Available online at www.sciencedirect.com
ScienceDirect
Journal of Computational Physics 225 (2007) 1175–1197
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

Consistent properties reconstruction on adaptive Cartesian meshes for complex fluids computations

Guoping Xia [*,1], Ding Li [2], Charles L. Merkle [3]

Purdue University, West Lafayette, IN 47907, USA

Received 19 September 2006; received in revised form 22 January 2007; accepted 29 January 2007
Available online 8 February 2007

Abstract

An efficient reconstruction procedure for evaluating the constitutive properties of a complex fluid from general or specialized thermodynamic databases is presented. Properties and their pertinent derivatives are evaluated by means of an adaptive Cartesian mesh in the thermodynamic plane that provides user-specified accuracy over any selected domain. The Cartesian grid produces a binary tree data structure whose search efficiency is competitive with that for an equally spaced table or with simple equations of state such as a perfect gas. Reconstruction is accomplished on a triangular subdivision of the 2D Cartesian mesh that ensures function continuity across cell boundaries in equally and unequally spaced portions of the table to C^0, C^1 or C^2 levels. The C^0 and C^1 reconstructions fit the equation of state and enthalpy relations separately, while the C^2 reconstruction fits the Helmholtz or Gibbs function enabling EOS/enthalpy consistency also. All three reconstruction levels appear effective for CFD solutions obtained to date. The efficiency of the method is demonstrated through storage and data retrieval examples for air, water and carbon dioxide. The time required for property evaluations is approximately two orders of magnitude faster with the reconstruction procedure than with the complete thermodynamic equations resulting in estimated 3D CFD savings of from 30 to 60. Storage requirements are modest for today's computers, with the C^1 method requiring slightly less storage than those for the C^0 and C^2 reconstructions when the same accuracy is specified. Sample fluid dynamic calculations based upon the procedure show that the C^1 and C^2 methods are approximately a factor of two slower than the C^0 method but that the reconstruction procedure enables arbitrary fluid CFD calculations that are as efficient as those for a perfect gas or an incompressible fluid for all three accuracy levels.

Keywords: Real fluid; Property reconstruction; Equation of state; Adaptive Cartesian method

* Corresponding author. Tel.: +1 765 496 2849; fax: +1 765 494 0530.
E-mail address: xiag@purdue.edu (G. Xia).
[1] Research Associate, School of Mechanical Engineering, Purdue University.
[2] Research Assistant Professor, School of Mechanical Engineering, Purdue University.
[3] Reilly Professor of Engineering, Purdue University.

0021-9991/$ - see front matter
doi:10.1016/j.jcp.2007.01.034

Nomenclature

E	flux vector
g	Gibbs function
H	source term in Euler equations
h	enthalpy
h^0	stagnation enthalpy
Q	primary transport variable
q	coefficient of a polynomial
T	temperature
t	physical time
u	Cartesian velocity component
x	Cartesian coordinate
Y	species mass fraction
ρ	density
τ	pseudo-time

1. Introduction

Computational fluid dynamics requires the coupled solution of the partial differential equations that comprise the basic conservation laws (mass, momentum and energy) in combination with an auxiliary set of constitutive relations, generally expressed in algebraic form, that describes the physical properties of the fluid of interest. Taken separately, the conservation laws constitute an underdetermined system with more unknowns than equations. The fluid property relations close the system. Two common types of fluids, perfect gases and incompressible fluids, are used in the bulk of all CFD applications. These fluid characteristics are frequently combined directly with the conservation relations before they are coded resulting in algorithms that are specialized to one of these two equations of state. There are, however, many engineering applications in which gases or vapors do not follow the perfect gas laws and liquids can not be treated as incompressible. For such applications, it is necessary to turn to more general equations of state.

One approach is to use the family of classical semi-theoretical expressions that include the van der Waals, Peng-Robinson [1,2] and Redlich–Kwong–Soave (RKS) [3] equations of state in which the perfect gas relation is modified by adding additional terms and constants to improve accuracy in regions where real gas effects become important. Although these methods provide pressure–temperature–density relations that are accurate over a broader range, they generally do not provide an analogous expression for the internal energy and relations analogous to those for perfect gases are often used for the specific heats.

A second alternative is to use more complete and accurate thermodynamic databases such as REFPROP [4,5] and SUPERTRAPP [6,7] which provide highly accurate properties across the liquid–vapor-supercritical regimes for a large variety of fluids including full pressure–density–temperature-internal energy relations. The REFPROP database is based upon nonlinear regression of all available data for approximately forty common fluids and refrigerants. The companion database, SUPERTRAPP provides predictive capabilities for the properties of more than 200 different hydrocarbon components and their mixtures. These databases provide consistent expressions for the pressure–density–temperature relation and the enthalpy as well as values for the transport properties. An additional widely used database is the SESAME tables which provide characteristics of materials at extreme conditions [8].

In using any of these equations of state, a variety of difficulties can be encountered that increases the cost and complexity of a CFD solution. First, the various property formulations may be expressed in terms of independent variables pairs that do not match those used in the CFD algorithm. This incompatibility requires iterative property evaluations that increase the cost and complexity of solutions. Second, the algebraic expressions in the more global formulations are so complex that their evaluation is far more expensive

than the CFD solution procedure itself. Real fluid solutions based directly upon such databases can be more than an order of magnitude slower than perfect gas solutions. A third issue is that the CFD code must be modified each time the properties formulation is changed leading to increased effort in switching between different fluids.

Previous efforts on the evaluation of real fluid properties in CFD simulations fall into four categories: methods that use of one of the above-referenced semi-theoretical formulations as seen in the work by Edwards et al. [9], Meng and Yang [10] and Zilliac and Karabeyoglu [11]; methods that use local curve/surface fitting of the properties as in the work by Coirier [12]; methods based upon tabular reconstruction of the properties as in the work by Hosangadi and Ahuja [13]; and direct evaluation of properties from a database (NIST [5]) as in the work by Davis and Campbell [14]. In addition to these, Cheng and Farmer use a combination of the semi-theoretical formulation and curve fitting methods for real fluid calculations [15].

In the present paper we first ascertain the impact of the constitutive relations on numerical solutions and then develop a generalized properties evaluation procedure based upon an adaptive table look-up method that provides flexibility, efficiency and accuracy for nearly any conceivable property formulation. The flexibility comes because the equation of state is inverted by separate software outside the CFD code thereby enabling the independent variables pair to be chosen by the user while simultaneously enabling the entire variety of equations of state to be incorporated in a CFD code in a common fashion. The efficiency comes from using a tree-based data structure that provides fast table look-up over large or small thermodynamic domains. The accuracy comes from an adaptive Cartesian mesh that adjusts the table resolution to user specified input criteria and evaluates both the thermodynamic functions and their derivatives. Although the procedure is applicable to any properties formulation, we base our examples on information obtained from REFPROP [5]. The same approach can be used with any other data base, with any of the semi-theoretical EOS formulations, or even with perfect gases or incompressible fluids.

The REFPROP database expresses fluid properties as the sum of several lengthy series formulations whose coefficients are based upon regression analyses of available experimental data coupled with theoretical insight. The end result is a single (though complex and lengthy) algebraic expression that retains high accuracy for a given fluid over the entire thermodynamic range (vapor, liquid, supercritical and two-phase regions) for which data is available. The fundamental thermodynamic variable used in REFPROP is the Helmholtz function which is expressed in terms of the density–temperature independent variables pair. The Helmholtz function and its partial derivatives enable both the pressure and the enthalpy to be determined from temperature and density in a consistent fashion. Separate equations are also provided for the viscosity and thermal conductivity as functions of temperature and density. The resulting routines are very effective when hundreds or thousands of property evaluations are needed, but they become prohibitively slow when billions of property evaluations are required as is typical for a three-dimensional CFD computation on even a modest grid of a few million cells. Because REFPROP is not effective for CFD solutions, more efficient methods must be devised if the flexibility and generality of the REFPROP database are to be incorporated into a CFD code. Our goal is to define a procedure that retains the flexibility of REFPROP (highly accurate properties over wide thermodynamic ranges for a large number of fluids) while vastly decreasing property evaluation time. Simultaneously, we wish to provide accuracy that is commensurate with the requirements of the problem of interest. This means that it must be possible to specify the desired degree of accuracy in the evaluation routines. In determining accuracy, we inherently assume that the REFPROP routines (or other thermodynamic source) are precise descriptions of the fluid. As REFPROP or other databases are updated, or if it is desired to switch to other equations of state, we need only re-tabulate our procedures to meet this updated physical representation.

In the following sections, we first describe the adaptive Cartesian methods that are used to develop the reconstruction procedure for the requisite fluid properties along with the techniques used for evaluating and populating the look-up table. We then briefly present the equations of motion to demonstrate the specific property requirements and to identify the property derivatives that are needed. Following this we describe appropriate triangulation methods that enable property reconstruction with three different degrees of function continuity and discuss the related issue of thermodynamic consistency. In the final sections we present timing comparisons for adaptive reconstruction with complete function evaluations and present two CFD examples that assess both consistency and computational cost.

2. Adaptive Cartesian grid methods

The simplest data structure for tabular reconstruction is one that spans the region of interest with equally spaced intervals. Equal spacing provides very efficient table look up because the reconstruction interval can be computed directly without searching, but results in very large table sizes. The level of accuracy is set by the region in which the function varies most rapidly so that slowly varying regions must be over-refined to accommodate high-gradient regions. Unequally spaced tables that refine locally as needed are much more efficient in storage, but require a time-consuming trial and error search to find the appropriate reconstruction interval. In the present paper we choose an alternative that matches the fast search capabilities of equally spaced tables while still retaining a local refinement capability. This method is a tree-based, Cartesian adaptive mesh structure. Although the method is capable of using any independent pair of thermodynamic variables, for definiteness, we use pressure and temperature.

