JOURNAL OF COMPUTATIONAL PHYSICS **134,** 150–168 (1997)
ARTICLE NO. CP975680

PDF Model Calculations of Compressible Turbulent Flows Using Smoothed Particle Hydrodynamics

Walter C. Welton and Stephen B. Pope

Department of Mechanical & Aerospace Engineering, Cornell University, Ithaca, New York 14850
E-mail: welton@mae.cornell.edu

Received February 26, 1996; revised February 20, 1997

A particle method which applies the probability density function (PDF) method to compressible turbulent flows is presented. Solution of the PDF equation is achieved using a Lagrangian/Monte Carlo approach. A unique feature of the method is its ability to calculate the mean pressure gradient directly from the particles using a grid-free approach. This is accomplished by applying techniques borrowed from the field of smoothed particle hydrodynamics. Furthermore, these techniques have been implemented using a recently discovered algorithm which greatly reduces the computational work in 1D. The particle method also incorporates a variance-reduction technique which can significantly reduce statistical error in first and second moments of selected mean flow quantities. When combined with a second-order accurate predictor/corrector scheme, the resulting particle method provides a feasible way to obtain accurate PDF solutions to compressible turbulent flow problems. Results have been obtained for a variety of quasi-1D flows to demonstrate the method's robustness. These include solutions to both statistically stationary and nonstationary problems, and use both periodic and characteristic-based inflow/outflow boundary conditions. Convergence of the method with respect to four different kinds of numerical errors has also been studied. Detailed results are presented which confirm the expected convergence behavior of each error.

1. INTRODUCTION

Accurate prediction of the properties of complex turbulent reacting flows is an important but challenging problem. For these types of flows probability density function (PDF) methods offer great potential [1]. Compared to conventional turbulence models, PDF methods have the primary advantage of being able to treat the important processes of convection and reaction exactly [2]; gradient-diffusion models for turbulent transport are unnecessary, plus complete treatment of arbitrarily complex finite-rate reactions is possible. PDF methods also offer significant advantages from a modeling standpoint [3].

Significant progress with PDF methods has been made in the last decade, including derivation of an evolution equation for the joint-pdf of velocity and composition [2, 4], and relatively recent development of a model for the velocity-dissipation joint-pdf [5, 6]. Solutions of these equations have been successfully obtained using Monte Carlo algorithms in which the fluid is represented by a large set of stochastic particles having properties that evolve in time. The majority of applications, however, are limited to incompressible flows, and only recently has work been done in extending and applying the method to compressible flows having pressure induced density variations [7, 8]. Related to this is the need to develop a general method for calculating the mean pressure field within the Monte Carlo algorithm. A few approaches exist, but these are limited in scope: for thin shear flows the mean pressure has been calculated by invoking boundary-layer approximations [9], and for statistically stationary incompressible flows the pressure has been obtained from a Poisson equation [10]. Recently, a pressure-correction algorithm has been devised and implemented which is applicable to either constant density flows or statistically stationary variable density flows [11]. A different approach is to couple the Monte Carlo code to a finite-volume code which calculates the mean pressure [8, 12–14]. While this approach can combine advantages of both methods, the coupling becomes complex for compressible reacting flows and this drives the need to develop a completely independent and robust Monte Carlo/PDF method.

The particle method described in this paper addresses these two areas. The method is applicable to compressible turbulent nonreacting quasi-1D flows in which changes in pressure lead to corresponding changes in density. Only isentropic flows are presently considered and, hence, flows with shocks are excluded. The flows, however, can be either statistically stationary or nonstationary, as well as laminar. In all cases, the mean pressure field and gradient are calculated directly from the particles to give a completely self-contained Monte Carlo/PDF code. Although emphasis has been placed on determination of the mean pressure field and only nonreactive flows are examined, it is expected that extension of the method to include reaction is straightforward. Extension of the method to 2D is also expected to be straightforward.

0021-9991/97 $25.00

The procedure used to obtain the mean pressure is simple. At each point in the flow the mean fluid density is proportional to the local particle number density, and the pressure is subsequently obtained from a simple isentropic equation of state. The problem of determining the local particle number density is solved by utilizing techniques borrowed from smoothed particle hydrodynamics (SPH).

SPH is a grid-free particle method that was originally developed for astrophysical problems [15] yet is also applicable to fluid flow problems encountered in industry, of which the majority are turbulent. This paper describes the first application of SPH to turbulent flows. In SPH each particle is assigned properties which evolve in time according to specified evolution equations. Appearing in these equations are coefficients which are interpolated from the moving and usually disordered set of particles [16]. In this sense SPH fits quite well into a Lagrangian Monte Carlo/PDF framework.

Some differences in SPH and Monte Carlo/PDF particle methods also exist. In all previous Monte Carlo/PDF calculations the mean fields have been represented as spline functions on a fixed grid [2], and mean properties at particle locations are interpolated from the spline functions. In SPH the mean quantities for each particle are calculated without reference to a grid. This particle method is therefore the first grid-free Monte Carlo/PDF method. While this allows direct particle–particle interactions, it results in the computational work scaling as $\mathcal{O}(N^2)$ (N being the number of particles) if implemented in a straightforward manner, whereas in all present Monte Carlo/PDF methods, the computational work is of order N. Since large numbers of particles are typically required to accurately model turbulent flows, the overall computational cost of using SPH in a Monte Carlo/PDF framework would quickly become prohibitive. But a recently developed algorithm for which the computational work scales as $\mathcal{O}(N)$ in 1D overcomes this problem, making accurate and meaningful results feasible. The algorithm can also be extended to 2D and 3D, and while the work scaling here is no longer $\mathcal{O}(N)$, it is still less than $\mathcal{O}(N^2)$.

The particle method also employs a variance-reduction technique to reduce the statistical error produced by the Monte Carlo algorithm. The technique, originally applied to 0D homogeneous flows [17], has been extended to accommodate SPH and is successfully applied to the 1D inhomogeneous flows considered here. Variance reduction is performed on first and second moments of calculated flow quantities, both on the initial condition and at each time-step.

The paper begins by briefly discussing the PDF method: the pdf applicable to this problem is introduced, and the modeled particle evolution equations along with the corresponding modeled pdf evolution equation are presented. Numerical implementation of the method is discussed in Section 3. The predictor/corrector scheme used to integrate the stochastic differential equations is described, and the basics of SPH are presented. This is followed by a description of the $\mathcal{O}(N)$ algorithm, its extension to 2D and 3D, and the variance-reduction technique. The implementation of boundary conditions is described separately in Section 4, and numerical results are presented in the final two sections. Section 5 includes solutions to some sample flow problems and a performance test of the variance-reduction technique, while Section 6 presents a detailed convergence study of the numerical errors arising in the method, and an estimate of CPU time needed to achieve an error level of 1%.

2. PDF METHOD

The particle method described in this paper is based on a Lagrangian formulation and is developed for compressible (variable density) inhomogeneous flows. The applicable pdf is therefore the Lagrangian *mass density function* (mdf) of velocity and position. Given a fluid particle with position $\mathbf{x}_0$ at a reference time t_0, its Lagrangian velocity and position at time t are defined to be $\mathbf{U}^+(t, \mathbf{x}_0)$ and $\mathbf{x}^+(t, \mathbf{x}_0)$, respectively. The Lagrangian mass density function $\mathscr{F}(\mathbf{V}, \mathbf{x}; t)$ gives the expected mass density at time t corresponding to the joint event $\mathbf{U}^+(t, \mathbf{x}_0) = \mathbf{V}$ and $\mathbf{x}^+(t, \mathbf{x}_0) = \mathbf{x}$ [2]. Two basic properties of $\mathscr{F}$ are

$$\int \mathscr{F}(\mathbf{V}, \mathbf{x}; t)\, d\mathbf{V} = \langle \rho(\mathbf{x}, t) \rangle \tag{1}$$

and

$$\int Q(\mathbf{V}, \mathbf{x}, t) \mathscr{F}(\mathbf{V}, \mathbf{x}; t)\, d\mathbf{V} = \langle \rho Q \rangle = \langle \rho \rangle \tilde{Q}, \tag{2}$$

where $\langle \rho \rangle$ is the mean fluid density, Q is a random variable, and $\tilde{Q}$ is by definition the density-weighted mean, or Favre average, of Q. Both integrals are over the velocity sample space.

In a Monte Carlo simulation of a fluid of constant mass $\mathscr{M}$, $\mathscr{F}$ is represented by an ensemble of N stochastic particles, each of mass $\Delta m = \mathscr{M}/N$, which model fluid particles. The particles, which are continuously distributed in the domain, have velocities $\{\mathbf{U}^{*(n)}\}$ and positions $\{\mathbf{x}^{*(n)}\}$, and from them the discrete Lagrangian mdf is defined as

$$\mathscr{F}_N(\mathbf{V}, \mathbf{x}; t) = \Delta m \sum_{n=1}^{N} \delta(\mathbf{V} - \mathbf{U}^{*(n)}) \delta(\mathbf{x} - \mathbf{x}^{*(n)}). \tag{3}$$

Asterisks indicate that these particles model fluid particles. The relationship between $\mathscr{F}_N$ and $\mathscr{F}$ is simply [2]

$$\langle \mathscr{F}_N(\mathbf{V}, \mathbf{x}; t) \rangle = \mathscr{F}(\mathbf{V}, \mathbf{x}; t), \quad \text{all } N \geq 1, \tag{4}$$

which establishes the validity of the particle representation. Although multitime statistics can be calculated from $\mathscr{T}$, it contains no two-point information since each particle is considered to be a sample from different and independent realizations of the flow.

2.1. *Evolution Equations*

The particle method presented here is applicable to quasi-1D flows through nozzles in which all properties are assumed to vary only in the streamwise (x_1) direction. Each particle is assigned a streamwise position x_1^* and three components of velocity $\{U_1^*, u_2^*, u_3^*\}$, where lower-case velocities denote fluctuations with respect to the Favre average. Only fluctuations in the cross-stream velocities need to be considered, since mean cross-stream velocities are zero at the centerline for quasi-1D flow. The modeled evolution equations for these properties apply at the nozzle centerline. The change in the streamwise particle position over a time interval dt is

$$dx_1^* = U_1^*\, dt. \tag{5}$$

Using the simplified Langevin model [3, 18] and taking into account area effects, the velocity evolution equations are

$$\begin{aligned} dU_1^* = &-\frac{1}{\langle\rho\rangle}\frac{\partial\langle P\rangle}{\partial x_1}\,dt - \beta\omega(U_1^* - \tilde{U}_1)\,dt \\ &+ (C_0\omega\tilde{k})^{1/2}\,dW_1(t) \end{aligned} \tag{6}$$

and

$$\begin{aligned} du_i^* = &-\beta\omega u_i^*\,dt - \left(\frac{1}{2}\frac{\tilde{U}_1}{A}\frac{dA}{dx_1}\right)u_i^*\,dt \\ &+ (C_0\omega\tilde{k})^{1/2}\,dW_i(t), \quad i = 2, 3. \end{aligned} \tag{7}$$

Appearing in the equations is the mean pressure $\langle P\rangle$, the Eulerian mean velocity $\tilde{U}_1$, the turbulent kinetic energy $\tilde{k}$, the turbulent frequency ω (defined as $\langle\varepsilon\rangle/\tilde{k}$, where $\langle\varepsilon\rangle$ is the mean dissipation), a universal constant C_0, a drift coefficient β, an isotropic vector Wiener process $\mathbf{W}(t)$, and the nozzle cross-sectional area $A(x_1)$. All means are evaluated at the particle position x_1^*. The drift coefficient β is defined to be $\frac{1}{2} + \frac{3}{4}C_0$, which causes turbulent kinetic energy to be dissipated at the rate $\langle\varepsilon\rangle$ in homogeneous isotropic turbulence.

The simplified Langevin model is a stochastic model for inhomogeneous, incompressible turbulence which, at the level of second moments, is equivalent to Rotta's model [19]. The model simulates return to isotropy, although the pressure gradient and area production terms in Eqs. (6) and (7) tend to make the turbulence anisotropic. The model also assumes high Reynolds number so that dissipation due to the mean velocities can be neglected. The value $C_0 = 2.1$ recommended for the simplified Langevin model [20] is used here. Furthermore, the turbulent frequency is assumed constant in this study. In more sophisticated models [5, 6], each particle has its own turbulent frequency which evolves according to an additional stochastic equation. Some models [4] also attempt to account for "rapid" pressure effects. Compressible turbulence has also been addressed in a recent model which solves a stochastic evolution equation for the full pressure of each particle [21]. For the purpose of testing the mean pressure algorithm, however, the simplified Langevin model is judged adequate and is therefore chosen over these more sophisticated models.

