(N_z^3)$ and the proposed strategies scale as $O(N_z^2)$, with the number of zones N_z . Unless stated otherwise, the performance gains discussed here refer to the different integration strategies using the original source code for the ODE integrator built with the older LINPACK library.

MODEL DESCRIPTION

The multizone chemical kinetics model for internal combustion represents a single cylinder of an engine as a set of N_z coupled well-mixed reactors or zones. A schematic of the model showing the zone interactions is given in Figure 1. Each zone ($i = 1, \dots, N_z$) has a temperature T_i , mass m_i and species composition associated with it. The species composition of the i^{th} zone is given as a set mass fractions $y_{i,1}, y_{i,2}, \dots, y_{i,N_s}$ for the N_s species considered in the chemical mechanism. The time evolution of these flow quantities for the trapped mass captured in the cylinder between intake valve closing and exhaust valve opening is given by the following set of coupled, first-order, ordinary differential equations for each zone ($i = 1, \dots, N_z$):

$$\frac{dm_i}{dt} = F_i^{(m)}, \quad (1)$$

$$\frac{dy_{i,j}}{dt} = \frac{\omega_{i,j} W_j}{\rho_i} + F_i^{(y)}, \text{ for } j = 1, \dots, N_s, \text{ and} \quad (2)$$

$$\frac{dT_i}{dt} = \frac{1}{\rho_i \bar{c}_{p,i}} \left(\frac{dp}{dt} - \sum_{j=1}^{N_s} \omega_{i,j} h_{i,j} W_j - Q_{w,i} \right) + F_i^{(T)}. \quad (3)$$

Equations (1,2,3) are the conservation laws for mass, species and energy, respectively. The terms $F_i^{(m)}$, $F_i^{(y)}$ and $F_i^{(T)}$ represent the transport of the mass, species and energy (respectively) to the i^{th} zone from adjacent zones ($i-1$ and $i+1$ - excluding the first and last zones). In the remaining terms, the subscript i refers to the property in the i^{th} zone, the subscript j refers to the property for the j^{th} species, and the overbar refers to a mixture averaged property. Here $\omega_{i,j}$ is the species molar production rate, W_j is the molecular mass, ρ is the fluid density, $c_{p,i}$ is the specific heat at constant pressure, p is the global cylinder pressure, and $h_{i,j}$ is the specific enthalpy. The term $Q_{w,i}$ in (3) represents the volumetric heat loss rate to the wall, which is modeled here using a simplified Woschni correlation [15]. The computational speedup achieved in this investigation for the multizone model is not affected by the choice of the wall heat transfer model; therefore, any of the more recent heat transfer models developed specifically for HCCI may be used as well (see [16] as an example).

By selecting only adjacent zones for transport coupling, as is done here, the in-cylinder distribution of the fluid properties is assumed to be one-dimensional. One can imagine the zones as either a set of disks stacked in the direction of the piston motion, or as set of concentric shells depending on the effective area and associated time constants governing the transport between the zones. The treatment of the zone-to-zone transport terms can vary significantly between different instances of the multi-zone model while still showing agreement with experimental engine performance data. As an example, the transport terms are taken to be zero in the HCCI cycle model validated by Aceves *et al.* in [2]. More detailed treatments of the transport phenomena, such as the conduction type model used by Komminos and Hountalas [7] for energy exchange between zones, are also found to perform similarly well. The computational speedup achieved in this investigation is expected to apply to any choice of the zone-to-zone transport model including higher dimension zone coupling schemes.

One of the main assumptions of the multizone model considered here is that the engine does not experience knock, and the in-cylinder pressure is therefore uniform across all the zones. This is equivalent to saying the flow induced by local pressure gradient has a minor impact on the overall engine performance, which is




Figure 1. Schematic of the fully-coupled multizone model.

consistent with omitting the conservation of momentum from the governing set of equations (1,2,3). Using the uniform pressure assumption and the ideal gas law, the pressure derivative appearing in (3) can be expressed in terms of quantities that are solved directly. Specifically,

$$\frac{dp}{dt} = \frac{R_u}{V} \dots, \quad \left(\sum_{i=1}^{N_z} m_i T_i \sum_{j=1}^{N_s} \frac{1}{W_j} \frac{dy_{i,j}}{dt} + \sum_{i=1}^{N_z} m_i \frac{dT_i}{dt} \sum_{j=1}^{N_s} \frac{y_{i,j}}{W_j} + \sum_{i=1}^{N_z} \frac{dm_i}{dt} T_i \sum_{j=1}^{N_s} \frac{y_{i,j}}{W_j} - \frac{1}{V} \frac{dV}{dt} \sum_{i=1}^{N_z} m_i T_i \sum_{j=1}^{N_s} \frac{y_{i,j}}{W_j} \right) \quad (4)$$

where the rate of change of the cylinder volume dV/dt is given in terms of the engine speed and geometry (*i.e.*, bore, stroke and compression ratio) using the slider-crank equation. It is important to note that as a consequence of the uniform pressure assumption, each reacting zone is coupled to all the others through the pressure derivative (4). It is this pressure coupling across all zones, and not the coupling from the transport terms, that leads to most of the computational cost of the multizone method using the basic solution tools.