Tree-based Cartesian methods are predicated upon resolving the features of a rectangular (or hexahedral) region by local subdivisions obtained by dividing all rectangles (hexahedra) into four (eight) smaller rectangles (hexahedra) [16,17]. For the two-dimensional problem, which is pertinent for thermodynamic surfaces of a pure fluid, we begin by defining a rectangular region in the x–y plane that includes the physical domain of interest. (For simplicity, we explain the method by reference to the x–y plane. The interpretation is then easily transferred to the p–T, or other thermodynamic plane). The requisite properties are then evaluated at each of the four corners of the square and an appropriate reconstruction is used to evaluate these properties at preselected points in the square. (Usually a grid of 25 equal spaced points are used in a square.) The reconstruction is accomplished by subdividing the square into triangles in a consistent manner (additional details are given later) and using bi-variate interpolation in each triangle. The reconstructed properties and their derivatives are then compared with their 'exact' values obtained from the complete property equations and used to assess the reconstruction error. If any error in a square exceeds a user-specified threshold, it is subdivided into four smaller squares with property evaluations again being computed and the trial reconstruction procedure repeated for each new square. As the interpolation grid is refined, relatively flatter portions of the surface will fall within the error tolerance and will not require further refinement while squares in regions of larger gradients will be refined further. Keeping track of these adaptively refined regions in a tree-based structure enables rapid search procedures that approach the speed of those for equally spaced tables. Results given later indicate that look-up times for this tree-based data structure are approximately 10% slower than those for an equally spaced table and approximately 100 times faster than those required for evaluating the complete thermodynamics formulation in REFPROP. Test results also indicate that data retrieval times are essentially independent of table size.

The primary advantage of the sub-divided cell structure is that a relatively few refinement levels results in very small subdivisions. The rapidly increasing resolution of powers of two implies that very accurate reconstructions can be done with table sizes that are easily accommodated on today's computers. For example, 15 levels of refinement correspond to dividing the original thermodynamic region into squares whose edges are 3×10^{-5} times smaller than those of the original interval while 25 levels reduce the edge lengths to intervals of approximately 3×10^{-8} of their original size. For a temperature range covering 3000 K, a refinement of twenty levels corresponds to interpolation over intervals as small as 0.003 K, while still producing practical table sizes. As a practical matter, limits must be placed on the total number of levels of refinement allowed, but these limits are typically not reached unless the domain crosses discontinuities such as the liquid–vapor line.

The Cartesian grid method which has been popularized for adaptive grids in CFD [16,17] is used to provide uniform accuracy over an arbitrarily selected reconstruction domain. A graphical example of a representative data structure is given in Fig. 1 and the typical binary tree that organizes the grids on a sequence of levels is shown in Fig. 2. The Cartesian map in Fig. 1 corresponds to carbon dioxide in a region of p–T space that spans the liquid–vapor line. The intersections in the figure show the grid locations where properties data are stored. The finished table contains relatively large interpolation cells far from the phase-change line and very small sizes adjacent to the phase-change line. To demonstrate the characteristics of the Cartesian reconstruction method, we have intentionally generated a coarse map based upon input parameters calling for an accuracy of 1% with a maximum six refinement levels. Practical applications for CFD calculations would typically involve additional refinement levels and considerably tighter accuracy tolerances. The binary

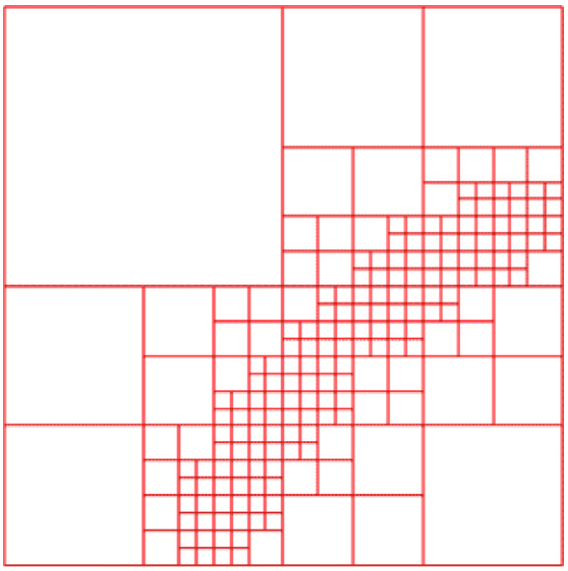

Fig. 1. Adaptive Cartesian grid.

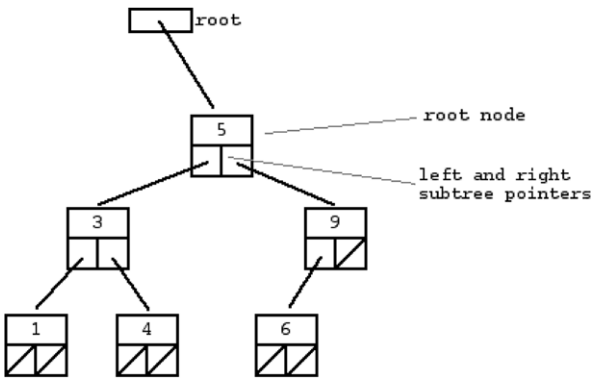

Fig. 2. Binary tree constructed by adaptive Cartesian method.

tree in Fig. 2 allows the repeated refinement of cells in regions where property variations dictate more refinement, as shown pictorially in Fig. 1. The apparently high irregularity of the grid is actually well organized and there are efficient methods for refining and searching in such grids.

A recursive procedure is usually used in traversing and refining a binary tree. For example, in order to find the cell within which a point is located, one starts from the root cell and visits every sub-cell until the one encompassing the point is found. This procedure is repeated in this new cell until the cell found is not further refined. The algorithms for a binary tree data structure are well documented in literature [18].

To provide flexibility for problems in which the domain of interest does not fall within a rectangular region in the thermodynamic plane, an arbitrary trapezoidal region can be mapped into a square on a computational plane as shown in Fig. 3 before generating the properties table and the tree structure. Two boundaries of the quadrilateral are specified as arbitrary functions of the chosen independent variables (e.g., the vertical boundaries in Fig. 3 are arbitrary functions of T while the remaining two sides are taken as constant values of the second variable ($\log(p)$) in Fig. 3). Alternatively, straight boundaries may be placed on the vertical sides and curved boundaries on the horizontal sides. Trapezoidal domains of this nature generate a linear mapping from the original to the transformed space, thereby offering flexibility without increasing the complexity of the look-up procedure. Provision can be made to specify the arbitrary

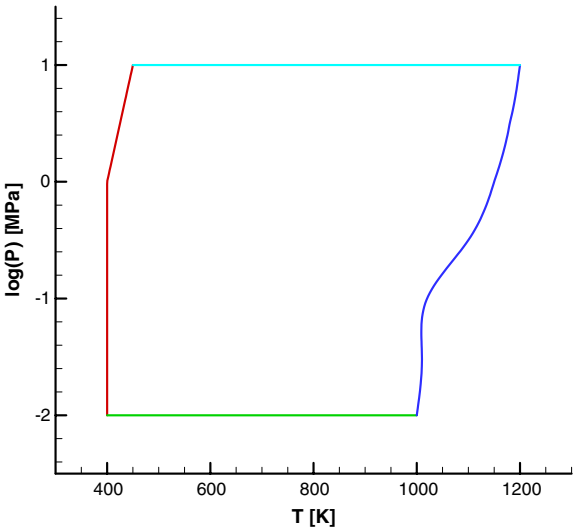

Fig. 3. General trapezoidal domain in p–T domain.

boundaries as a series of straight line segments, or as a continuous (e.g., a spline or polynomial) curve. Numerous transformation methods are available for mapping an arbitrary quadrilateral into a square. For the present applications we have used transfinite interpolation [19]. The Cartesian adaptive method is then applied in this mapped space.

3. The conservation equations and fluid properties

The computational solution of fluid dynamics problems requires both fluid properties and their derivatives. To identify the precise quantities needed, we introduce the equations of motion to indicate the manner in which these property derivatives appear. The conservation form of the mass, momentum and energy equations for an arbitrary, Newtonian fluid are:

$$\frac{\partial Q}{\partial t}+\frac{\partial E_i}{\partial x_i}=\frac{\partial V_i}{\partial x_i}+H \tag{1}$$

For a flow with multiple species, the vectors, Q and E, are given by

$$Q=\begin{pmatrix}\rho\\ \rho u_j\\ \rho h^0-p\\ \rho Y_k\end{pmatrix},\quad E_i=\begin{pmatrix}\rho u_i\\ \rho u_i u_j\\ \rho u_i h^0\\ \rho u_i Y_k\end{pmatrix} \tag{2}$$

The quantities; x_i and u_i represent Cartesian coordinates and velocity components; p and ρ represent the pressure and density; h^0 is the stagnation enthalpy; and Y_k is the species or phasic mass fraction. As a short hand notation, we have symbolically combined the three components of the momentum equation into one equation and included a single species conservation equation to represent any number greater than or equal to zero. No species equation is needed if a single species or phase is involved since the global continuity equation is given as the first conservation relation. The vector, H, represents a source term while V_i represents the viscous fluxes. Because the fundamental conservation laws contain no dependence upon the fluid, the relations given in Eq. (1) are identical for all types of fluids.

Before introducing the constitutive relations that close the system, we add a pseudo-time term to Eq. (1) to enable time-marching. By analogy with the physical time derivative, we express this pseudo-time term in the form, $\partial Q/\partial\tau$, where τ represents the pseudo-time. In the limit as pseudo-time goes to infinity, this added derivative vanishes and so has no impact on the final solution. Accordingly, it is not necessary to express it in

conservation form. The chain rule can be used to transform from the conservative variables, Q, to any other complete set of primary dependent variables. Here we work with the primitive variables, $Q_p = (p, u_i, T, Y_k)^T$, in which the pressure and temperature appear as the thermodynamic variables pair along with the velocity components and the mass fractions. After employing the chain rule transformation in the pseudo-time term, Eq. (1) becomes:

$$\frac{\partial Q}{\partial Q_p}\frac{\partial Q_p}{\partial \tau} + \frac{\partial Q}{\partial t} + \frac{\partial E_i}{\partial x_i} = \frac{\partial V_i}{\partial x_i} + H \tag{3}$$

The eigenvalues of this system, which control both convergence and accuracy of the numerical solution [20], are determined by the product of the inverse of the pseudo-time matrix and the flux Jacobian of the convective term. For general fluid dynamic applications, the coefficient matrix multiplying the pseudo-time derivative must be modified to ensure the eigenvalues are well conditioned in all regimes [20,21].