Corresponding to the modeled stochastic equations for the particles is the modeled transport equation for the Lagrangian mdf $\mathscr{T}^*$. Using standard techniques [2] the derived equation is

$$\begin{aligned} \frac{\partial\mathscr{T}^*}{\partial t} + \frac{\partial}{\partial x_1}[V_1\mathscr{T}^*] &= \frac{C_0\omega\tilde{k}}{2}\frac{\partial^2\mathscr{T}^*}{\partial V_j\,\partial V_j} \\ &+ \frac{\partial}{\partial V_1}\left[\left\{\frac{1}{\langle\rho\rangle}\frac{\partial\langle P\rangle}{\partial x_1} + \beta\omega(V_1 - \tilde{U}_1)\right\}\mathscr{T}^*\right] \\ &+ \sum_{j=2}^{3}\frac{\partial}{\partial V_j}\left[\left\{\frac{1}{2}\frac{\tilde{U}_1}{A}\frac{dA}{dx_1} + \beta\omega\right\}V_j\mathscr{T}^*\right]. \end{aligned} \tag{8}$$

When compared to the exact evolution equation for $\mathscr{T}$ obtained from the Navier–Stokes equations (for quasi-1D flow), it may be seen that modeled terms include the contributions due to fluctuating pressure gradient and viscous dissipation.

3. NUMERICAL IMPLEMENTATION

3.1. *Predictor/Corrector Scheme*

The stochastic system of particle evolution equations (5)–(7) is integrated forward in time using a second-order accurate weak predictor/corrector (P/C) scheme. For a general diffusion process of the form

$$d\mathbf{x}(t) = \mathbf{U}(t)\,dt, \tag{9}$$

$$d\mathbf{U}(t) = \mathbf{a}(\mathbf{x}[t], \mathbf{U}[t], t)\,dt + b(\mathbf{x}[t], t)\,d\mathbf{W}(t), \tag{10}$$

with drift and diffusion coefficients $\mathbf{a}$ and b, respectively, the predictor step has the explicit form

$$\hat{\mathbf{x}}^{k+1} = \mathbf{x}^k + \mathbf{U}^k\,\Delta t, \tag{11}$$

$$\hat{\mathbf{U}}^{k+1} = \mathbf{U}^k + \mathbf{a}^k\,\Delta t + b^k\,\Delta t^{1/2}\boldsymbol{\xi}^k, \tag{12}$$

and the corrector step has the averaged form

$$\mathbf{x}^{k+1} = \mathbf{x}^k + \frac{(\mathbf{U}^k + \hat{\mathbf{U}}^{k+1})}{2}\Delta t, \tag{13}$$

$$\mathbf{U}^{k+1} = \mathbf{U}^k + \frac{(\mathbf{a}^k + \hat{\mathbf{a}}^{k+1})}{2}\Delta t + \frac{(b^k + \hat{b}^{k+1})}{2}\Delta t^{1/2}\,\boldsymbol{\xi}^k. \tag{14}$$

Superscripts k and $k+1$ denote values at time t^k and $t^{k+1} = t^k + \Delta t$, respectively, while $\mathbf{a}^k$ and $\hat{\mathbf{a}}^{k+1}$ are written for $\mathbf{a}(\mathbf{x}^k, \mathbf{U}^k, t^k)$ and $\mathbf{a}(\hat{\mathbf{x}}^{k+1}, \hat{\mathbf{U}}^{k+1}, t^{k+1})$ (similarly for b^k and $\hat{b}^{k+1}$). The quantity ξ_i^k is a standardized random variable with zero mean and unit variance. Note that the same value of ξ_i^k is used for both the predictor and corrector steps, but that different and independent values of ξ_i^k are used for each time step. The time step in Eqs. (11)–(14) has a maximum value determined by a CFL constraint, the precise form of which is given in the next section.

It has been verified that this P/C scheme is equivalent to the Runge–Kutta form of Mil'shtein's general second-order weak scheme for stochastic systems [22, 23], and hence, second-order time accuracy is expected, in theory, for this particle method.

3.2. *SPH Method*

Coefficients in the particle evolution equations, including $\langle\rho\rangle$, $\partial\langle P\rangle/\partial x_1$, $\tilde{U}_1$, and $\tilde{k}$, are evaluated directly from the particles using the kernel estimation techniques of SPH. In 1D the integral form of the kernel estimate for the quantity Q is

$$\langle Q(x_1)\rangle_h = \int_{-\infty}^{\infty} Q(x_1')K(x_1 - x_1', h)\,dx_1', \tag{15}$$

where h, termed the "smoothing length," is a measure of the width of the interpolating kernel function K. This function has the properties

$$\int_{-\infty}^{\infty} K(x_1 - x_1', h)\,dx_1' = 1 \tag{16}$$

and

$$\lim_{h\to 0} K(x_1 - x_1', h) = \delta(x_1 - x_1'). \tag{17}$$

The integral in Eq. (15) can be estimated from the particles using the discrete form

$$\langle Q(x_1)\rangle_{N,h} = \Delta m \sum_{n=1}^{N} \frac{Q^{(n)}}{\overline{\rho}^{(n)}} K(x_1 - x_1^{*(n)}, h), \tag{18}$$

where $\overline{\rho}^{(n)}$ is the mean linear density of the nth particle. In quasi-1D the linear density is just equal to the mean fluid density multiplied by the cross-sectional area. Assuming a differentiable kernel, gradients of functions are obtained by direct differentiation of Eq. (18):

$$\frac{\partial\langle Q(x_1)\rangle_{N,h}}{\partial x_1} = \Delta m \sum_{n=1}^{N} \frac{Q^{(n)}}{\overline{\rho}^{(n)}} \frac{\partial K(x_1 - x_1^{*(n)}, h)}{\partial x_1}. \tag{19}$$

In both Eqs. (18) and (19), smaller values of h give a more local estimate to the mean, but result in fewer particles giving significant contributions and hence more statistical error.

The kernel used in this study is symmetric, piecewise-quartic, and has finite support. The latter implies that only a subset of the particles contributes to each kernel estimate. The functional form of $K(r, h)$ is

$$K(r,h) = \begin{cases} c(1 + 3|r|/h)(1 - |r|/h)^3 & \text{for } |r| \le h \\ 0 & \text{for } |r| > h, \end{cases} \tag{20}$$

where $c = 5/(4h)$ is a constant obtained from the normalization constraint Eq. (16). From Eq. (20) it is easy to show that K possesses continuous first and second derivatives.

Implementation of the method is simple. The first step is to estimate the linear density $\overline{\rho}^{(i)}$ for each particle by replacing $Q^{(n)}$ with $\overline{\rho}^{(n)}$ in Eq. (18):

$$\overline{\rho}^{(i)} = \overline{\rho}(x_1^{*(i)}) = \Delta m \sum_{n=1}^{N} K(x_1^{*(i)} - x_1^{*(n)}, h). \tag{21}$$

Having these, all other quantities can be obtained from Eqs. (18) and (19). The Favre average $\widetilde{U_1}$, for example, is obtained from

$$\widetilde{U_1}^{(i)} = \frac{\Delta m}{\overline{\rho}^{(i)}} \sum_{n=1}^{N} U_1^{*(n)} K(x_1^{*(i)} - x_1^{*(n)}, h). \tag{22}$$

The pressure gradient is obtained in a modified way. For the isentropic flows considered here, the mean pressure and density are related through the equation of state

$$\left(\frac{\langle P\rangle}{P_0}\right) = \left(\frac{\langle\rho\rangle}{\rho_0}\right)^{\gamma}, \tag{23}$$

where γ is the ratio of specific heats ($\gamma = 1.4$ in this study) and P_0 and ρ_0 are reference values. The pressure gradient can then be calculated via the relation

$$-\frac{1}{\langle\rho\rangle}\frac{\partial\langle P\rangle}{\partial x_1} = -\gamma\frac{\langle P\rangle}{\overline{\rho}^2}\left[A\frac{\partial\overline{\rho}}{\partial x_1} - \overline{\rho}\frac{\partial A}{\partial x_1}\right], \tag{24}$$

which has the advantage of having the linear density inside the gradient operator [15] as opposed to the pressure.

When SPH is combined with the previously described predictor/corrector scheme, the time step constraint has the form

$$\Delta t \le C_t \frac{h}{\max_i(|U_1^{*(i)}| + a^{(i)})}, \tag{25}$$

where $a^{(i)}$ is the local mean sound speed for the ith particle, and C_t is a constant. Numerical tests show that the upper limit on C_t for stable solutions is about 0.75.

3.3. $\mathcal{O}(N)$ Algorithm

Straightforward calculation of the $\mathcal{O}(N)$ coefficients using the discrete sums in Eqs. (18) and (19) for each of the N particles results in the computational work scaling as $\mathcal{O}(N^2)$. For kernels with compact support, this work is readily reduced to $\mathcal{O}(N\overline{N})$, where $\overline{N} = Nh/L$ is the average number of particles contributing to each kernel estimate. Since accurate modeling of turbulent flows requires that $\overline{N}$ (and hence N) be large, calculation of the coefficients using a direct summation algorithm is clearly unacceptable. This difficulty has been overcome, however, with the development of an algorithm whose computational work scales as $\mathcal{O}(N)$ in 1D.

The algorithm takes advantage of the piecewise polynomial property of the kernel to expand each kernel estimate in a finite Taylor series. To simplify notation, define

$$S^{(p)}(x_1) = \Delta m \sum_{n=1}^{N} \frac{Q^{(n)}}{\overline{\rho}^{(n)}} K^{(p)}(x_1 - x_1^{*(n)}, h) \tag{26}$$

as the pth derivative of the kernel estimate to the quantity Q at x_1. The Taylor series expansion of Eq. (26) about x_1 using the piecewise-quartic kernel is

$$\begin{aligned} S^{(p)}(x_1 + \Delta x) &= \Delta m \sum_{n=1}^{N} \frac{Q^{(n)}}{\overline{\rho}_{(n)}} \sum_{q=0}^{4-p} \frac{\Delta x^q}{q!} K^{(p+q)}(x_1 - x_1^{*(n)}, h) \\ &= \sum_{q=0}^{4-p} \frac{\Delta x^q}{q!} S^{(p+q)}(x_1), \quad p = 0 \cdots 4. \end{aligned} \tag{27}$$

This expansion is valid as long as $K(x, h)$ and its derivatives are smooth in $x \in [x_1 - x_1^{*(n)}, x_1 + \Delta x - x_1^{*(n)})$ for all n. Jumps in the expansion are caused when $x_1 + \Delta x - x_1^{*(n)}$ is at a point where the kernel has discontinuous derivatives. The effect of such jumps can be taken into account exactly though. Letting subscripts + and − denote quantities to the right and left of the discontinuity, respectively, the jump due to the nth particle is then

$$\begin{aligned} &S_+^{(p)}(x_1 + \Delta x) - S_-^{(p)}(x_1 + \Delta x) \\ &\quad = \Delta m \frac{Q^{(n)}}{\overline{\rho}^{(n)}} [K_+^{(p)}(x_1 + \Delta x - x_1^{*(n)}, h) \\ &\quad - K_-^{(p)}(x_1 + \Delta x - x_1^{*(n)}, h)]. \end{aligned} \tag{28}$$

These ideas can be implemented to calculate a set of N 1D kernel estimates with $\mathcal{O}(N)$ work. Before proceeding, however, the particles must be sorted by position, e.g., $x_1^{*(1)} \le x_1^{*(2)} \le \cdots \le x_1^{*(N)}$. Assuming the particles are continuously distributed, this sorting can also be done in $\mathcal{O}(N)$ work by "binning" the particles into M cells, where M is $\mathcal{O}(N)$, and then doing a heapsort in each bin. Since the average number of particles per bin is N/M, the total work for the sort scales as $M(N/M)\log(N/M)$, which is $\mathcal{O}(N)$.