The remaining terms in the governing set of equations (1,2,3) are determined from the system variables (zone mass m_i , temperature T_i and species mass fractions $y_{i,j}$) using the ideal gas law, thermodynamic relationships and the user-defined reactions in the chemical kinetics mechanism. The cylinder pressure is obtained from the sum of partial pressures exerted on the system by each zone; that is,

$$p = \sum_{i=1}^{N_z} \frac{m_i R_u T_i}{V} \sum_{j=1}^{N_s} \frac{y_{i,j}}{W_j}. \quad (5)$$

The mixture density in the i^{th} zone is then given by

$$\rho_i = \frac{p}{R_u T \sum_{j=1}^{N_s} y_{i,j} / W_j}, \quad (6)$$

and the specific heat at constant pressure of the mixture (mass units) is

$$\bar{c}_{p,i} = \sum_{j=1}^{N_s} y_{i,j} c_{p,j}. \quad (7)$$

The mixture is assumed to be thermally perfect; that is, the individual specific heats $c_{p,j}$ and enthalpies $h_{i,j}$ of the species are only a function of the zone temperature T_i . The exact temperature dependence for the thermodynamic variables $c_{p,j}$ and $h_{i,j}$ is typically supplied as a database of empirical fitting functions, such as the one calculated for the NASA chemical equilibrium code CEA, developed by Gordon and McBride [17-18]. The final term needed to solve the governing equations (1,2,3) is the molar production rate $\omega_{i,j}$. The molar production rate in each zone can be calculated as a function of pressure, temperature and species mass fractions with the user-specified chemical kinetics mechanism and the associated reaction rates (see Chapter 2 of [20] for more details on the calculation).

BASIC SOLUTION METHODOLOGY

The governing equations (1,2,3) are a system of coupled, first-order, ordinary differential equations (ODEs). The molar production rate $\omega_{i,j}$ appearing in (2-3) typically varies by several orders of magnitude across the chemical system for engine combustion problems. As a consequence, the system of ODEs is referred to as being “stiff” because of the large range timescales that must be resolved in a solution. Explicit ODE integrators are straightforward to develop, but are computationally inefficient for stiff systems because the explicit time step is limited to the smallest timescale for numerical stability. Implicit ODE integrators are able to reduce the total number of time steps by removing the stability limit and taking larger steps. However, obtaining the updated system variables at the end of the larger, implicit time step requires a solution to a system of non-linear equations based on the governing ODE system (1,2,3) and the discretization scheme for the time derivatives. Implicit ODE integrators for stiff systems are much more complicated to




Figure 2. Hierarchal organization of the numerical tasks in the implicit ODE integrator used for the basic multizone calculation in this investigation. The values given in the parentheses are the fraction of the total simulation time spent performing a given tasks and all included subtasks.

develop than their explicit counterparts. Fortunately, a number of well-established and freely-distributed stiff ODE integrators are available, such as the EPSIODE/ODEPACK/VODE collection developed at Lawrence Livermore National Laboratory [10,11,12,13].

The organization of the numerical tasks needed by an implicit ODE integrator is given in Figure 2. The value appearing in the parentheses below each task name is the fraction of the total simulation time spent on the task and any associated subtasks for the basic coupled multizone model with 20 reacting zones. Although only the DVODE integrator (a double precision variant of the VODE integrator described in [12]) is shown in Figure 2, the basic tasks are essentially the same for most implicit ODE integrators. The only differences would be in the choice of the numerical method to solve the linear system, and the means to construct the Jacobian matrix of the system.

In the coupled multizone model, almost all the computation effort is spent in the ODE integrator. Most robust integrators are built on a function that integrates the ODE system in a single time step, in DVODE this function is called DVSTEP. The function DVSTEP is a general time stepper that manages the necessary function calls (or subtasks) to complete the

integration and controls the method order and time step size to maintain the solution accuracy to within user-specified tolerances. The computational cost of an implicit time step is due almost entirely to the nonlinear system solver - the tasks associated with the error control account for only 0.1% of the total simulation cost.

<figure 2 here>

To illustrate why a nonlinear system solver is needed and the role of the Jacobian matrix in its solution, consider a general system of N ordinary differential equations with N state variables,

$$\begin{aligned}\frac{\partial x_1}{\partial t} &= f_1(t, x_1, \dots, x_N) \\ \frac{\partial x_2}{\partial t} &= f_2(t, x_1, \dots, x_N) \\ &\vdots \\ \frac{\partial x_N}{\partial t} &= f_N(t, x_1, \dots, x_N).\end{aligned}$$

(8)