The inclusion of the pressure and temperature in the primary dependent variable, Q_p, is consistent with basing the fluid properties on the Gibbs function, $g(p, T) = h - Ts$, where s represents the entropy. In differential form the Gibbs function is given by:

$$\mathrm{d}g = -s\mathrm{d}T + \frac{1}{\rho}\mathrm{d}p \tag{4}$$

Comparing the partial derivatives of this differential form with those from the differential obtained from expressing the Gibbs function as an arbitrary function of pressure and temperature, $g = g(p, T)$, immediately yields relations for the density and the entropy

$$\rho(p, T) = \frac{1}{g_p} \quad \text{and} \quad s(p, T) = -g_T \tag{5}$$

Here g_p and g_T are the partial derivatives of Gibbs function with respect to pressure and temperature respectively. We can then obtain the enthalpy from the Gibbs function, its derivatives and the thermodynamic definition,

$$h(p, T) = g - Tg_T \tag{6}$$

Finally, we define the stagnation enthalpy in terms of the enthalpy and the velocity components,

$$h^0 = h + \frac{1}{2}\sum u_i^2 \tag{7}$$

The procedure for ideal mixtures of species and/or phases is identical except that the Gibbs function must be supplied for each component $g_k = g_k(p, T)$. Operations are simplified if the density is defined in terms of Amagat's law rather than the more traditional Dalton's law. The partial density of component (species or phase) k is defined as the mass of fluid k divided by the volume it would occupy if placed in an isolated container at the local temperature and pressure. For notational clarity, we define this partial density as $\tilde{\rho}_k$ (where the tilde signifies that this is the partial density defined by Amagat's law), while the corresponding partial enthalpy is h_k. Combining these gives the mixture density and mixture enthalpy, $\rho = \tilde{\rho}_l\alpha_l + \tilde{\rho}_v(1 - \alpha_l)$ and $h = h_l Y_l + h_v(1 - Y_l)$ where the volume fraction, α_k is related to the mass fraction as: $Y_l = \tilde{\rho}_l\alpha_l/\rho$. Real mixtures require additional terms [4].

The Jacobian matrix that multiplies the pseudo-time term in Eq. (3) and the product matrix from which the eigenvalues of the system can be found are important parameters in a numerical solution. For compactness of notation, these matrices are given for the one-dimensional equations with only a single species as:

$$\frac{\partial Q}{\partial Q_p} = \begin{pmatrix} \rho_p & 0 & \rho_T \\ \rho_p u & \rho & \rho_T u \\ \rho_p h^o + (1 - \rho h_p) & \rho u & \rho_p h^o + \rho h_T \end{pmatrix} \quad \text{and} \quad A_p = \frac{\partial Q_p}{\partial Q}\frac{\partial E}{\partial Q_p} = \begin{pmatrix} u & \rho h_T/\Delta & 0 \\ 1/\rho & u & 0 \\ 0 & (1 - \rho h_p)/\Delta & u \end{pmatrix} \tag{8}$$

where $\Delta = \rho_p h_T + \rho_T(1 - \rho h_p)/\rho$. Note that both matrices contain the density plus four thermodynamic derivatives: $\rho_{p,k}$, $\rho_{T,k}$, $h_{p,k}$ and $h_{T,k}$ (where $\rho_{p,k} = (\partial\tilde{\rho}_k/\partial p)_T$, *etc.*). These derivatives directly affect the fluids solution and must also be specified from the fluids database. The corresponding multi-dimensional/multi-component

expressions follow directly from these. Note that these four thermodynamic derivatives, ρ_p, ρ_T, h_p and h_T, can be evaluated in terms of the first and second partial derivatives of the Gibbs function:

$$\rho_p = -\frac{g_{pp}}{g_p^2}, \quad \rho_T = -\frac{g_{pT}}{g_p^2}, \quad h_p = g_p - Tg_{pT}, \quad \text{and} \quad h_T = -Tg_{TT} \tag{9}$$

Therefore, all quantities except the transport properties are contained within the Gibbs function and are interrelated.

As an aside, we note that by choosing the density and temperature (the variables corresponding to the Helmholtz function) as the primary thermodynamic variables in conjunction with the velocity components and mass fractions to give the variable set, $Q_\rho = (\rho, u_j, T, Y_k)^T$, again introduces four thermodynamic derivatives: $(\partial p/\partial\rho)_T$, $(\partial p/\partial T)_\rho$, $(\partial h/\partial\rho)_T$ and $(\partial h/\partial T)_\rho$, into the Jacobian matrix. In general, for any thermodynamic variables pair, and for any iterative solution method, two thermodynamic functions and four thermodynamic derivatives always appear in the equations governing the convergence process. Having accurate values for the two thermodynamic functions and their four partial derivatives is an important requirement for the properties routine.

In addition to these fluid properties data, the transport properties, molecular viscosity, μ, and thermal conductivity, k, must also be evaluated and are again taken as arbitrary functions of the two primary thermodynamic variables. The thermodynamic variables, their derivatives and the transport properties are all available from complete fluid properties databases such as REFPROP [5]. The challenge is to define an efficient procedure for evaluating these quantities.

4. Continuity and consistency in reconstruction

For a fluid computation based upon primitive variables, the six thermodynamic properties, ρ, ρ_p, ρ_T, h, h_p, and h_T, and the two transport properties, μ and k, must be evaluated as a function of p and T at every cell and every iteration. There are several potential ways to employ Cartesian adaptive methods to store and recover the data. The most straightforward approach is to store all eight properties at each node of the Cartesian grid and reconstruct all eight properties independently. While this procedure can provide continuous functions across cell boundaries, it produces inconsistent results because it ignores the interdependency of the density and the enthalpy and the four thermodynamic derivatives. To incorporate this characteristic in CFD solutions, the consistency of the property evaluations must be considered.

4.1. Thermodynamic consistency and reconstruction continuity

Thermodynamic consistency in property evaluations has been discussed by Swesty [22], who used the Helmholtz free energy as a basis function. His results show that failing to include proper consistency conditions in thermodynamic reconstruction can lead to significant errors in flowfield solutions, however, his reconstructions were based on much coarser meshes (and correspondingly lower accuracies) than the ones contemplated here. Our experience (discussed below) indicates that fine grid interpolation generally keeps the inconsistencies to acceptable levels. As a means for guaranteeing that thermodynamic consistency, he devised a higher-order reconstruction procedure that maintains second-order consistency throughout a piece-wise continuous domain of quadrilaterals. In analogous fashion, we consider three different levels of reconstruction involving C^0, C^1, and C^2 continuity across cell boundaries whose results range from partially to fully consistent. In place of the quadrilateral sub-domains used by Swesty, we decompose all squares in the mapped domain into triangles and reconstruct the property functions over these. With the global structure maintained through the quadrilateral grid, properties are retrieved by first locating the triangle in which their given p and T conditions lie and then using stored information at the vertices of the squares to interpolate to the desired degree of accuracy during numerical simulations.

In the C^0 method, values for each of the six thermodynamic fluid properties plus the two transport properties are stored at the vertices of each square in the mapped plane. The squares are then subdivided into triangles and the values of each function at the three vertices are used to construct independent bi-linear reconstruction functions for each property. The bi-linear reconstruction used in the C^0 method therefore

ensures function continuity along all faces of adjacent cells of the same size (the reason for limiting this statement to cells of the same size is given later) for each of the eight fluid properties but does not provide internal consistency inside the triangles. For example with C^0 reconstruction, both the density and its first derivatives vary linearly over each triangular region. Clearly a linear derivative is not consistent with a linear function. Further, the enthalpy and density have no interdependence on each other. Nevertheless, because the six properties stored at the vertices are themselves taken from a consistent thermodynamic database (as is the case for REFPROP), the internal inconsistencies will be small on a fine grid.

The alternative of using linear reconstruction over all triangles for the density and enthalpy and differentiating these linear functions to obtain the property derivatives results in improved consistency, but also produces discontinuities in the derivatives at the cell faces. Experience shows that these derivative discontinuities are much more detrimental to CFD solutions than the inconsistencies.

Incorporating property consistency in the reconstructed results requires the use of higher order functions. Higher order polynomials represent the most natural and straightforward manner in which to express such a local function. Here we use bivariate polynomials of the form [23,24],

$$z(x,y) = \sum_{i=0}^{n} \sum_{j=0}^{n-i} q_{i,j} x^i y^j \tag{10}$$

to represent the local solution. The number of undetermined constants in a polynomial of degree n is $(n+1)(n+2)/2$ where n is the degree of the polynomial. One advantage of the Cartesian subdivision is that these polynomials are especially suitable for triangles. The theory and implementation of polynomials on triangular meshes is well documented in finite element references (see, for example, Akima [23] and Preusser [24]) and both C^1 and C^2 reconstruction methods based upon the polynomial expansion of Eq. (10) have been considered. The C^1 reconstruction provides consistency between the density and its partial derivatives and the enthalpy and its partial derivatives, but does not provide consistency between density and enthalpy. We refer to the C^1 reconstruction as partially consistent. The C^2 reconstruction provides fully consistent results. Details of both methods are given below.

The C^1 method treats the density and enthalpy as unrelated functions that are each handled in analogous fashion. Using the density as an example, the goal of C^1 reconstruction is to reconstruct the density function in such a manner that the values of ρ, ρ_p and ρ_T are consistent within any given cell and that all three quantities are continuous across cell boundaries (i.e., the density is C^1 continuous while the derivatives are C^0 continuous). This level of consistency and continuity can be achieved by using a bi-quintic (fifth-order) polynomial of the form given in Eq. (10). The construction of a bi-quintic polynomial requires the determination of 21 coefficients. Storing the values of the density and its derivatives up to second order at each of the three vertices provides 18 conditions. The remaining three come from the requirement that the derivatives in the direction normal to each edge be continuous. This can be done either by prescribing the normal derivatives or by reducing the order of the polynomials. The later choice is the known as the 'condensation of parameters' method [24]. Since the density and enthalpy are treated independently, the enthalpy and its first two derivatives must also be stored at the vertices of all squares. Similar consistency and continuity are then ensured for the enthalpy and its first derivatives. The enthalpy and density are, however, unrelated on a given region so that consistency between these two functions is not guaranteed. Nevertheless, if both density and enthalpy are taken from a database such as REFPROP that is itself consistent, the degree of inconsistency in the C^1 approximation is quite small.