Having the sorted set of particles, the first step is to use Eq. (26) to directly calculate $S^{(p)}(x_1)$ for $p = \{0, 1, 2, 3, 4\}$ at the first particle location $x_1^{*(1)}$. This involves $\mathcal{O}(\overline{N})$ work. The next step is to alternately use Eqs. (27) and (28) to obtain the values of $S^{(p)}$ at $x_1 = x_1^{*(2)}$: Eq. (27) is used to extrapolate the values of $S^{(p)}$ out to the next discontinuity, and Eq. (28) gives the jump across the discontinuity. Each intervening discontinuity is therefore crossed until the point $x_1^{*(2)}$ is reached. This cyclic process is continued for all remaining kernel estimates until the final particle location $x_1^{*(N)}$ is reached. The total work is $\mathcal{O}(N)$ since there are $\mathcal{O}(N)$ discontinuities and stopping points.

The extension of the algorithm to 2D and 3D is analogous to the 1D algorithm. The first step is to arrange the particles in a tree (for example, a Euclidean minimum spanning tree) which gives an efficient sequence for generating the kernel estimates. A direct sum is then used to calculate the kernel estimate and its derivatives at the root of the tree. The remaining kernel estimates are obtained by alternately extrapolating via a 2D or 3D Taylor series and adding jumps due to discontinuities. For this extension to work the kernel must also be a tensor-product kernel, $\overline{K}(x_1, x_2) = K(x_1)K(x_2)$, where $K(x)$ is again a piecewise polynomial.

In higher dimensions the computational work no longer scales purely as $\mathcal{O}(N)$. In general, the work in D dimensions can be shown to scale as

$$\begin{aligned} W &\sim N^{2-1/D}(h/L)^{D-1} \\ &\sim N\overline{N}^{1-1/D}, \end{aligned} \tag{29}$$

where $\overline{N} = N(h/L)^D$. In 2D this gives $W \sim N\overline{N}^{1/2}$ compared to $W \sim N\overline{N}$ for the direct sum implementation. From the latter form of Eq. (29), it should be clear that

the algorithm offers a computational improvement for all D in the asymptotic limit of $\overline{N}$ approaching infinity.

3.4. *Variance Reduction*

Although Monte Carlo algorithms offer a feasible way to solve evolution equations for pdfs of large dimensionality, they have the drawback of introducing statistical error into the calculations. This numerical error, which is usually dominant over other errors, can arise from many sources, including initial conditions, Wiener processes, and coefficients in the particle evolution equations. Straightforward ways of reducing statistical error include using large numbers of particles, or averaging results from multiple independent simulations. While these methods reduce statistical error in all moments, they are inefficient since reducing the error by a factor of p requires that the computational work increase by a factor of at least p^2. A different way to reduce statistical error is to incorporate a variance-reduction technique into the Monte Carlo algorithm. Variance-reduction techniques have the potential to greatly reduce (and in some cases completely remove) the statistical error in selected moments. Such techniques are therefore essential in Monte Carlo/PDF calculations.

The variance-reduction technique used in this particle method is one developed by Pope [17] and originally applied to 0D homogeneous flows, and it has been extended here to accommodate SPH and inhomogeneous 1D flows. Variance reduction is performed on all first and second moments of the particle positions and velocities, which include

$$\langle x_1^* \rangle, \langle U_i^* \rangle, \langle x_1^{*2} \rangle, \langle x_1^* u_i^* \rangle, \langle u_i^* u_j^* \rangle. \tag{30}$$

Clearly it is important to do variance reduction on the velocities, since these quantities appear as coefficients in the simplified Langevin equation. It is also important, however, to do variance reduction on the positions since they solely determine the mean pressure gradient.

The method is applied to individual ensembles each consisting of N_e particles, where N_e is a specified parameter. For each ensemble, the variance-reduction process involves two steps: obtaining ensemble estimates to these moments that are free of primary statistical error, and then adding small corrections to the particle properties so that the moments calculated from the ensemble of particles equal the statistical-error-free estimates. These two steps are now briefly described; further details may be found in the reference.

At each time step, a set of isotropic random vectors $\boldsymbol{\xi}^{(n)}$ is generated and used to advance the particle properties on both the predictor and corrector steps (Eqs. (11)–(14)). At the end of each predictor/corrector step, the quantities in Eq. (30) are estimated using ensemble averages over the set of N_e particles, and these contain error, in part because the ensemble averages $\langle \boldsymbol{\xi} \rangle_{N_e}$ and $\langle \xi_i \xi_j \rangle_{N_e}$ are not exactly zero and δ_{ij}, respectively.

To obtain the statistical-error-free estimates of the quantities in Eq. (30), four temporary samples are generated for each particle n using specially constructed random vectors $\boldsymbol{\eta}^{(n,\alpha)}$ ($\alpha = 1 \ldots 4$) in place of $\boldsymbol{\xi}^{(n)}$, and ensemble averages are then formed from these samples. For each particle, the four random vectors are geometrically defined to be the vertices of a randomly oriented regular tetrahedron lying on the sphere of radius $\sqrt{3}$ centered at the origin. The following properties then hold:

$$\langle \boldsymbol{\eta}^{(n,\alpha)} \rangle = 0, \quad \langle \eta_i^{(n,\alpha)} \eta_j^{(n,\alpha)} \rangle = \delta_{ij}, \tag{31}$$

and

$$\frac{1}{4} \sum_{\alpha=1}^{4} \boldsymbol{\eta}^{(n,\alpha)} = 0, \quad \frac{1}{4} \sum_{\alpha=1}^{4} \eta_i^{(n,\alpha)} \eta_j^{(n,\alpha)} = \delta_{ij}. \tag{32}$$

For each of the quantities in Eq. (30), these properties ensure that ensemble averages formed from these four sets of samples are free of primary statistical error. For example, consider determining $\langle U_1^* \rangle$ on the predictor step. By replacing the original random vectors $\boldsymbol{\xi}^{(n)}$ with the special random vectors $\boldsymbol{\eta}^{(n,\alpha)}$, four predictor values, $\hat{U}_1^{*(n,\alpha)}$, are generated for each particle from the streamwise velocity evolution equation, Eq. (6), via the predictor, Eq. (12). The ensemble average $\langle U_1^* \rangle_{4N_e}$, defined as

$$\langle U_1^* \rangle_{4N_e} = \frac{1}{4\boldsymbol{N}_e} \sum_{n=1}^{N_e} \sum_{\alpha=1}^{4} \hat{U}_1^{*(n,\alpha)}, \tag{33}$$

then has a primary statistical error of

$$\langle U_1^* \rangle_{4N_e} - \langle U_1^* \rangle = \frac{1}{4\boldsymbol{N}_e} \sum_{n=1}^{N_e} \sum_{\alpha=1}^{4} \eta_1^{(n,\alpha)} b^{(n)} \, \Delta t^{1/2} \tag{34}$$

which is zero because of the first property in Eq. (32). It can be similarly verified that the statistical error in the ensemble averages for the other quantities is also zero.

Having the statistical-error-free estimates, the correction to the particle properties can be easily performed. Define the vector of particle properties to be $\mathbf{q}^{(n)} = \{x_1^{*(n)}, U_1^{*(n)}, u_2^{*(n)}, u_3^{*(n)}\}^{\mathrm{T}}$; these properties are generated using the original set of random vectors $\boldsymbol{\xi}^{(n)}$ in either the predictor or corrector step, and they have ensemble mean $\mathbf{b} = \langle \mathbf{q} \rangle_{N_e}$ and variance $C_{ij} = \langle q_i q_j \rangle_{N_e}$. Both $\mathbf{b}$ and $\mathbf{C}$ contain statistical error. The correction applied to each particle in

the ensemble which yields the desired statistical-error-free mean $\overline{\mathbf{b}} = \langle \mathbf{q} \rangle_{4N_e}$ and variance $\overline{C}_{ij} = \langle q_i q_j \rangle_{4N_e}$ has the form

$$\overline{\mathbf{q}}^{(n)} = \overline{\mathbf{b}} + \mathbf{L}(\mathbf{q}^{(n)} - \mathbf{b}), \quad n = 1 \cdots N_e, \tag{35}$$

where $\mathbf{L}$ is the lower triangular matrix obtained from the matrix equation $\overline{\mathbf{C}} = \mathbf{L}\mathbf{C}\mathbf{L}^{\mathrm{T}}$. The correction to the particle velocities is of order $N_e^{-1/2}\Delta t^{1/2}$, while for the particle positions it is of order $N_e^{-1/2}\Delta t^{3/2}$. This type of correction is also performed on the initial condition, for which the desired means and variances are known explicitly.

Implementing these ideas into the current particle method is straightforward. The first step is to divide the domain into local ensembles of N_e particles each. In the 0D case for which the variance-reduction technique was originally tested, the entire set of particles was treated as a single ensemble. For the inhomogeneous 1D flows considered here, however, the technique must be applied to many ensembles. The next step is to perform the variance reduction in each ensemble. This involves calculating the statistical-error-free estimates from the temporary samples obtained using the special random variables, and then correcting the particle properties obtained using the original random variables. The final step is to calculate new coefficients for the Langevin equation using kernel estimates from the corrected set of particles. The kernel estimates can be expected to have reduced statistical error since the ensemble means and variances are statistical-error-free (this is quantitatively studied in Section 5.3). The whole process is done for the predictor step and repeated for the corrector step using the same groups of particles.

4. BOUNDARY CONDITIONS

Both periodic and characteristic-based inflow/outflow boundary conditions are used in this study, and this section describes their implementation within the framework of this particle method. In finite-volume/finite-difference methods boundary conditions are applied at grid points on or near boundaries. Particle methods, however, require that boundary conditions also be applied on particles. The particle method of this study is grid-free, and hence all boundary conditions are implemented solely through the particles.

4.1. *Periodic BCs*

Periodic boundary conditions have the advantage of being simple to implement. Their effect on the accuracy and stability of the numerical method is also minimal; the convergence study in Section 6 uses periodic boundary conditions throughout for this reason. Implementation has two parts: (1) a particle which exits either side of the periodic duct automatically reenters on the opposite side with the same properties, and (2) kernel estimates are calculated using a periodically extended kernel. Section 5.1 presents a sample flow calculation using periodic boundary conditions.

4.2. *Inflow/Outflow BCs*

Inflow/outflow boundary conditions allow simulation of more realistic flows such as the configuration shown in Fig. 1. The setup consists of two large reservoirs connected by a convergent–divergent nozzle. One reservoir contains fluid at stagnation conditions, while the other contains fluid held at a lower back pressure P_B, traditionally normalized by the upstream stagnation pressure P_o. The pressure difference between the reservoirs causes fluid to flow through the nozzle. Turbulence is introduced into the problem by positioning an idealized ''turbulence generator'' just upstream of the nozzle test section (that portion of the nozzle being simulated using the particle method). In an actual experiment the generator might consist of a wire grid or an array of fans.

It is important to relate flow quantities across the generator to each other and to the upstream stagnation conditions. Denote quantities just upstream and downstream of the generator with subscripts a and b, respectively. The inviscid mass and momentum jump conditions across the generator are

$$\text{continuity:} \quad \rho_a \widetilde{U}_{1a} A = \rho_b \widetilde{U}_{1b} A, \tag{36}$$

$$\text{momentum:} \quad P_a + \rho_a \widetilde{U}_{1a}^2 = P_b + \rho_b \widetilde{U}_{1b}^2 + \rho_b \widetilde{u_{1b}^2} + F/A, \tag{37}$$

where A is the cross-sectional area at the generator position. The quantity $\widetilde{u_{1b}^2}$ determines the turbulence intensity, and its value is an input to the problem (as are $\widetilde{u_{2b}^2}$ and $\widetilde{u_{3b}^2}$). The drag force F acting on the generator is also

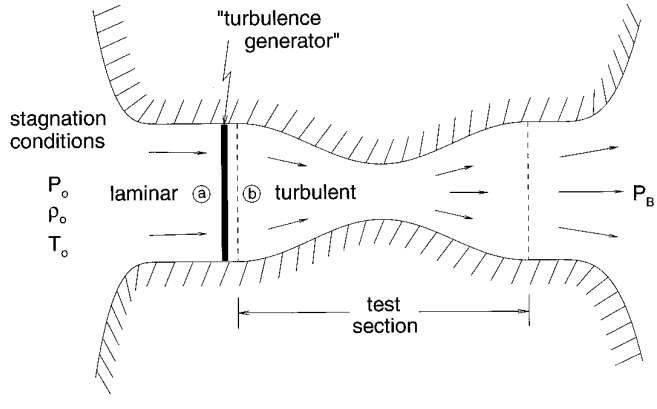

FIG. 1. Physical setup of nozzle using inflow/outflow boundary conditions.

assumed known. Assuming steady isentropic flow upstream of the generator, this leads to the additional relations:

$$\text{energy:}\quad h_o = \frac{\gamma}{\gamma-1}\frac{P_o}{\rho_o} = \frac{\gamma}{\gamma-1}\frac{P_a}{\rho_a} + \frac{\widetilde{U_{1a}^2}}{2}, \tag{38}$$

$$\text{eq. of state:}\quad \left(\frac{P_a}{P_o}\right) = \left(\frac{\rho_a}{\rho_o}\right)^{\gamma}. \tag{39}$$

The assumption of steady flow upstream is reasonable since this will be the case at large time if P_B is held fixed. A final relationship comes by assuming that the entropy does not increase across the generator, hence the term ''idealized.'' The equation of state then has the same form downstream of the generator:

$$\text{eq. of state:}\quad \left(\frac{P_b}{P_o}\right) = \left(\frac{\rho_b}{\rho_o}\right)^{\gamma}. \tag{40}$$

With $\widetilde{u_{1b}^2}$ and F as inputs, Eqs. (36)–(40) give five equations in six unknowns. The final input, which will be evident shortly, will be an extrapolated value from the interior of the nozzle.