The C^2 method interpolates all properties and their derivatives by reconstructing the Gibbs function as a ninth-order polynomial which requires the determination of 55 coefficients in each triangular sub-element. Fifteen values are determined at every vertex of every square by storing the Gibbs function and its derivatives up to fourth order. After subdividing the squares into triangles, each triangular patch therefore encompasses three vertices with fifteen parameters at each vertex for a total of 45 conditions. Requiring C^2-continuity on each of the edges provides another nine conditions. The final condition can be specified at will without violating the smoothness constraints, while still keeping the fit to the vertex data. Additional details of the formulation are given in the literature [24]. The density, enthalpy and their four first derivatives are then obtained from the Gibbs functions using relations in Eqs. (5), (6) and (9). Thus, in the

C^2 method, only the Gibbs function is reconstructed and all six thermodynamic properties are computed from this single function resulting in properties that are fully consistent with each other and the underlying Gibbs function. Specifically, all six thermodynamic properties are based upon higher order variations of the Gibbs function across each triangle with first- and second-derivative continuity across all adjoining faces. This results in complete consistency and continuity among all six properties. The resulting density and enthalpy reconstructions are consistent and C^1 continuous, while their derivatives are consistent and C^0 continuous.

For all three reconstruction levels, the coefficients of the piecewise function are stored for each triangle once they have been calculated and need not be updated in a CFD calculation. In the C^0 method, a total of 24 coefficients are stored in each triangle for the linear interpolation of the six thermodynamic properties ρ, ρ_p, ρ_T, h, h_p and h_T, and the two transport properties, μ and k. The transport properties are always interpolated linearly in our work. In the C^1 method, 48 coefficients are stored, including 21 each for density and enthalpy, and 3 each for the two transport properties. In C^2 method, there are 61 coefficients stored in each triangle, with 55 for the Gibbs function and 6 for the transport properties. In the following section we discuss appropriate triangulation methods for effectively implementing the C^0, C^1, and C^2 methods on the adaptive Cartesian grid.

4.2. Triangulation on unequally sized squares

In equally spaced regions of a Cartesian mesh where all neighbors are the same size, the squares can be split into triangles by simply dividing along either diagonal. Any of the three reconstruction levels described above can then be applied on the resulting triangles. The situation, however, is somewhat more complex when the sizes of two neighboring cells are unequal. Consider, for example, the line ***l*** shown in Fig. 4 which crosses face AF between cells of two different refinement levels. The simple diagonal triangulation indicated in Fig. 4 results in a property discontinuity at this interface for either the C^0, C^1 or C^2 reconstruction procedures because the properties on the left side of the interface (point, P^{L}) are determined by values of the function at vertices **A** and **F**, while those on the right side (point P^{R}) are determined by vertices **A** and **D**.

This difficulty is demonstrated in Fig. 5 which compares the reconstructions along line ***l*** of the density (left) and its pressure derivative, ρ_p (right), based upon the triangulation shown in Fig. 4 with the exact solution as functions of temperature. In addition, the reconstruction based upon an alternative triangulation discussed below is also presented. The discontinuity at the interface for the simple diagonal triangulation (Fig. 4) is readily visible. The reconstructions in Fig. 5 were obtained by using the C^0 method, but reconstructions based on C^1 or C^2 exhibit similar discontinuities. Further, the discontinuity exists regardless of which diagonal is used to bisect the squares.

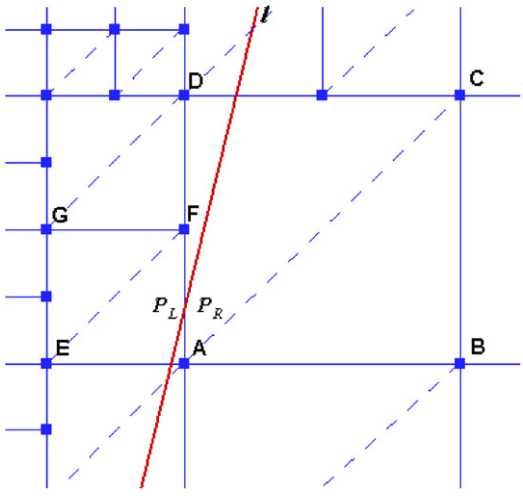

Fig. 4. Details of diagonal triangulation on a square mesh at a change in cell size.

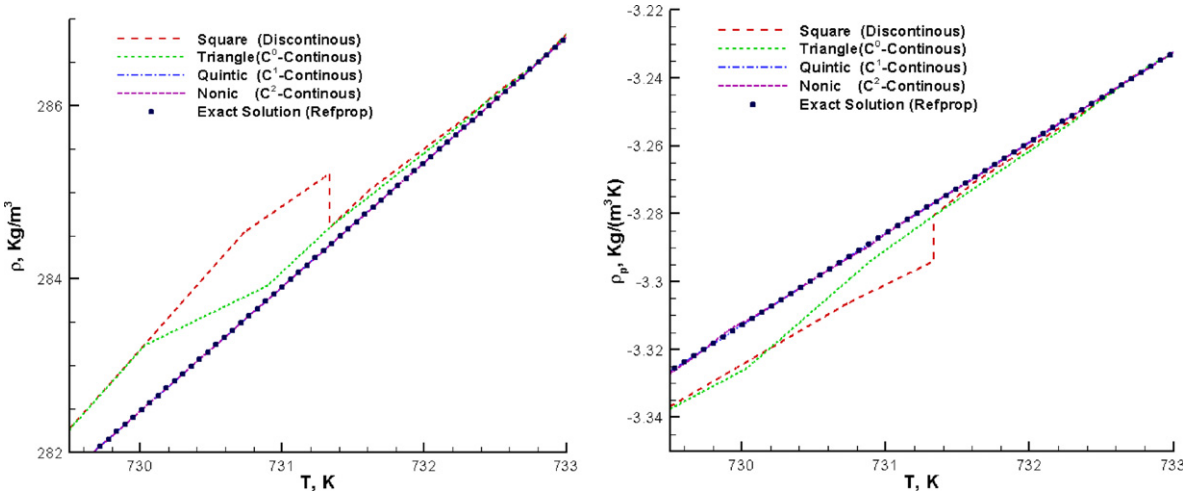

Fig. 5. Reconstructed values of ρ and ρ_p for triangulations in Fig. 4 (long dashed line) and Fig. 8 (short-dashed and dotted lines) along line l in Figs. 4 and 8. Exact solution given by dotted points.

To demonstrate the impact that discontinuities of this nature have on CFD calculations, we show the results of one-dimensional CFD solution in Fig. 6 and the corresponding convergence plots in Fig. 7. The axial variation of the temperature in Fig. 6 indicates that the discontinuity in properties at a cell size change produces an unphysical wiggle in the solution. The corresponding convergence curve in Fig. 7 also indicates that the convergence stalls when it reaches a level consistent with the magnitude of the discontinuity in the reconstructed functions. The remaining convergence curves are discussed below.

4.3. Modified triangulation for ensuring continuity

The difficulties arising from unequally sized squares can be rectified by triangulating the Cartesian mesh in a manner that ensures the reconstructions on both sides of the interface are based on the same data. One acceptable triangulation pattern for a cell with two refined and two unrefined neighbors is demonstrated in Fig. 8. As shown by the subdivision in the figure, the properties on the left and right sides of the boundary at point **P** are each based upon similar values and result in the specified degree of function continuity. The

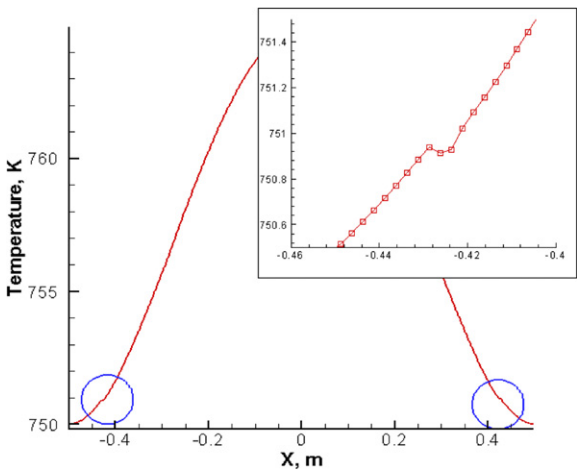

Fig. 6. Temperature variation in 1D numerical simulation of flow through a converging–diverging nozzle with reconstruction based on diagonally divided squares (Fig. 4) illustrating wiggles created in solution by property discontinuities at cell faces.

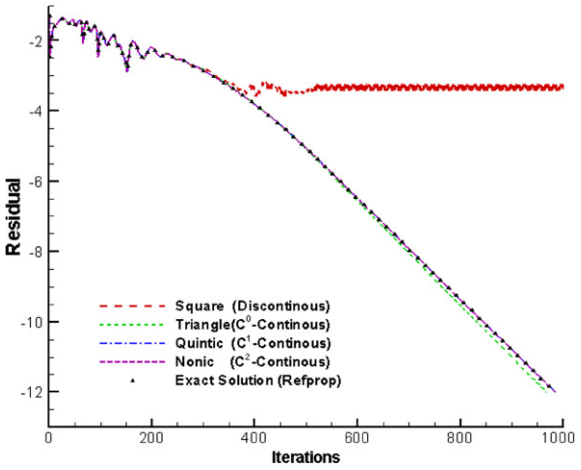

Fig. 7. Exemplary plot showing effect of triangulation on convergence of a specific CFD calculation. Red: Triangulation taken from Fig. 4; Green, blue and violet: Triangulation taken from Fig. 8; Triangles: properties evaluated directly from REFPROP. (For interpretation of the references in colour in this figure legend, the reader is referred to the web version of this article.)