Implementation of inflow/outflow boundary conditions is significantly more complex than periodic boundary conditions, in part because of the compressible nature of the flow. These boundary conditions must be designed to allow traveling pressure waves to escape, which implies some type of feedback mechanism. One approach which accomplishes this is characteristic-based boundary conditions, discussed in detail in Refs. [24, 25]. While this approach is primarily aimed at grid-based numerical schemes, it has been successfully implemented in the current particle method.

In a turbulent flow simulation, the characteristic-based approach begins by transforming the system of governing equations for mean quantities into equations for characteristic variables. In the particle method the governing equations are ultimately the particle equations of motion (Eqs. (5)–(7)), but these directly imply an evolution equation for the mass density function $\mathscr{F}$ (Eq. (8)). This equation for $\mathscr{F}$ together with Eq. (2) subsequently yields governing equations for any desired mean involving functions of density and velocities.

The level of closure determines which mean quantities are used in the computation of boundary conditions. Consistent with the overall method, the density is represented only by its mean value. The velocity pdf of incoming and outgoing fluid is assumed to be Gaussian, and hence, it is completely characterized by its mean and variance. For quasi-1D flow, both $\widetilde{U_2}$ and $\widetilde{U_3}$ are zero. Furthermore, we impose $\widetilde{u_i u_j} = 0$ for $i \neq j$ at the inflow boundary. It can then be shown that these off-diagonal components remain zero everywhere in the domain. The resulting vector of required mean quantities is therefore $\mathbf{V} = \{\langle\rho\rangle, \widetilde{U_1}, \widetilde{u_1^2}, \widetilde{u_2^2}, \widetilde{u_3^2}\}^{\mathrm{T}}$. The evolution equations for these quantities, obtained from the equation for $\mathscr{F}$, can be written in matrix form as

$$\frac{\partial \mathbf{V}}{\partial t} + \mathbf{A}\frac{\partial \mathbf{v}}{\partial x_1} = \mathbf{S}, \tag{41}$$

where

$$\mathbf{A} = \begin{bmatrix} \widetilde{U_1} & \langle\rho\rangle & 0 & 0 & 0 \\ \dfrac{\langle a\rangle^2 + \widetilde{u_1^2}}{\langle\rho\rangle} & \widetilde{U_1} & 1 & 0 & 0 \\ 0 & 2\widetilde{u_1^2} & \widetilde{U_1} & 0 & 0 \\ 0 & 0 & 0 & \widetilde{U_1} & 0 \\ 0 & 0 & 0 & 0 & \widetilde{U_1} \end{bmatrix}. \tag{42}$$

Here, $\langle a\rangle$ is the mean local sound speed $[\gamma\langle P\rangle/\langle\rho\rangle]^{1/2}$. The characteristic speeds, given by the eigenvalues of $\mathbf{A}$, are

$$\lambda_{1\cdots5} = \{\widetilde{U_1} - \sqrt{\langle a\rangle^2 + 3\widetilde{u_1^2}}, \widetilde{U_1} + \sqrt{\langle a\rangle^2 + 3\widetilde{u_1^2}}, \widetilde{U_1}, \widetilde{U_1}, \widetilde{U_1}\}, \tag{43}$$

and the corresponding characteristic variables are

$$w_1 = -\langle\rho\rangle\left[\frac{1}{\langle\rho\rangle}\frac{\langle a\rangle^2 + \widetilde{u_1^2}}{\sqrt{\langle a\rangle^2 + 3\widetilde{u_1^2}}}\bigg|_0\right] + \widetilde{U_1} - \widetilde{u_1^2}\left[\frac{1}{\sqrt{\langle a\rangle^2 + 3\widetilde{u_1^2}}}\bigg|_0\right],$$

$$w_2 = +\langle\rho\rangle\left[\frac{1}{\langle\rho\rangle}\frac{\langle a\rangle^2 + \widetilde{u_1^2}}{\sqrt{\langle a\rangle^2 + 3\widetilde{u_1^2}}}\bigg|_0\right] + \widetilde{U_1} + \widetilde{u_1^2}\left[\frac{1}{\sqrt{\langle a\rangle^2 + 3\widetilde{u_1^2}}}\bigg|_0\right],$$

$$w_3 = -\langle\rho\rangle\left[2\frac{\widetilde{u_1^2}}{\langle\rho\rangle}\bigg|_0\right] + \widetilde{u_1^2}, \quad w_4 = \widetilde{u_2^2}, \quad w_5 = \widetilde{u_3^2}. \tag{44}$$

Here, the system of equations has been linearized about local values at the end of the previous time step (in brackets and with the subscript 0). The applied boundary values for the primitive variables are extracted from these characteristic variables, for example $\widetilde{U_1}^{\mathrm{BV}} = \frac{1}{2}(w_1 + w_2)$.

Since the flows considered in this study are subsonic, λ_1 is less than zero and consequently the characteristic w_1 is outgoing at the inlet and incoming at the outlet. At the inlet, w_1 is evaluated using values for $\langle\rho\rangle$, $\widetilde{U_1}$, and $\widetilde{u_1^2}$ extrapolated from the interior, while w_{2-5} are evaluated using the upstream quantities $\{\rho_b, \widetilde{U_{1b}}, \widetilde{u_{ib}^2}\}$ (computed from Eqs. (36)–(40) using extrapolated density as the additional input). The converse applies at the outlet, with the specified back pressure used as incoming information. The extrapolated values at boundaries are obtained from the particles using kernel estimates. The extrapolating kernel has the

same form as Eq. (20), except that any portion of the kernel which extends outside the boundary is set to zero (along with the appropriate renormalization).

Having computed the boundary values for $\langle \rho \rangle$, $\widetilde{U_1}$, and $\widetilde{u_i^2}$, the final step is to enforce this at the particle level. This is accomplished by inserting new particles in buffer zones at both the inlet and outlet. These zones extend outside the nozzle and have widths proportional to the kernel smoothing length used in the simulation. New particles are uniformly distributed in each buffer zone so as to yield the required boundary value for mean density, and their velocities are sampled from a Gaussian pdf having means and variances given by the boundary values. For the new particles, both initial and predicted values for mean density and velocity are set to the mean boundary values, and the mean pressure gradient force is set to zero (consistent with having a constant mean density in each zone).

The use of such buffer zones has a number of advantages. First, the computation of kernel estimates in the interior of the nozzle requires no modifications to the algorithm. A kernel estimate near the boundary will receive contributions from both interior and buffer particles. Second, particles will flow into and out of the nozzle naturally. Not every particle in the inlet buffer zone will enter the nozzle in the course of one time step. It is also possible for interior particles to move upstream into the inlet buffer zone, or for particles in the outlet buffer zone to enter the nozzle, although the probability of these events is low. Finally, this approach has been found to not affect the overall stability of the method. Other approaches [26] have been tried, but these were found to allow unphysical pressure waves to build up and not properly exit the computational domain.

The computation of boundary values and insertion of particles into buffer zones is performed once at the beginning of each time step. At the end of each time step, all particles remaining in buffer zones are discarded. Results using inflow/outflow boundary conditions are presented in Section 5.2.

5. NUMERICAL RESULTS

The numerical properties of the method have been studied through a variety of quasi-1D test problems. These problems range from simple steady and unsteady laminar flows in a constant area duct to nonstationary inhomogeneous anisotropic turbulence in a varying area duct. Flows with shocks are excluded since only isentropic flows are presently considered. This restraint may be removed, however, by solving an additional evolution equation for the internal energy or enthalpy of each particle and using a different equation of state.

The purpose of presenting these results is to show that the combined PDF/SPH method is robust and can be readily applied to a wide range of problems. Although other methods exist for the solution of these problems, a performance comparison between them and the current method is not presented. The results shown here are limited to inert flows, and any fair comparison involving PDF methods ought to also include results for reacting flows.

5.1. *Sample Flow Calculation Using Periodic BCs*

This example involves unsteady isentropic turbulent flow through a periodic nozzle of sinusoidally varying area, with a maximum to minimum area ratio of 2 and a period L of 1. For the initial condition, particles are deterministically positioned with uniform interparticle spacing $\Delta x = L/N$ to give constant linear density. The initial velocity pdf is everywhere Gaussian with mean $\widetilde{U_1} = 1$ and variance $\widetilde{u_i u_j} = (0.1a_o)^2 \delta_{ij}$, where $a_o = (\gamma P_o/\rho_o)^{1/2}$ is a reference speed of sound ($P_o = \rho_o = 1$). The turbulent frequency ω is set equal to 1 everywhere. Note that the pressure can exceed P_o in this simulation—P_o is a reference pressure and not a stagnation pressure.

Figures 2 and 3 show evolution of the mean flow quantities $\langle P \rangle$, $\widetilde{U_1}$, and $\widetilde{u_i u_j}$. Each field is shown at times $t = 0$, $t = \frac{1}{4}$, and $t = \frac{1}{2}$. The mean flow is from left to right, and the nozzle throat is located at $x_1/L = \frac{1}{2}$. In this simulation 16,384 particles were used, with a time step $\Delta t = \frac{1}{200}$ and a smoothing length $h = \frac{1}{16}$. Variance reduction was also performed using 64 particles per ensemble. Solid lines are the estimated fields obtained by averaging 16 independent

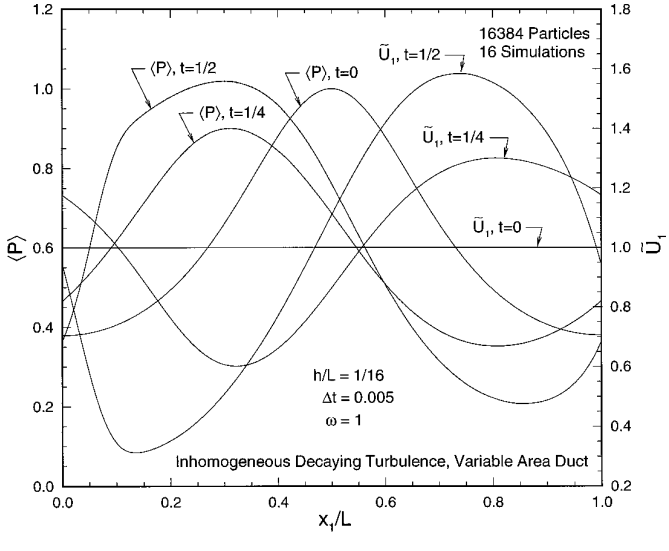

FIG. 2. Temporal evolution of $\langle P \rangle$ and $\widetilde{U_1}$ for periodic test case. Solid lines are fields obtained by averaging 16 independent simulations, and dotted lines (barely visible) show corresponding 95% confidence intervals. The nozzle cross-sectional area varies sinusoidally, with the throat located at $x/L = 0.5$.