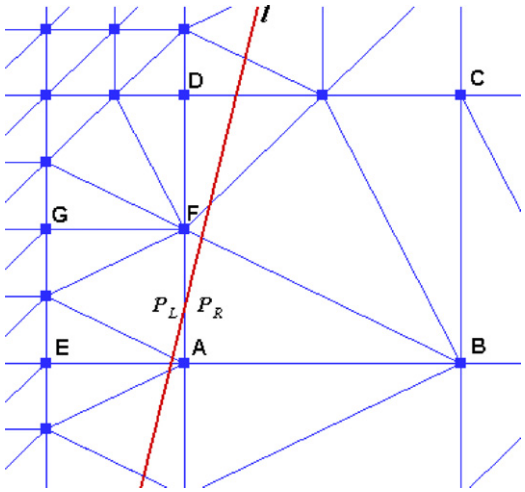

Fig. 8. Triangulation method chosen to ensure continuous reconstruction in a region between two cell sizes.

reconstructions of the properties along line ***l*** across the boundary are now both determined by the values at **A** and **F**. The continuity of the piecewise function across the interface is thus guaranteed as exhibited by the C^0, C^1 and C^2 reconstruction curves in Fig. 5. The C^0 reconstruction results in a small local error, but this has no adverse impact on the convergence rate (Fig. 7) or the solution (not shown). The reconstructed functions for the C^1 and C^2 reconstructions also provide excellent convergence as Fig. 7 shows. Finally, Fig. 7 shows the convergence rate for a solution for which the properties were taken directly from REFPROP without any reconstruction method. This 'exact' property evaluation method resulted in exactly the same convergence as any of the three reconstruction levels (although it took much more CPU time). Consequently, we see that a proper triangulation results in both efficient convergence and accurate solutions.

Similar triangulations can be found for nearly any type of mesh topology, but to limit the number of patterns to a countable level, we prohibit cell level changes greater than one in neighboring cells during the initial database setup. Thus the ratio of face length in two adjacent cells will be either one or two. This limitation, which is also commonly used in tree-based adaptive grid methods [16,17], reduces the number of patterns to six: cells with four undivided faces, cells with one divided and three undivided faces; cells with two adjacent faces divided; cells with two opposite faces divided; cells with three faces divided; and cells with all four faces

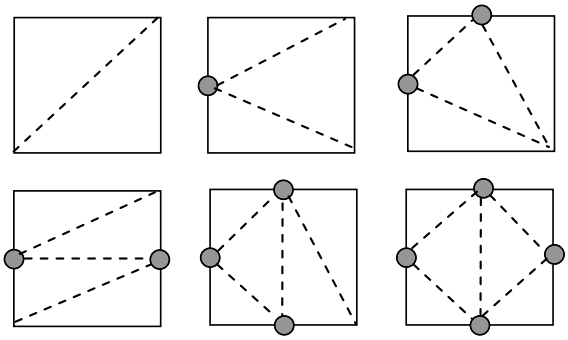

Fig. 9. Six possible cell refinement patterns indicating acceptable triangulation patterns for each.

divided. These six patterns are shown in Fig. 9 along with appropriate triangulations. Because only six patterns must be recognized, the logic for the triangulation is straightforward and the overhead for locating the triangles within a given square is small so that the search advantage of the Cartesian mesh is not lost. The reconstruction process applies equally to any triangle whether it is a portion of a regular region of the mesh or one with unequal sides. Consequently, the modified triangulation adds no complexity to the reconstruction procedure and allows the overall structure of the Cartesian grid to be retained.

5. Storage and timing comparisons

In the present section we present some timing comparisons for the adaptive reconstruction procedure and some representative reconstruction results. To provide perspective, we consider properties evaluations for three fluids, CO_2, H_2O and air. Statistically significant information for timing is obtained by considering several different zones in the p–T domain for each fluid. For each of the three fluids we start by showing the density contours on the pressure–temperature plane with the several zones used for reconstruction identified on the plane. The reconstruction regions chosen include some that exclude the liquid–vapor line as well as some that cross it. For these two-phase regions, we also compare thermodynamic procedures that treat the two phases as a single fluid and procedures that treat them as two distinct fluids. Following the definition of these various zones, we show the tree-structure grid for reconstruction to both 1% and 0.1% accuracy. Finally we compare the time required for the table look-up procedure and the complete REFPROP solution.

Fig. 10 shows the density of CO_2 as a function of pressure and temperature as obtained from REFPROP. The plot clearly shows the liquid vapor line and the global manner in which the density varies. The four colored squares represent four rectangular zones for which reconstruction databases have been computed. Zone 1 comprises the temperature and pressure ranges, $500\ \mathrm{K} \leqslant T \leqslant 1600\ \mathrm{K}$, $0.01\ \mathrm{MPa} \leqslant p \leqslant 1000\ \mathrm{MPa}$ and lies entirely within the vapor region. The grid structures for this zone are shown in Fig. 11 for accuracies of 1% and 0.1%. The grid color is keyed to the magnitude of the density. The 1% accuracy case requires a total of ten levels of refinement and results in 22,000 reconstruction points. The 0.1% case requires 11 refinement levels and 225,000 points. Similar results are obtained for Zones 2 and 3. The table fit and reconstruction are done on the basis of the logarithm of the pressure. The size of the reconstruction maps is summarized in Table 1.

Zones 3 and 4 of the CO_2 map cross the liquid vapor line where the density is discontinuous. Maps with accuracies of 1% and 0.1% are shown in Fig. 12. The presence of the discontinuity implies that the refinement criterion can never be satisfied at the liquid–vapor line. We have arbitrarily terminated the refinement process for this case at 15 levels. This implies that the errors will be less than the stated values at all locations except those immediately adjacent to the liquid–vapor line.

An alternative reconstruction map for the CO_2, Zone 4 domain, is given in Fig. 13 but with the vapor and liquid regions broken into two separate curve fits. Because it is no longer necessary to attempt to interpolate across the discontinuity, the number of reconstruction cells for a given accuracy is reduced dramatically. At

the 1% accuracy level, the number of cells reduces from 22,000 for the continuous fit across the discontinuity to 5336 points in the vapor and 225 points in the liquid. Thus a total of 5561 points (and only ten levels of refinement) give the same (or slightly improved) accuracy in the 'two-phase' description as in the 'single-phase'

Table 1
Storage and timing comparisons for properties evaluation

Fluid	Zone	Max Err (%)	No. Levels	No. Points	Time (s) Cartes	Time (s) EEM[a]	Time (s) REFPROP	Time ratio
CO_2	1	1	10	22,246	0.438	0.375	77.686	177.4
		0.1	11	225,121	0.438	0.375	77.297	176.5
	2	1	17	25,990	0.422	0.375	76.532	181.4
		0.1	20	257,841	0.422	0.375	76.250	180.7
	3	1	15[b]	159,325	0.422	0.375	73.022	173.1
		0.1	15[b]	801,151	0.438	0.375	73.250	167.2
	4	1	15[b]	72,703	0.422	0.375	87.329	207.1
		0.1	15[b]	108,384	0.422	0.375	87.563	207.5
	4 (Liquid)	1	5	86				
		0.1	7	1102				
	4 (Vapor)	1	10	5336				
		0.1	12	52,296				
H_2O	1	1	10	28,391	0.438	0.375	36.234	82.7
		0.1	12	225,121	0.438	0.375	36.343	83.0
	2	1	9	871	0.406	0.375	186.107	458.4
		0.1	11	8049	0.422	0.375	186.796	442.6
Air	1	1	11	63,740	0.422	0.375	202.672	480.3
		0.1	12	793,657	0.453	0.375	202.203	446.4
	2	1	8	459	0.391	0.375	989.203	2530.0
		0.1	10	4184	0.406	0.375	989.532	2437.3

[a] Equivalent equal-sized mesh (a uniform mesh refined to the deepest level).
[b] The lowest level allowed (for cases where discontinuity presents).

Table 2
Column information in Table 1

Column	Heading	Information
1	Fluid	Fluid type
2	Zone	Zone
3	Max Err	Maximum error of curve fit
4	No. Levels	Number of refinement levels needed (or allowed in cases with property discontinuity)
5	No. Points	Number of points in the adaptive database
6	Cartes Time	Time required for 100,000 property evaluations using adaptive database
7	Time (s) EEM	Time required for 100,000 property evaluations using equally spaced database.
8	REFPROP Time	Time required for 100,000 property evaluations using REFPROP
9	Time Ratio	Ratio of REFPROP to adaptive database property evaluation time

5.1. CPU comparisons with C^0 reconstruction method

An assessment of the property evaluation times required for the adaptive reconstruction procedure as compared to that for the original REFPROP routine is given in the bar chart in Fig. 16. Timing comparisons are shown for four zones of carbon dioxide, two zones of water and two zones of air. The evaluation times are plotted on a logarithmic scale with three bars given for each fluid zone. The bars represent, respectively, the time for reconstruction accuracies of 1% and 0.1% along with the time for the REFPROP evaluations. All timings are based on making 100,000 property evaluations at equally spaced points on diagonal lines across the respective fluid zones (as indicated by the dashed lines on Fig. 11 for CO_2). The reconstruction timings are for the C^0 method. The added cost for the C^1 and C^2 methods is indicated later.

The timing evaluations in Fig. 16 immediately show the advantage of the reconstruction procedure. In all cases, reconstruction is at least two orders of magnitude faster than the complete REFPROP routines. The evaluation times for reconstruction are essentially the same for all fluid types, all fluid regions, and for the two accuracy levels. This insensitivity of evaluation time to accuracy level (table size) is a clear indication

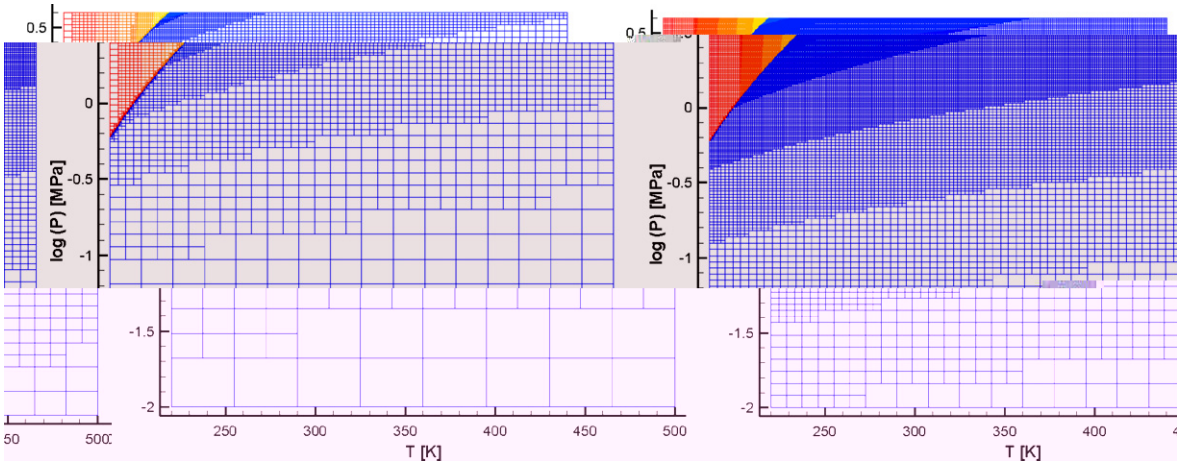

Fig. 12. Adaptive reconstruction maps for Zone 4 (Fig. 10) of CO_2. Left: 1% accuracy, 72,703 points; Right: 0.1% accuracy, 108,384 points. Domain location: 220 K $\leqslant T \leqslant$ 500 K; 0.01 MPa $\leqslant p \leqslant$ 4 MPa.