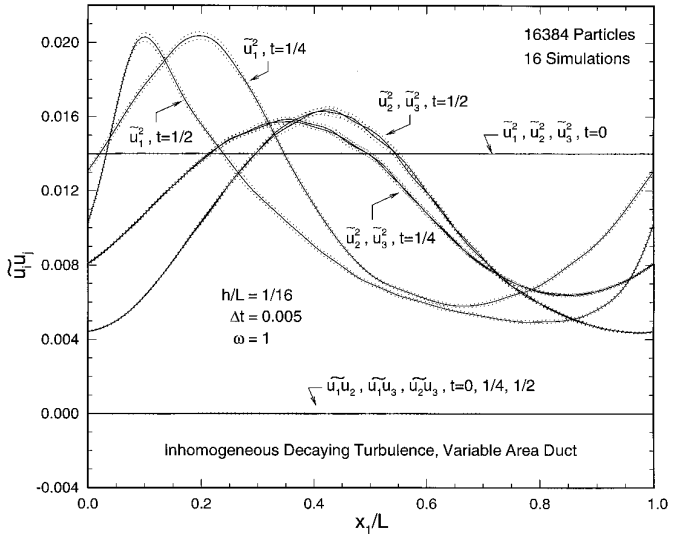

FIG. 3. Temporal evolution of $\widetilde{u_iu_j}$ for periodic test case; dotted lines show 95% confidence intervals for mean fields. Sinusoidal nozzle area distribution, with throat at $x/L = 0.5$.

simulations, and dashed lines (visible with the Reynolds stresses) give 95% confidence intervals for these fields. In other words, any field obtained by averaging 16 independent simulations will, with 95% probability, lie within the given interval.

Although the turbulence is initially homogeneous and isotropic, Fig. 3 shows that it quickly becomes inhomogeneous and anisotropic due to production from area effects and the streamwise mean pressure gradient.

Computational requirements are quite modest for this calculation. Approximately 1.2 CPU seconds per time step are needed in each simulation on an IBM RS/6000 Model 590. The total CPU time needed to obtain the solution at $t = \frac{1}{2}$ using 16 simulations is therefore about 30 min. A cluster of machines was actually used in this calculation, with each simulation run on a separate node. The actual wall-clock time required to obtain the solution was therefore only 2 min.

5.2. *Sample Flow Calculations Using Inflow/Outflow BCs*

This section presents results using inflow/outflow boundary conditions. Three types of results are presented: (1) the effect of varying back pressure P_B on distributions of various mean quantities; (2) the predicted steady-state mass flow rate versus P_B, including the phenomenon of choked flow; and (3) the effect of varying turbulent frequency on stationary distributions of $\widetilde{u_iu_j}$. The particle method's results are also compared to predictions from a simple finite-difference code and are found to be very satisfactory.

In these calculations using inflow/outflow boundary conditions the back pressure is held fixed, and consequently, the velocity pdf relaxes to a stationary distribution. The relaxation time-scale is on the order of the residence time T_R—the average time needed for a fluid particle to flow completely through the nozzle. Given the stationary mean density field, T_R can be obtained through the integral relation

$$T_R = \frac{1}{\dot{m}_{ss}} \int_0^L \langle \rho(x) \rangle \, dx, \tag{45}$$

where $\dot{m}_{ss}$ is the steady-state mass flow rate through the nozzle. The laminar steady-state solution provides a good initial estimate to T_R. Each simulation is integrated out several residence times to insure that all information from the initial conditions has convected out of the nozzle and that the flow has reached a stationary distribution.

For stationary distributions the statistical error may be further reduced by using time-averaging. In this procedure, flow quantities are sampled once per time step and averaged together over a finite time interval. Although the procedure is straightforward, care must be taken when computing confidence intervals for the time-average since these must take into account autocorrelation of the data. For a stationary random variable q with zero mean, the autocorrelation coefficient ρ_q is defined as

$$\rho_q(s) = \frac{\langle q(t)q(t+s) \rangle}{\langle q^2 \rangle}. \tag{46}$$

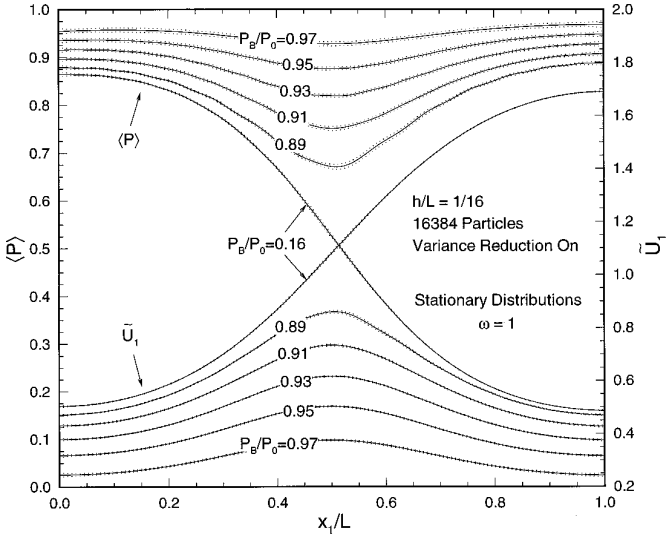

FIG. 4. Stationary distributions of $\langle P \rangle$ and $\widetilde{U_1}$ for varying P_B using inflow/outflow boundary conditions; 95% confidence intervals given by dotted lines. Variance reduction with $N_e = 64$ was used. Flow is from left to right.

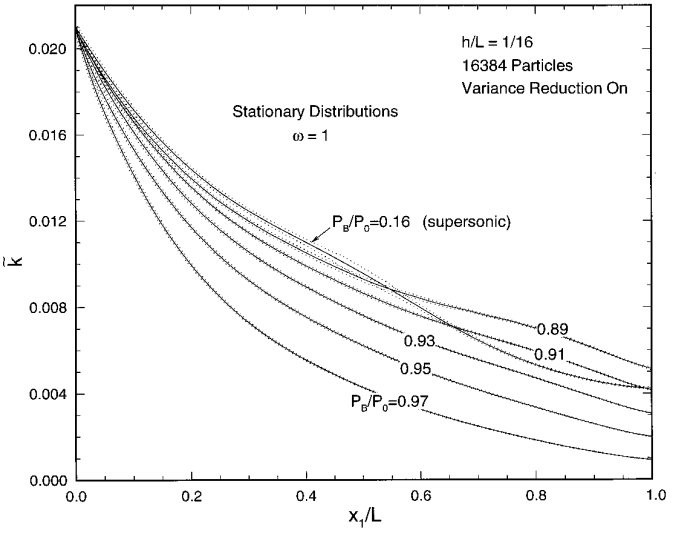

FIG. 5. Stationary distributions of $\tilde{k}$ for varying P_B using inflow/outflow boundary conditions; dotted lines show 95% confidence intervals. Variance reduction used here with $N_e = 64$.

The integral of ρ_q over all time,

$$T_q = \int_0^\infty \rho_q(s)\, ds, \tag{47}$$

yields a time-scale for the quantity q, and samples of q taken $\mathcal{O}(T_q)$ intervals apart are approximately independent. The equivalent number of independent samples in the time-averaging interval is thus known and can be used to calculate the correct confidence interval widths.

In each simulation, the particle distribution is initialized to yield $\langle\rho\rangle$ and $\widetilde{U_1}$ equal to the steady-state laminar solution for the given back pressure. Only small changes will then occur to these mean quantities as they relax to their stationary distributions. The variances $\widetilde{u_i u_j}$ are initialized to $(0.1a_o)^2\delta_{ij}$, as before. This initial condition is far from the stationary distribution so that large changes occur in the variances as they relax.

Figures 4 and 5 show stationary distributions of $\langle P\rangle$, $\widetilde{U_1}$, and $\tilde{k}$ for different values of P_B. The nozzle cross-sectional area is sinusoidal with an area ratio of 1.5. Each simulation used 16,384 particles, a time step of $\frac{1}{200}$, and a smoothing length $h = \frac{1}{16}$. Time-averaging was performed from $t = 7.5$ until $t = 15$, and variance reduction (with $N_e = 64$) was used for the entire simulation. As the nozzle back pressure decreases from 0.97 to 0.89, the mean velocity gradually rises and the pressure drops, as expected. Each quantity hits its extreme value at the throat located at the center. Careful inspection reveals that the distributions are not symmetric about the throat: in the divergent section pressure is slightly higher and mean velocity slightly lower. This effect is simply a consequence of momentum conservation in the presence of the decaying turbulence visible in Fig. 5. The accumulated amount of dissipation decreases with P_B since $\widetilde{U_1}$ consequently increases and the residence time is reduced.

Figures 4 and 5 also show mean distributions at the isentropic (shock-free) design condition, $P_B = 0.16$. At this condition the flow is sonic at the throat and supersonic in the divergent section. The outlet Mach number is 1.855 which is within 0.05% of the laminar value. The dissipation of $\tilde{k}$ in the divergent section is somewhat greater for the supersonic case as compared to the subsonic distributions. This is actually due to negative production: $\partial\widetilde{U_1}/\partial x_1$ remains positive everywhere in the nozzle for the supersonic case.

The steady-state mass flux through the nozzle can be found by computing the average particle flux across the boundaries. This quantity is plotted versus P_B/P_o in Fig. 6. The plot is divided into four regions, each corresponding to a different flow regime for the nozzle: fully subsonic, nonisentropic with shock, overexpanded, and underexpanded. Since the particle method cannot be applied to flows containing shocks, the second region (shaded in gray) is not allowed. Over- and underexpanded flows pose no problem for the particle method. In these regimes the flow is choked, and the particle method predicts this phenomenon accurately.

Figure 6 also shows mass flow rates computed using a finite-difference code. This code solves the steady-state governing equations for $\langle P\rangle$, $\widetilde{U_1}$, $\widetilde{u_i u_j}$, and all triple-velocity correlations using a space-marching approach. The governing equations are those derived directly from the mass density function evolution equation, Eq. (8). By them-

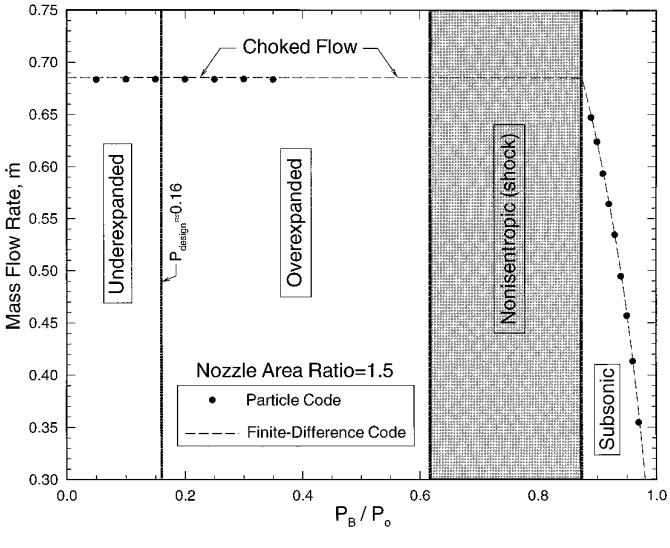

FIG. 6. Computed steady-state mass flow rates for varying P_B.

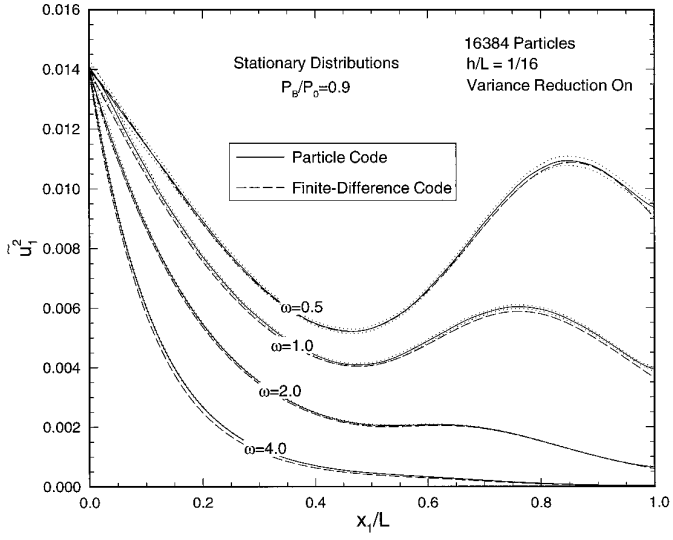

FIG. 7. Stationary distributions of $\widetilde{u_1^2}$ for varying ω using inflow/outflow boundary conditions; dotted lines show 95% confidence intervals. Variance reduction turned on, with $N_e = 64$. Flow is from left to right.

selves, however, the mean equations are unclosed at the level of fourth moments. Closure is achieved by invoking the Millionschchikov hypothesis [27], which states:

$$\langle u_i u_j u_k u_l \rangle = \langle u_i u_j \rangle \langle u_k u_l \rangle + \langle u_i u_k \rangle \langle u_j u_l \rangle + \langle u_i u_l \rangle \langle u_j u_k \rangle. \quad (48)$$

This equation is exact for Gaussian distributions. Results from the particle code have shown that the flows considered in this study are reasonably close to Gaussian, and therefore, meaningful comparisons can be made between the particle and finite-difference codes. As can be seen in Fig. 6, the mass flow rates are in excellent agreement.

The third and final set of results shows the effect of varying turbulent frequency on stationary distributions of $\widetilde{u_i u_j}$. Figure 7 gives distributions of the streamwise normal stress $\widetilde{u_1^2}$, and Fig. 8 shows the cross-stream normal stresses $\widetilde{u_2^2}$ and $\widetilde{u_3^2}$. Since $1/\omega$ is the dissipation time-scale, the choice of ω controls the balance between production and dissipation in the evolution of these stresses. The production time-scales are just reciprocals of the corresponding strain-rates $S_1 = \partial \widetilde{U_1}/\partial x_1$ and $S_{2,3} = (\widetilde{U_1}/A)(dA/dx_1)$. These are functions of the mean flow and do not vary significantly in this set of results. The parameter S_i/ω, the ratio of dissipation and production time-scales, measures the relative importance of each process. A flow with large $|S_i/\omega|$ is dominated by production, and one with small $|S_i/\omega|$ is dominated by dissipation.

The curves for $\omega = 4.0$ show dissipation dominating over production, with the turbulence decaying rapidly along the full length of the nozzle. The parameters S_1/ω and $S_{2,3}/\omega$ range only between ± 0.59 and ± 0.20, respectively. The curves for $\omega = 0.5$, on the contrary, clearly show the influence of production. In this case, the parameters range between ± 4.2 and ± 1.5. The streamwise stress $\widetilde{u_1^2}$ shows strong growth in the divergent section, while $\widetilde{u_2^2}$ and $\widetilde{u_3^2}$ remain essentially flat in the convergent section before rapidly decaying beyond the throat. This behavior can be physically interpreted as the stretching of streamwise vortices. As the flow accelerates in the convergent section and squeezes through the throat, the vorticity increases and creates higher cross-stream turbulence intensities. In the particle evolution equations, this effect is represented by the area source term:

$$du^*_{2,3} = \cdots - \left(\frac{1}{2}\frac{\widetilde{U_1}}{A}\frac{dA}{dx_1}\right) u^*_{2,3}\, dt + \cdots. \quad (49)$$

Figures 7 and 8 also include finite-difference solutions for comparison with the particle code results. To ensure meaningful comparisons, the flatness distributions of the particle code results were all measured to be within 8% of the Gaussian value of 3. Once again, the agreement between the particle and finite-difference results is excellent.

5.3. *Variance Reduction Performance*

The performance of the variance-reduction technique is controlled by the ratio $N_e/\overline{N}$, hereafter called Λ. Physically, $1/\Lambda$ can be interpreted as the average number of ensembles contributing to each kernel estimate. Although there are no computational restrictions on Λ, it is expected that the

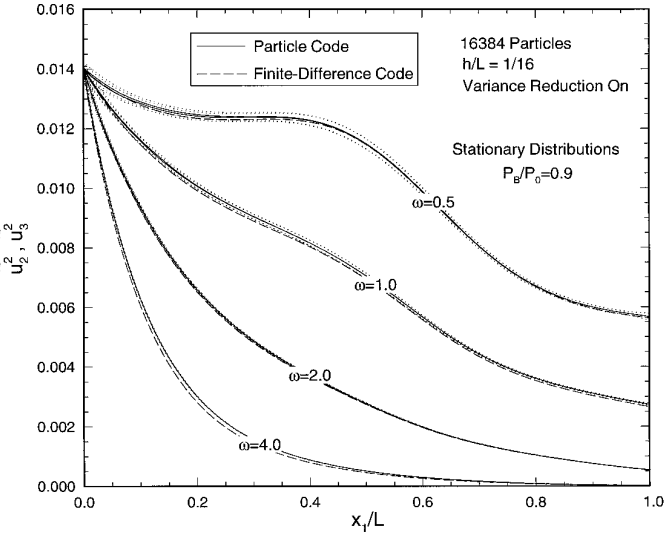

FIG. 8. Stationary distributions of cross-stream normal stresses $\widetilde{u_2^2}$ and $\widetilde{u_3^2}$ for varying ω using inflow/outflow boundary conditions; 95% confidence intervals given by dotted lines.

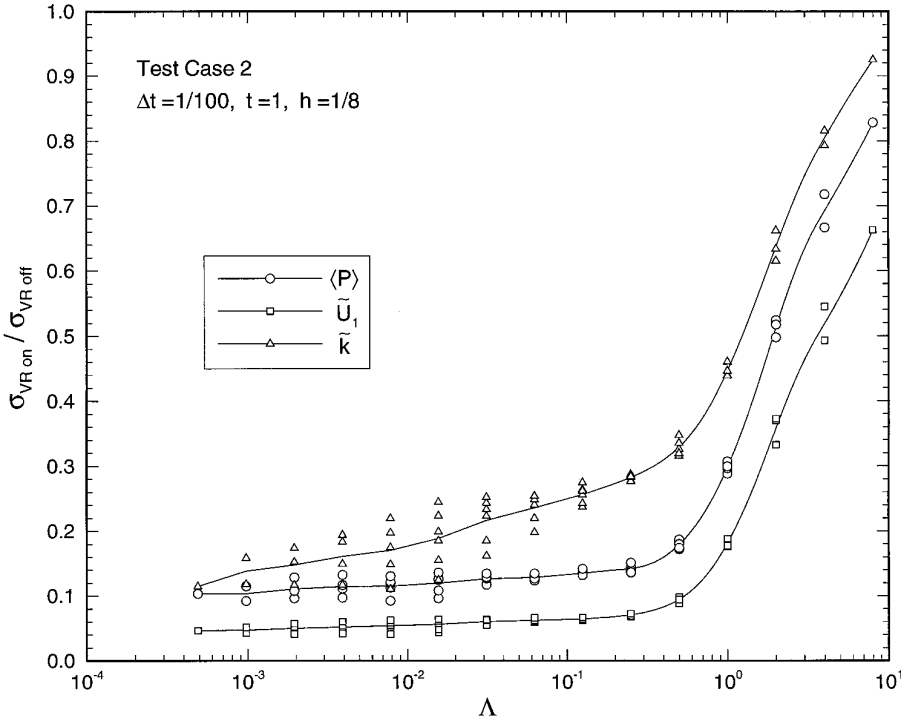

FIG. 9. Variance reduction performance vs. Λ for test case 2. Solid curves are cubic-spline fits to each set of data. The parameter $\overline{h}$ is fixed at $\frac{1}{8}$, while N varies between 4,096 and 131,072 and N_e ranges between 8 and 4,096.

variance-reduction technique is most effective for $\Lambda \leq 1$, approximately.

Figure 9 shows the performance of the method for a selected test problem (this is test case 2, fully described in Section 6.1). The vertical axis gives the ratio of root-mean-square statistical errors, $\sigma_{\text{VR on}}/\sigma_{\text{VR off}}$, where $\sigma_{\text{VR on}}$ is the rms error obtained with variance reduction on, and $\sigma_{\text{VR off}}$ is that obtained with it off. The ratio is plotted versus Λ for the three quantities $\langle P\rangle$, $\widetilde{U_1}$, and $\tilde{k}$. The rms statistical errors for each point were estimated from 128 independent simulations, each of which was integrated out to $t = 1$ using a time step of $\frac{1}{100}$. The smoothing length h was kept fixed at $\frac{1}{8}$ for all points, while N was varied between 4,096 and 131,072, and N_e was varied between 8 and 4,096. To improve readability, a spline was passed through the data for each flow quantity.

The data clearly show a decrease in statistical error for each quantity as Λ decreases (and the number of ensembles per kernel estimate increases). For $\Lambda = 8$, the largest value tested, $\sigma_{\text{VR on}}/\sigma_{\text{VR off}}$ is still less than 1. For small Λ, the statistical error has been decreased by approximately a factor of 10 for $\langle P\rangle$, 20 for $\widetilde{U_1}$, and 5 for $\tilde{k}$.

Computational performance of the variance reduction technique is measured through the *quality* q [17]. The quality is related to the rms statistical error by $q = T\sigma^2$, where T is the CPU time needed to achieve the statistical error level σ. Smaller values of q are better, and the computational performance can therefore be measured by the ratio

$$\frac{q_{\text{VR off}}}{q_{\text{VR on}}} = \frac{T_{\text{VR off}}}{T_{\text{VR on}}} \frac{\sigma^2_{\text{VR off}}}{\sigma^2_{\text{VR on}}}. \tag{50}$$

The CPU time is found to increase by approximately a factor of 1.7, independent of N_e, when using this variance-reduction technique. But as Fig. 9 shows, the corresponding decrease in σ is significant. Thus, given a specified statistical error tolerance, the decrease in required CPU time with variance reduction on can be estimated from Eq. (50) to be a factor of up to 15 for $\tilde{k}$, 60 for $\langle P\rangle$, and 230 for $\widetilde{U_1}$. Clearly then this variance-reduction technique is essential for obtaining accurate solutions using a minimal amount of CPU time.

6. CONVERGENCE OF PARTICLE METHOD

In the particle method described in this paper, various numerical errors arise due to finite values of the number of particles N, the time step Δt, and the smoothing length h. Success of the method requires that these errors converge to zero as the appropriate parameters approach their limiting values. Numerical stability is also essential for the convergence of these errors and, hence, for the success of

the method. A detailed numerical study has therefore been performed in which the different types of errors arising in the particle method are identified and shown to converge at the expected rates.

Convergence of stochastic methods can be interpreted in either a strong or weak sense [28]. With PDF-based particle methods it is appropriate to require weak convergence. Specifically, the discrete Lagrangian mdf $\mathscr{F}_N^*$ represented by the particles should converge *in distribution* to the actual modeled mdf $\mathscr{F}^*$ [3, 17], and consequently any mean quantity evaluated from the particles should converge in mean-square sense to the actual mean, if it exists. Using this criterion for convergence, four different types of errors can be identified by considering estimating a mean quantity $\langle Q(x,t)\rangle_{\overline{N},\overline{h},\Delta t}$ at a fixed time t using an average number of particles per kernel $\overline{N} = Nh/L$, a normalized smoothing length $\overline{h} = h/L$, and a time step Δt. Given the exact value $\langle Q\rangle$, the error in this estimate is a random variable and can be decomposed as

$$\begin{aligned}\langle Q\rangle_{\overline{N},\overline{h},\Delta t} - \langle Q\rangle &= [\langle Q\rangle_{\overline{N},\overline{h},\Delta t} - \langle\langle Q\rangle_{\overline{N},\overline{h},\Delta t}\rangle] \\ &\quad + [\langle\langle Q\rangle_{\overline{N},\overline{h},\Delta t}\rangle - \langle Q\rangle] \\ &= \Sigma_Q(t,\overline{N},\overline{h},\Delta t) + D_Q(t,\overline{N},\overline{h},\Delta t) \\ &= \Sigma_Q + T_Q + H_Q + B_Q,\end{aligned} \tag{51}$$

where Σ_Q is the statistical error and D_Q is the deterministic error. As mentioned earlier, statistical error arises from the initial conditions, the Wiener processes, and the coefficients in the Langevin equation. The deterministic error can be further broken down into temporal error T_Q, spatial error H_Q, and bias B_Q. Temporal error results from using a finite time step in the time integration scheme, while spatial error is due to the finite smoothing length used in the kernel estimates. Finally, bias is the deterministic error resulting from using a finite number of particles. These four errors are studied and discussed individually.