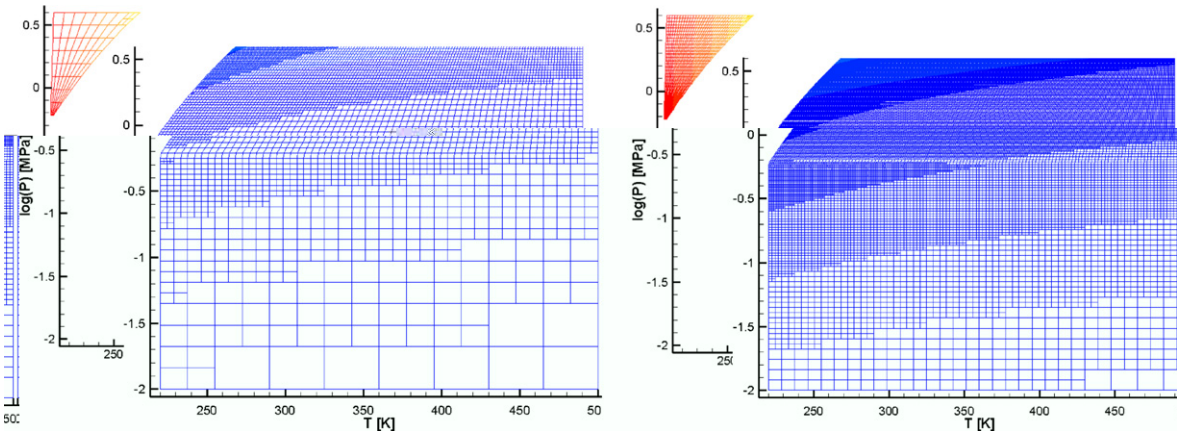

Fig. 13. Adaptive interpolation maps for Zone 4 of CO_2 (see Fig. 10). Left: 1% accuracy, 225 points in liquid zone, 5336 points in vapor zone; Right: 0.1% accuracy, 1102 points in liquid zone, 52,296 points in vapor zone. Domain location: 220 K $\leqslant T \leqslant$ 500 K; 0.01 MPa $\leqslant p \leqslant$ 4 MPa.

of the effectiveness of the tree structure in the reconstruction table. By contrast, the evaluation times for the REFPROP routines vary considerably with fluid type and fluid region because the complexity of the underlying equations changes. Air, which is treated as a mixture, is considerably more expensive to evaluate than the pure fluids. Because of the timing variability in REFPROP, the smallest time ratio occurs in Zone 1 for H_2O where reconstruction is 85 times faster while the maximum advantage occurs for air where reconstruction is as much as 2500 times faster. The savings achieved by the reconstruction procedure would clearly be smaller for equations of state based upon simpler algebraic equations, but even for the simplest equations, the procedure remains competitive.

The timing results in Fig. 16 are also tabulated in Table 1 along with the number of refinement levels and the total number of cells in each zone. For comparative purposes, the times required for making 100,000 property evaluations in an equally spaced table are also given to indicate the penalty incurred by the tree-structure. The cost of a tree-based reconstruction is approximately 12% larger than that for the equally spaced database. A numerical value of the ratio of the adaptive reconstruction time to the REFPROP evaluation time is also included. The columns in Table 1 are defined in Table 2.

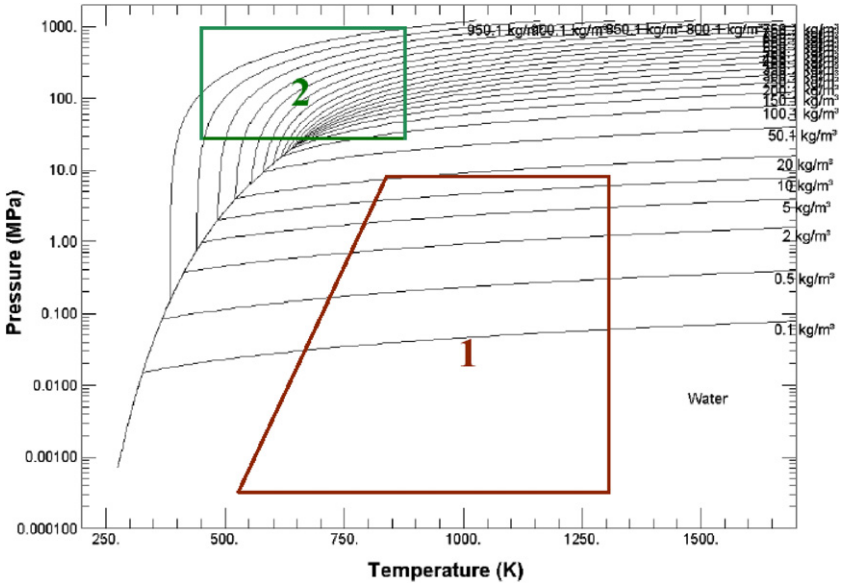

Fig. 14. Density of H_2O as a function of pressure and temperature (taken from REFPROP). Interpolation databases have been obtained for the two zones indicated.

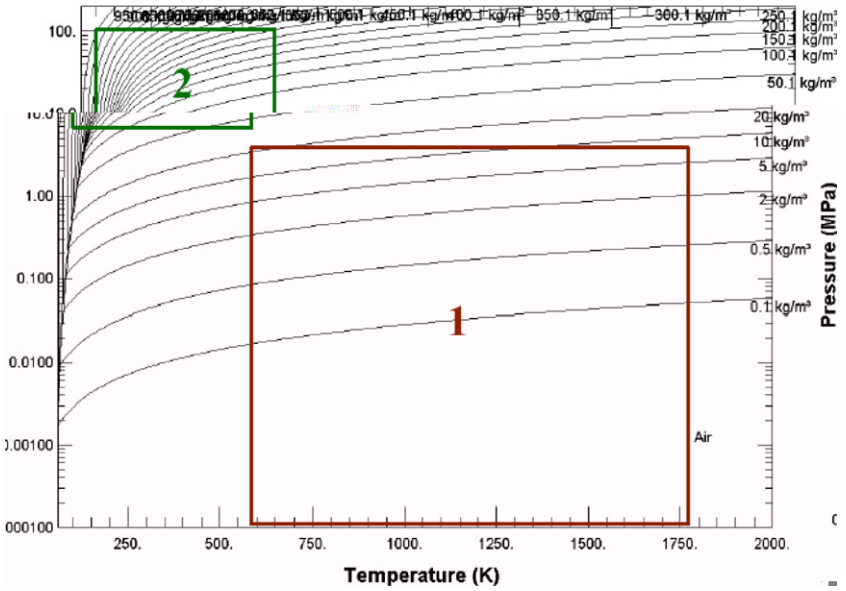

Fig. 15. Density of air as a function of pressure and temperature (taken from REFPROP). Interpolation databases have been obtained for the two zones indicated.

One caution in interpreting these comparisons is that REFPROP uses the Helmholtz function which expresses fluid properties as functions of temperature and density as the fundamental thermodynamic variable. Property evaluations in REFPROP based upon pressure and temperature will therefore incur some penalty. To demonstrate that this is not the major reason for the advantage of the adaptive method, the CO_2 properties evaluation in Zone 2 was chosen as a representative comparison. The time required for 100,000 property evaluations at equally spaced points based upon pressure and temperature was 76.5 s as noted in Table 1. Similar evaluations, based upon density and temperature required 45.1 s, a savings of 42% over

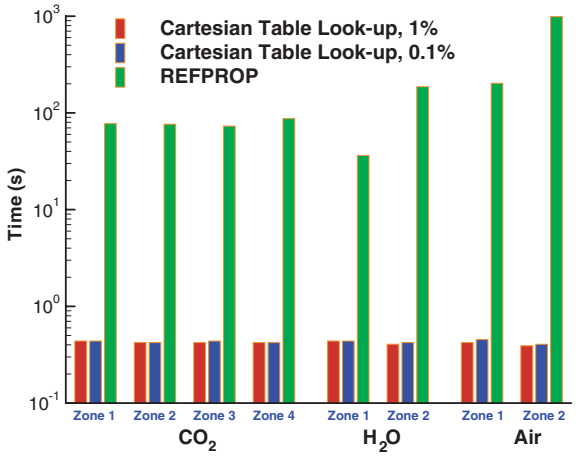

Fig. 16. Comparison of CPU time of adaptive Cartesian property reconstruction and REFPROP calculation.

the pressure and temperature pair. Therefore for CFD codes that are based upon density and temperature, the CPU advantages of the adaptive Cartesian method would be reduced by approximately a factor of two, still providing a major savings. For codes based upon other variables pairs, the speed-ups in Table 1 remain appropriate.