6.1. *Test Cases*

Convergence of the method is studied through two test cases. The first is incompressible homogeneous isotropic turbulence decaying in a constant area duct, for which all mean properties are independent of position. Initial conditions are similar to those for the previously described sample flow: particles are deterministically positioned with uniform spacing, and the velocity pdf is specified to have mean $\widetilde{U_1} = 0$ and variance $\widetilde{u_i u_j}\big|_0 = (0.1a_o)^2\delta_{ij}$. Periodic boundary conditions are imposed at the duct inlet and outlet, and the turbulent frequency ω is set equal to 1. An analytic solution to this problem exists and facilitates measuring errors. The solution is given by

$$\langle\rho\rangle = \langle P\rangle = 1, \quad \widetilde{U_1} = 0, \quad \widetilde{u_i u_j}(t) = \widetilde{u_i u_j}\big|_0 e^{-\omega t}. \tag{52}$$

Although the mean density, pressure, and velocity remain constant for all time, the Reynolds stresses decay exponentially at the constant rate ω, and hence, this test case is well-suited for measuring the temporal convergence of the method. This test case is subsequently referred to as test case 1.

The second test case is more comprehensive and consists of inhomogeneous anisotropic turbulent flow in a constant area duct. The initial velocity pdf is again specified to have mean zero and constant variance $\widetilde{u_i u_j}\big|_0 = (0.1a_o)^2\delta_{ij}$, while the initial mean density distribution is specified to be a sine wave of period 1 and amplitude $\frac{1}{4}$ about $\langle\rho\rangle = 1$. This is achieved by deterministic positioning of the particles. Boundary conditions are again periodic, and the turbulent frequency is kept equal to 1. As in the sample flow calculations, the turbulence is initially homogeneous and isotropic, but quickly becomes inhomogeneous and anisotropic due to production of kinetic energy from the streamwise mean pressure gradient. Despite the lack of an analytic solution for this test case (subsequently called test case 2), convergence of errors can still be studied by making use of error extrapolation techniques.

6.2. *Statistical Error*

Of the four errors present in the method, statistical error is usually dominant. The statistical error Σ_Q has the form $\overline{N}^{-1/2}\xi S_Q$, where ξ is a standardized random variable with zero mean and unit variance and S_Q is the standard error defined as

$$S_Q(t,\overline{N},\overline{h},\Delta t) = [\overline{N}\,\mathrm{var}(\langle Q\rangle_{\overline{N},\overline{h},\Delta t})]^{1/2}. \tag{53}$$

The rms statistical error σ_Q is therefore $\overline{N}^{-1/2}S_Q$. As $\overline{N}$ approaches infinity S_Q becomes independent of $\overline{N}$. In Σ_Q the inverse square root dependence is on $\overline{N} = N\overline{h}$ and not N, since $\overline{N}$ is the average number of particles contributing to any kernel estimate.

Figure 10 shows results for test case 2 of the dependence of S_Q on $\overline{N}$ for the three quantities $\langle P\rangle$, $\widetilde{U_1}$, and $\tilde{k}$. Variance reduction was not used here, although qualitatively the results are the same. The data were obtained by estimating the variance of each quantity along the length of the duct from 128 statistically independent simulations, and then averaging the variance over the duct. A time step of $\frac{1}{100}$ was used in each simulation to obtain the solution at $t = 1$, which allows significant flow evolution to occur. The parameter $\overline{N}$ was varied from 256 to 65,536, and over this range S_Q remains fairly constant for each of the flow quantities. The actual number of particles N ranged between 4,096 and 131,072. The observed slight vertical scatter in σ_Q occurs due to a weak dependence of S_Q on $\overline{h}$. With $\overline{h}$ ranging between $\frac{1}{2}$ and $\frac{1}{16}$, the variation in S_Q is between 5–20%.

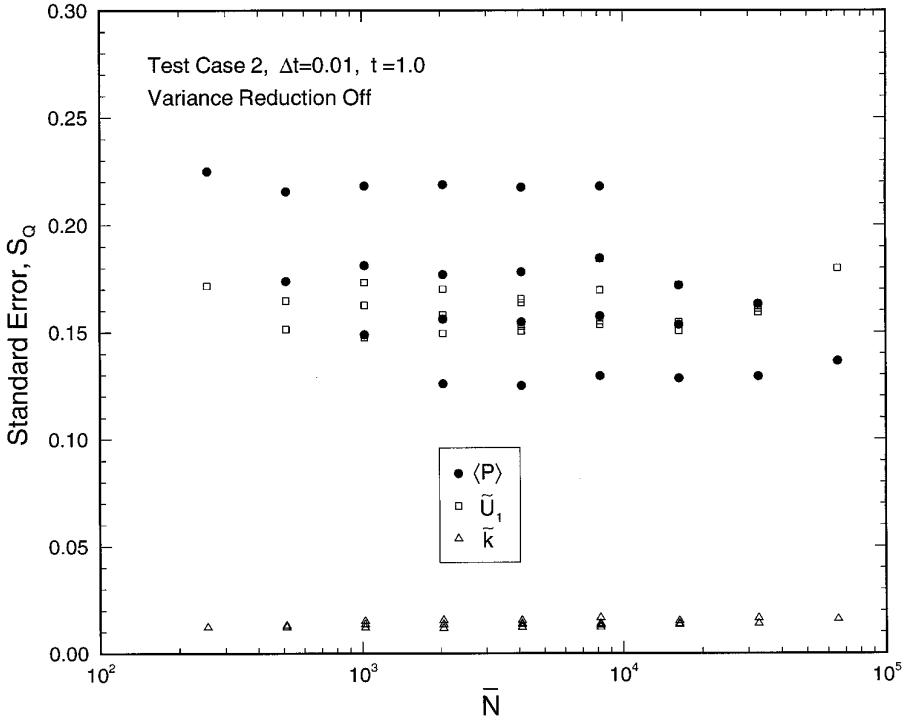

FIG. 10. Standard error vs. $\overline{N}$, with variance reduction off. The parameter N ranges between 4,096 and 131,072, while $\overline{h}$ varies from $\frac{1}{2}$ to $\frac{1}{16}$.

By doing a set of M independent simulations and averaging these together, the statistical error can be reduced by a factor of $M^{-1/2}$. Thus the total rms statistical error scales as $(M\overline{N})^{-1/2}S_Q$. By normalizing S_Q with respect to each flow quantity, it is possible to estimate what the product $M\overline{N}$ must be in order to achieve a normalized rms statistical error level of 1%, say. For test case 2 the requirements with variance reduction off are $M\overline{N} \approx 300$ for $\langle P\rangle$, $M\overline{N} \approx 14{,}500$ for $\widetilde{U_1}$, and $M\overline{N} \approx 42{,}000$ for $\tilde{k}$. With variance reduction on and Λ small, each requirement is scaled down by the corresponding ratio $\sigma^2_{\text{VR on}}/\sigma^2_{\text{VR off}}$ given previously. This ideally gives $M\overline{N} \approx 3$ for $\langle P\rangle$, $M\overline{N} \approx 35$ for $\widetilde{U_1}$, and $M\overline{N} \approx 1{,}700$ for $\tilde{k}$. Clearly the requirement for $\langle P\rangle$ is unrealistic since the corresponding minimum value of Λ is not small. Overall, in order to achieve an rms statistical error level of 1% for all three quantities, the product $M\overline{N}$ must therefore roughly be 1,700.

6.3. *Bias*

Bias is the contribution to the deterministic error caused by using a finite number of particles. Using the error decomposition form of Eq. (51), the bias B_Q in the ensemble average $\langle Q\rangle_{\overline{N},\overline{h},\Delta t}$ can be written

$$B_Q(t,\overline{N},\overline{h},\Delta t) = \langle\langle Q\rangle_{\overline{N},\overline{h},\Delta t}\rangle - \langle Q\rangle_{\infty,\overline{h},\Delta t}, \qquad (54)$$

with $\langle Q\rangle_{\infty,\overline{h},\Delta t} = \lim_{\overline{N}\to\infty}\langle Q\rangle_{\overline{N},\overline{h},\Delta t}$. Two sources of bias are present in this particle method. The first is the spatial integration error caused by replacing the integral form of the kernel estimate, Eq. (15), by the discrete sum over the particles, Eq. (18). For the laminar case in which the particle distribution has no randomness, this error can be shown to scale as $\overline{N}^{-4}$ in 1D; for the turbulent case in which there is particle disorder, it can be expected to scale as $\overline{N}^{-1}\log\overline{N}$ [15, 16]. The other source is statistical fluctuations in the coefficients, and this contribution can be expected to scale as $\overline{N}^{-1}$ [17]. Since turbulent flows are considered here, the combination of these two sources of bias is expected to yield a $\overline{N}^{-1}$ scaling.

Figure 11 is a log–log plot showing the scaling of bias with $\overline{N}$ for test case 2 at $t = 1$ for $\langle P\rangle$, $\widetilde{U_1}$, and $\tilde{k}$. This is without using variance reduction. The expected inverse scaling is observed over the range $\overline{N} = 4 \cdots 512$. Slope values given in the figure come from linear least-squares fits to each set of data. The dotted lines above and below each solid curve show 95% confidence intervals. Bias is more difficult to measure than statistical error. The data in Fig. 11 were obtained by comparing to a highly accurate solution with $\overline{N} = 8{,}192$, or 16 times the largest value shown in the figure. The plotted values are the rms errors averaged over the duct. Statistical error was reduced in all the data using multiple independent simulations; the maximum value of the ratio σ_Q/B_Q is 15%. For this test

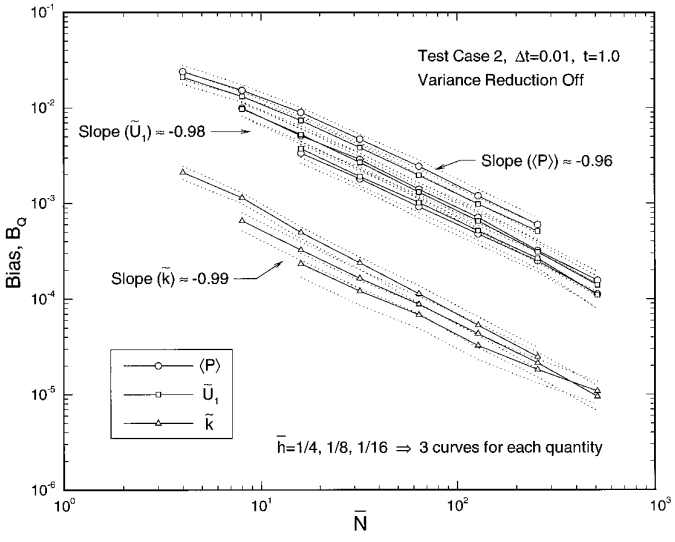

FIG. 11. Bias vs. $\overline{N}$, with variance reduction off. The total number of particles N varies between 16 and 8,192. Slopes are calculated from linear least-squares fits to each set of data, and dotted lines give 95% confidence intervals.

case, using the value $\overline{N} = 64$ keeps the bias for all three quantities below 1%.

The variance reduction technique introduces an additional contribution to the bias. This new contribution results from using a finite number of particles N_e in each ensemble, just as the previous contribution to the bias results from using a finite number of particles $\overline{N}$ in each kernel estimate. Preliminary numerical tests suggest that this additional contribution scales as N_e^{-p}, where p is close to the expected value of 1, but that the magnitude of this contribution is small, i.e.,

$$B_Q = \frac{b_1}{\overline{N}} + \frac{b_2}{N_e}, \quad b_2 \ll b_1. \tag{55}$$

These results, however, have not yet been fully quantified.

6.4. *Spatial Error*

Spatial error is caused by using a finite smoothing length $\overline{h}$ in the kernel estimates. A large value of $\overline{h}$ gives a more spatially averaged estimate to the desired quantity, whereas a small value of $\overline{h}$ gives a more local and, hence, more accurate, estimate. The spatial error H_Q in the ensemble average $\langle Q\rangle_{\overline{N},\overline{h},\Delta t}$ is

$$H_Q(t,\overline{h},\Delta t) = \langle Q\rangle_{\infty,\overline{h},\Delta t} - \langle Q\rangle_{\infty,0,\Delta t}, \tag{56}$$

where $\langle Q\rangle_{\infty,0,\Delta t} = \lim_{\overline{N}\to\infty,\overline{h}\to 0}\langle Q\rangle_{\overline{N},\overline{h},\Delta t}$. The accuracy of the kernel estimate given by Eq. (15) depends on the kernel. One can easily show using Taylor expansions that the spatial error scales as $\overline{h}^2$ for the symmetric kernel used in this method [16].