5.2. Table size comparison of C^0, C^1 and C^2 reconstruction methods

The above results have shown the CPU time advantages of the adaptive reconstruction method. In the present subsection, we compare the table sizes required for the C^0, C^1 and C^2 reconstruction procedures. In general, the storage requirements reflect the competing effect of the complexity of the interpolant and the number of grid points needed for accuracy. For instance the C^1 and C^2 methods require the storage of more values per grid point, but need fewer grid points due to the increased accuracy of the interpolating functions. The comparisons are made for the rectangular temperature-pressure region, $500\ \text{K} \leqslant T \leqslant 1000\ \text{K}$, $1\ \text{MPa} \leqslant p \leqslant 100\ \text{MPa}$, in CO_2 using an error tolerance of 1%. In the lower pressure portion of this region, CO_2 behaves as a perfect gas, while in the upper regimes it is strongly supercritical (see Fig. 10). The sizes of the storage tables for the three interpolation levels are given in Fig. 17. In computing these tables, the Gibbs function and its first two derivatives were obtained from analytical expressions within REFPROP but the high order derivatives (greater than two) were obtained by numerical differentiation. The number of cells required for the C^0 reconstruction method was 1290, while the more accurate C^1 interpolation method

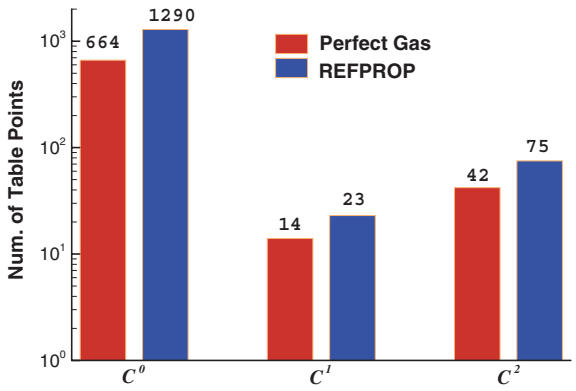

Fig. 17. Storage Comparison of different reconstruction methods. CO_2, 1% accuracy. Domain location: $500\ \text{K} \leqslant T \leqslant 1000\ \text{K}$; $1\ \text{MPa} \leqslant p \leqslant 100\ \text{MPa}$.

required only 23 cells to obtain the same accuracy. The C^2 reconstruction which uses even higher order interpolation however, required 75 cells suggesting a sensitivity in the C^2 reconstruction procedure to the high order derivatives. This phenomenon may arise from two issues. First, the Gibbs function, rather than the density and enthalpy themselves, is fitted in the C^2 reconstruction and the density and enthalpy are obtained by differentiating the Gibbs function twice. Second, the C^2 reconstruction uses higher order polynomials that are more prone to oscillate and this tendency may require increased resolution to realize a given accuracy.

To demonstrate that this trend in the storage size is not related to the method in which the REFPROP evaluations are made, we have added in Fig. 17 companion reconstruction results for a perfect gas using the same pressure–temperature region in CO_2. Here, all partial derivatives were computed analytically and the specific heats were taken as constant. The perfect gas tables for C^0, C^1 and C^2 reconstruction have 664, 14 and 42 mesh points respectively. Although the number of coefficients stored at each triangle of the table increases from 24 in the C^0 reconstruction to 48 in the C^1 reconstruction and to 61 in the C^2 reconstruction, the increment obviously is not enough to offset the savings here. This indicates that a significant reduction in database size can be achieved when moving from C^0 reconstruction to its C^1 counterpart. Going from C^1 reconstruction to C^2 reconstruction will result in a modest increase of the database size.

6. Sample CFD calculations

In the present section we present some representative computational results to demonstrate the method. We begin with a one-dimensional computation that is used to verify the method and to provide timings. We then present a two-dimensional example that demonstrates some of the thermodynamic characteristics that are enabled by the present formulation. These computations are based on the results of an in-house CFD code [25,26]. The code uses a second-order accurate, approximate Riemann solver in space with a two equation turbulence model [27,28]. A line relaxation iteration is used to solve the linear equation system. The focus in these examples is on demonstrating the efficiency and practicality of real fluid computations.

6.1. One-dimensional flow computation

As an initial example of the timing realized with the adaptive reconstruction method, we use a simple one-dimensional example of inviscid subsonic flow through a convergent–divergent nozzle. The nozzle is symmetric about the throat with the inlet and exit areas 1.25 times the throat area. The working fluid is chosen as water with an inlet total pressure of 60 MPa and stagnation temperature of 750 K. The back pressure at the outlet is 50 MPa. These parameters place the solution in Zone 2 of the H_2O property map on Fig. 14. A total of 200 cells are used in the calculation. The convergence rates for this case using the C^0, C^1 and C^2 methods were originally shown in Fig. 7 where we noted that the convergence with all three of these reconstruction methods was identical to the convergence obtained by coupling the REFPROP routines directly into the CFD code, so long as the triangulation depicted in Fig. 8 is used. Correspondingly we showed that convergence on the triangulation of Fig. 4 stalled. Although the number of iterations is essentially identical for the different property evaluation methods, their CPU costs are significantly different.

Fig. 18 shows the cumulative CPU time required for the computation as a function of iteration number for 1000 iterations. The C^0 reconstruction method results in the fastest execution. Comparison of the results for the inconsistent and the consistent triangulation methods indicates that the subdivided triangular mesh increases the CPU time by about 5% as compared to the diagonally divided triangles, indicating a minor overhead for locating a particular triangle inside each square. The costs of the C^1 and C^2 methods are nearly equal and are approximately twice that of the C^0 method. The C^2 method evaluates all properties from a single ninth-order polynomial while the C^1 method evaluates properties from separate fifth-order polynomials for density and enthalpy. The calculation based on the exact property evaluation from REFPROP is approximately 160 times slower than the C^0 method and 80 times slower than the C^1 and C^2 methods, again in keeping with the timing of the properties evaluation themselves.

The CPU time for a CFD calculation can be broken into two parts: the property evaluation time and the equation solution time. The relative advantages of faster property evaluation times clearly will decrease as the

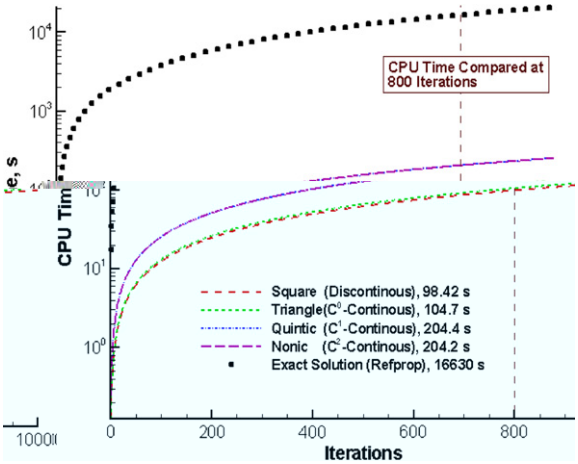

Fig. 18. CPU time comparison for different interpolation methods.

complexity of the equation solution increases, and in particular as we move from one to three dimensions. The nominal cost of the property evaluation by REFPROP in this region of the H_2O map is about 450 times that of the C^0 integration method (see Table 1). For this one-dimensional case, the CFD calculation time ratio is 160, suggesting that the ratio of equation solution time to property evaluation time in a one-dimensional calculation is approximately 1.8. Numerical computations show that the one-dimensional solver requires about 61.25 μs/(cell × iteration). Corresponding costs for two- and three-dimensional solutions are 100.2 and 200.9 μs/(cell × iteration) respectively. Assuming the property evaluation time does not change when going from one to three dimensions, this suggests that a two-dimensional calculation with C^0 interpolation method will be 115 times faster than the direct REFPROP evaluation calculation and the three-dimensional calculation will be 66 times faster. Solutions based upon C^1 and C^2 reconstruction will be about half this amount.

6.2. Two-dimensional real-fluid example

The geometrical configuration for the 2D viscous computation is again a C–D nozzle patterned after one designed for hypersonic flow testing at ultra-high pressures [29]. The working fluid is air with upstream stagnation conditions of 1700 MPa and 750 K. The density of air at these conditions is approximately equal to that of water and the resulting acceleration through the choked throat shows dramatic real-fluid characteristics. The convergent section of the nozzle is very strongly converging and the inflow starts at low speeds.

We limit our results here to a Mollier chart for air showing the physical regime of interest and comparisons between solutions of the real-fluid and corresponding perfect gas computations at the same pressure and temperature. The Mollier diagram is shown on the left half of Fig. 19 along with the h–s domain covered by the flow in the nozzle. Specifically, the enthalpy and entropy values in all cells in the computational domain are plotted on the h–s diagram superimposed upon the thermal database from REFPROP. The solution forms a roughly triangular region. The left edge of this triangular region corresponds to the thermodynamic path of the fluid on the centerline. As can be seen, the fluid on the centerline undergoes an approximately constant entropy process. The upper edge of the triangle corresponds to the fluid adjacent to the wall and shows its entropy continues to increase while the enthalpy decreases slightly in agreement with classical results for the adiabatic recovery temperature on a nozzle wall. The horizontal constant pressure lines to the right side of the Mollier diagram correspond to the perfect gas region where enthalpy is independent of pressure. The nearly vertical constant pressure lines on the left side indicate a strong real-fluid effect of pressure on enthalpy. These lead to a very different temperature pattern in the nozzle as the contours in the right half of Fig. 19 show.

The temperature distribution in the nozzle for a perfect gas calculation is compared with that for the real-fluid calculation on the right half of Fig. 19. Here, the temperature in every computational cell has been plotted

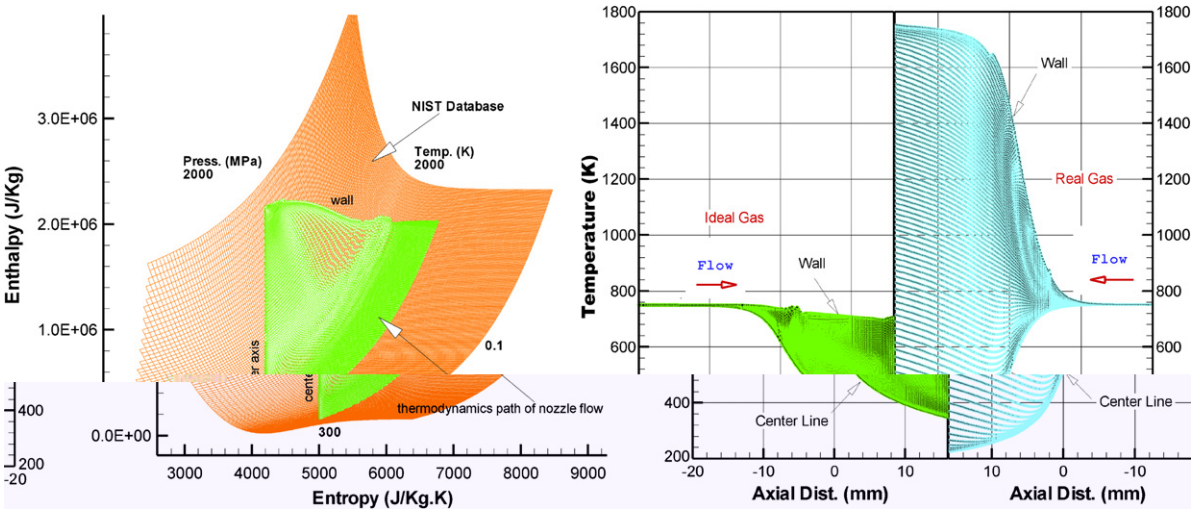

Fig. 19. 2D calculation of ultra-high pressure flow through hypersonic nozzle. Left: Mollier diagram and h–s domain of solution; Right: Temperature along the nozzle for calculation with real fluid properties and perfect gas assumption.

at its appropriate axial location in the nozzle. The left side of the plot shows the temperature for the perfect gas solution with the axial coordinate running from left to right. The right side of the plot shows the results for the real-fluid case with the axial coordinate running from right to left. There is clearly a major difference between the two solutions. In both solutions, the upper bound on the temperature at any axial location corresponds to the wall, while the minimum temperature represents the value in the free stream. For the perfect gas, the temperature on the wall (the upper bound) decreases slightly from its upstream value of 750 K, while for the real fluid, the wall temperature rises rapidly throughout, reaching over 1400 K at the throat ($x = 0$) and continuing to above 1800 K. Contrasting the increased wall temperature with the nearly constant stagnation enthalpy for the real fluid calculation in the Mollier diagram on the left plot of Fig. 19 demonstrates that the wall enthalpy for the real fluid behaves like that for the perfect gas in that it remains essentially constant throughout the nozzle. The wall temperature for the real-fluid solution, however, increases rapidly because of the strong, real-fluid, pressure dependence of enthalpy. The minimum temperature at any axial location is also lower for the real fluid than for the perfect gas, again indicating how much the real-fluid thermodynamics differ from the perfect gas assumption. As a final point, we note that both the convergence rates and the CPU times for the perfect gas and real fluid computations were nearly identical, thus verifying that the reconstruction procedure provides an effective thermodynamic interface for real-fluid computations.

7. Summary and conclusions

An adaptive reconstruction method for fluid dynamics computations of complex fluids with general equations of state is presented. The technique is based upon a Cartesian-grid approach that uses a binary tree data structure to store property information in an unequally spaced table whose resolution is automatically chosen to provide user-specified accuracy. The efficiency and accuracy of the method are assessed by comparing with the thermodynamics properties obtained from the REFPROP database. REFPROP constitutes one of the most accurate, detailed sources of thermodynamics properties for a large number of fluids. The resulting property reconstructions are nominally two orders of magnitude faster than property evaluations from the original database. This properties evaluation advantage translates into solution time enhancements of approximately one to two orders of magnitude for 3D CFD computations. The adaptive data structure allows resolutions that would be impractical for equally spaced tables of similar accuracy.

Reconstructions in regions where the adaptive procedure leads to adjacent cell of different sizes are shown to create discontinuities unless the local squares are broken into multiple (more than two) triangles. Simple diagonal bisection of squares into triangles works effectively when adjacent squares are the same size. When

appropriate triangulation procedures for unequally spaced meshes are used, reconstructions providing C^0, C^1 and C^2 function continuity are equally effective. Inconsistent properties obtained from the C^0 method appear not to have a detrimental effect on CFD solutions, but the C^1 or C^2 methods can be used as a check to verify that inconsistency is not an issue. The cost of C^0 reconstruction on an adapted Cartesian grid is about 10% larger than that on a uniform grid. This overhead arises from the triangulation inside a square, while the searching time is negligible, making property reconstruction nearly independent of table size. The semi-consistent C^1 reconstruction method is the most economical in terms of storage and provides a substantial savings over the C^0 method for the same accuracy. The C^2 method requires slightly larger storage size than C^1.

The convergence rate in example calculations is identical for the three reconstruction methods and for direct property evaluations from REFPROP, but the CPU costs are approximately two orders of magnitude smaller. The CPU cost of the C^0 reconstruction method is approximately half that of the C^1 method, while the C^1 and C^2 methods are about the same because the C^1 method constructs two fifth-order polynomials as opposed to one ninth-order polynomial in the C^2 method. Solutions based upon real gas properties are provided as evidence of the practicality of the method.

Acknowledgement

This work was supported by US Air Force under contract FA 9550-04-1-0089.

References

[1] O. Redlich, J.N.S. Kwong, On the thermodynamics of solutions. V. An equation of state. fugacities of gaseous solutions, Chem. Rev. 44 (1949) 233.
[2] G. Soave, Equilibrium constants from a modified Redlich–Kwong equation of state, Chem. Eng. Sci. 27 (1972) 1197.
[3] D.Y. Peng, D.B. Robinson, Ind. Eng. Chem. Fundam. 15 (1) (1976) 59.
[4] E.W. Lemmon, R.T. Jacobsen, S.G. Penoncello, D.G. Friend, Thermodynamic poperites of air and mistures of nitrogen, argon and oxygen from 60 to 2000 K at pressures to 2000 Mpa, J. Phys. Chem. Ref. Data 29 (3) 331-385.
[5] NIST Standard Reference Database 23, NIST Thermodynamic properties of refrigerant mixtures database (REFPROP), Version 4.0, Gaithersburg, MD, (1993).
[6] M.L. Huber, H.J.M. Hanley, The corresponding-states principle: dense fluids Chapter 12, in: J. Millat, J.H. Dymond, C.A. Nieto de Castro (Eds.), Transport Properties of Fluids, Cambridge University Press, 1996.
[7] M.L. Huber, NIST Thermophysics Properties of Hydrocarbon Mixtures Database (SUPERTRAPP), Version 3.1 User's Guide. U.S. Department of Commerce, Gaithersburg, MD, Feb., 2003.
[8] http://www.llnl.gov/science_on_lasers/09MatProps/MP-0_main.html.
[9] J.R. Edwards, R.K. Franklin, M.-S. Liou, Low-diffusion flux-splitting methods for real fluid flows at all speeds, in: 14th AIAA Computational Fluid Dynamics Conference, Norfolk, VA, June 28–July 1 1999, AIAA 1999-3327.
[10] H. Meng, V. Yang, A unified treatment of general fluid thermodynamics and its application to a preconditioning scheme, J. Comput. Phys. 189 (2003) 277–304.
[11] G. Zilliac, M.A. Karabeyolu, Modeling of propellant tank pressurization, in: 41st AIAA /ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, Tucson, Arizona, 10–13 July 2005, AIAA 2005-3549.
[12] W.J. Coirier, Efficient real gas Navier–Stokes computataions of high speed flows using an LU Scheme, in: 28th AIAA Aerospace Sciences Meeting, Reno, Nevada, Jan 8–11, 1990, AIAA, pp. 1990-0391.
[13] A. Hosangadi, V. Ahuja, A generalized multi-phase framework for modeling cavitation in cryogenic fluids, in: 33rd AIAA Fluid Dynamics Conference and Exhibit, Orlando, Florida, 23–26 June 2003, AIAA 2003-4000.
[14] R.L. Davis, B.T. Campbell, Quasi-2D Unsteady Flow Procedure for Real Fluids, in: 36th AIAA Fluid Dynamics Conference and Exhibit, San Francisco, California, 5–8 June 2006, AIAA 2006-3911.
[15] G. Cheng, R. Farmer, Development of Linearized Real-Fluid Model in Simulating Spray Combustion Flows, in: 43rd AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada, 10–13 January 2005, AIAA 2005-735.
[16] R.B. Pember, J.B. Bell, P. Colella, W.Y. Crutcheld, M.L. Welcome, An adaptive cartesian grid method for unsteady compressible flow in irregular regions, J. Comput. Phys. 120 (1995) 278–304.
[17] M.J. Aftosmis, J.E. Melton, M.J. Berger, Adaptation and surface modeling for Cartesian mesh methods, in: 12th AIAA Computational Fluid Dynamics Conference, San Diego, CA, June 19–22 1995, AIAA-95-1725-CP.
[18] E. Horowitz, S. Sahni, D. Mehta, Fundamentals of data structures in C++, Computer Science Press, 1995.
[19] W. Gordon, C. Hall, Transfinite element methods: blending-function interpolation over arbitrary curved element domains, Numer. Math. 21 (1) (1973) 109–129.
[20] C.L. Merkle, J.Y. Sullivan, P.E.O. Buelow, S. Venkateswaran, Computational of flows with arbitrary equation of state, AIAA J. 36 (4) (1998) 515–521.

[21] D. Li, V. Sankaran, J. Lindau, C.L. Merkle, A unified computational formulation for multi-component and multi-phase flows, in: 43rd AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada, January 10–13 2005, AIAA 2005-1391.
[22] F.D. Swesty, Thermodynamically consistent interpolation for equation of state tables, J. Comput. Phys. 127 (1996) 118–127.
[23] H. Akima, A method of bivariate interpolation and smooth surface fitting for irregularly distributed data points, ACM Trans. Math. Softw. 4 (2) (1978) 148–159, June.
[24] A. Preusser, Efficient formulation of a bivariate Nonic C2-Hermite polynomial on triangles, ACM Trans. Math. Softw. 16 (3) (1990) 246–252, September.
[25] D. Li, G. Xia, C.L. Merkle, Analysis of real fluid flows in converging diverging nozzles, in: 33rd AIAA Fluid Dynamics Conference and Exhibit, Orlando, Florida, June 23–26 2003, AIAA 2003-4132.
[26] D. Li, G. Xia, C.L. Merkle, Large-scale multidisciplinary computational physics simulations using parallel multi-physics zone methods, in: The International Conference on Parallel Computational Fluid Dynamics, University of Maryland, College Park Campus, May 24–27, 2005.
[27] D.C. Wilcox, Turbulence Modeling for CFD, 2nd Ed., DCW Industries, La Canada, CA, 1998.
[28] P. Cinnella, Roe-type schemes for dense gas flow computations, Comput. Fluids 35 (10) (2006) 1264–1281.
[29] I.G. Girgis, G.L. Brown, R.B. Miles, R.J. Lipinski, Fluid Mechanics of a Mach 7-12 Electron Beam Driven Missile Scale Hypersonic Wind Tunnel: Modeling and Predictions, in: 31st AIAA Fluid Dynamics Conference and Exhibit, Anaheim, CA, June 11–14, 2001, AIAA 2001-2777.