The scaling of H_Q with $\overline{h}$ for the particle method is shown in Fig. 12 for test case 2 without variance reduction. The spatial error scales as expected for each of the three quantities considered. Statistical error was reduced using multiple independent simulations, and 95% confidence intervals are shown in the figure. At $\overline{h} = \frac{1}{16}$ the rms statistical error in $\tilde{k}$ is 17% of $H_{\tilde{k}}$, and at $\overline{h} = \frac{1}{32}$ it is 33%. In each case bias was kept small by fixing $\overline{N}$ at 8,192. The spatial error was then measured relative to the solution obtained using Richardson extrapolation in the limit $\overline{h} = 0$. The choice $\overline{h} = \frac{1}{16}$ keeps the normalized spatial error for all three quantities at less than 1% for this test case.

With variance reduction turned on, precisely the same results as in Fig. 12 are obtained. The use of variance reduction therefore does not affect the spatial error, as expected.

6.5. *Temporal Error*

The fourth and final error present in the particle method is the temporal error, caused by using a finite time step in the predictor/corrector scheme. The temporal error T_Q in the ensemble average $\langle Q\rangle_{\overline{N},\overline{h},\Delta t}$ is given by

$$T_Q(t,\Delta t) = \langle Q\rangle_{\infty,0,\Delta t} - \langle Q\rangle, \tag{57}$$

with $\langle Q\rangle_{\infty,0,\Delta t}$ as previously defined. Direct empirical measurement of T_Q is difficult since all the other errors must be removed. Since the goal is to determine the dependence

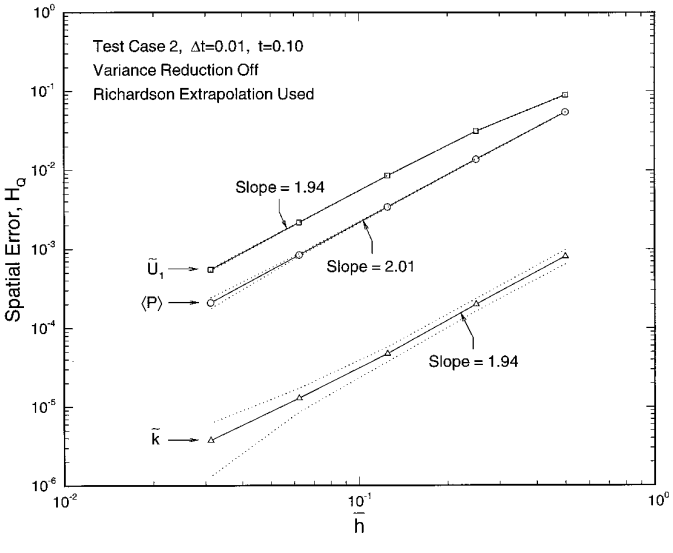

FIG. 12. Spatial error vs. $\overline{h}$, with variance reduction off. The parameter $\overline{N}$ is fixed at 8,192, while N ranges between 16,384 and 262,144. Slopes calculated from linear least-squares fits to data, and dotted lines give 95% confidence intervals.

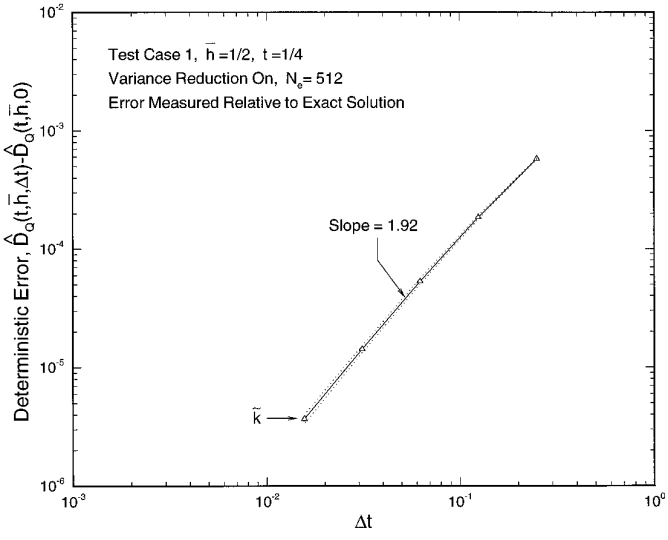

FIG. 13. Deterministic error vs. Δt for test case 1, with variance reduction on. Results shown use fixed values $\overline{N} = 65{,}536$, $N = 131{,}072$, and $\bar{h} = \frac{1}{2}$. Dotted lines give 95% confidence intervals.

of T_Q on Δt, this can still be done assuming that the dominant effect of Δt on the deterministic error is through the temporal error T_Q. In the empirical tests performed, both bias and statistical error were kept negligibly small. The error is therefore deterministic and of the general form

$$\begin{aligned}\hat{D}_Q(t,\bar{h},\Delta t) &= H_Q(t,\bar{h},\Delta t) + T_Q(t,\Delta t)\\ &= c_1(t,\Delta t)\bar{h}^2 + c_2(t,\bar{h})\Delta t^q,\end{aligned} \tag{58}$$

where $T_Q = c_2(t, 0)\Delta t^q$ and q is expected to be 2. Thus, by studying the dependence of $\hat{D}_Q$ on Δt it is possible to determine the value of q.

Results were obtained for both test cases, and these are shown in Figs. 13 and 14, respectively. In each test case the bias was kept negligibly small by fixing $\overline{N}$ at 65,536. In the results of test case 1, the bias is estimated to be less than 10% of the smallest value of $\hat{D}_Q$. In test case 2, the largest ratio of B_Q to $\hat{D}_Q$ is estimated at less than 2%. Although the results in Figs. 13 and 14 were obtained with variance reduction turned on ($N_e = 512$), further reduction of statistical error was achieved using multiple independent simulations. In both test cases the rms statistical error is less than 10% of $\hat{D}_Q$ over the entire range of Δt considered.

The deterministic error $\hat{D}_Q$ was calculated after having removed the bias and statistical error. In test case 1 the calculated solutions were compared to the exact analytic solution, while in test case 2 the solutions were compared to the one obtained using Aitken extrapolation in the limit $\Delta t = 0$. Richardson extrapolation was not used since the exact form of the scaling of the error with Δt in Eq. (58) was assumed unknown.

In both figures it is reasonably clear that the deterministic error scales approximately as Δt^2 for each of the considered quantities. The results are shown at $t = \frac{1}{4}$ with Δt ranging from $\frac{1}{8}$ to $\frac{1}{64}$. A smoothing length of $\bar{h} = \frac{1}{2}$ was used for both sets of results. In test case 1 the errors in $\langle P\rangle$ and $\widetilde{U_1}$ were not considered since these quantities do not evolve in time. The slope of the curve for $\tilde{k}$ using linear least-squares is 1.92, or very close to 2. In test case 2 the estimated scaling with Δt varies slightly depending on the quantity. For the three quantities $\langle P\rangle$, $\widetilde{U_1}$, and $\tilde{k}$, the scalings are $\Delta t^{1.88}$, $\Delta t^{2.25}$, and $\Delta t^{2.06}$, respectively. All of these are quite close to 2. Together these results indicate that the value of q in Eq. (58) must be at least 2, and hence, the temporal error T_Q scales at least as Δt^2.

For the two test cases considered here, the normalized temporal error can be kept below 1% by choosing $\Delta t = \frac{1}{10}$ in test case 1 and $\Delta t = \frac{1}{20}$ in test case 2. The actual time step used in a given problem, however, will most likely be less due to the CFL constraint.

6.6. *Computational Requirements for 1% Error*

The convergence results show that for test cases 1 and 2 all errors in the quantities $\langle P\rangle$, $\widetilde{U_1}$, and $\tilde{k}$ may be reduced to 1% by choosing $\bar{h} = \frac{1}{16}$, $\Delta t = \frac{1}{20}$, $\overline{N} = 64$, and either $M\,\overline{N} = 42{,}000$ (variance reduction off) or $M\,\overline{N} = 1{,}700$ (variance reduction on). The exact choice of M (and consequently $\overline{N}$) is not fixed. Knowing the requirement on $M\,\overline{N}$, it is good to keep $\overline{N}$ as large as possible since this minimizes the bias for the given statistical error level, although using large values for M has the advantage of giving more accurate confidence intervals.

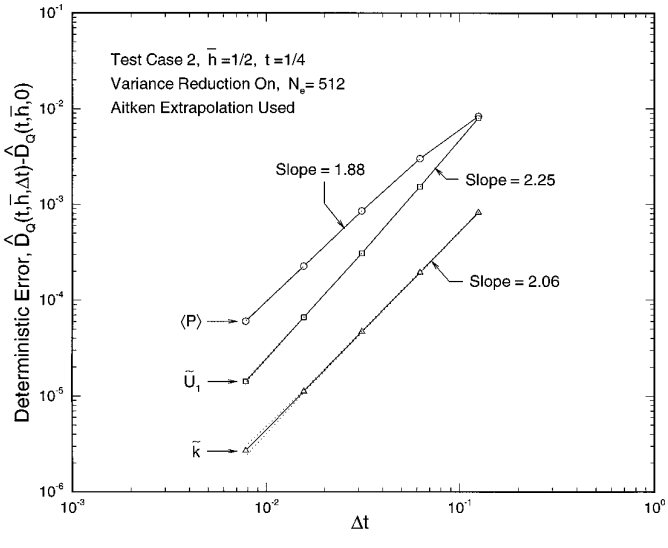

FIG. 14. Deterministic error vs. Δt for test case 2, with variance reduction on. Results shown use fixed values $\overline{N} = 65{,}536$, $N = 131{,}072$, and $\bar{h} = \frac{1}{2}$. Dotted lines show 95% confidence intervals.

Tests using these parameters were performed on an IBM RS/6000 Model 590 to determine the CPU time required for such a calculation, which scales as $M N = M \overline{N}/\overline{h}$. With variance reduction off, the choice $M = 6$ and $\overline{N} = 7000$ was used, and with $\overline{h} = \frac{1}{16}$ this gives $N = 112{,}000$. It also turns out that the CFL constraint limits Δt to about $\frac{1}{50}$. The solution at $t = 1$ was therefore obtained in 50 time steps and required 41 min of CPU time. With variance reduction on, the choice $N_e = 64$, $\overline{N} = 1700$, and $M = 1$ was used. This guarantees that Λ is small while also ensuring that the additional contribution to the bias, Eq. (55), is negligible. In this case the solution at $t = 1$ was obtained using only 165 s of CPU time, clearly showing the computational savings due to the variance-reduction technique. In both cases though, the solution was obtained in a very reasonable amount of CPU time.

In must be emphasized that these are conservative estimates since the computational work is inversely proportional to the square of the statistical error. Increasing the error level to 5%, say, decreases computational requirements by a factor of 25.

7. CONCLUSIONS

A PDF-based Monte Carlo particle method for compressible turbulent flows has been presented. The method is unique in that the mean pressure gradient is obtained directly from the particles using a grid-free approach based on kernel estimation techniques of SPH. The method has been successfully applied to a variety of quasi-1D flows, although flows with shocks or reaction have not yet been examined. The extension of the method to include these classes of flows, however, is expected to be straightforward.

The implementation of SPH within the Monte Carlo algorithm has been described, including a description of an algorithm which can calculate a set of N 1D SPH kernel estimates in $\mathcal{O}(N)$ work as opposed to $\mathcal{O}(N^2)$ work. The algorithm thus makes feasible the use of SPH in turbulent flow calculations. A variance-reduction technique which reduces the statistical error in first and second moments of velocity and position has also been described. The technique is shown to be very successful, decreasing computational requirements by more than an order of magnitude. Time integration of the particle evolution equations is done using a second-order accurate predictor/corrector scheme. This scheme is found to give high temporal accuracy with modest time step sizes.

The implementation of periodic and characteristic-based inflow/outflow boundary conditions has been described in detail, and a variety of numerical results using both types of boundary conditions has been presented. Stationary distributions of mean quantities have been compared to finite-difference solutions, and the agreement is excellent. The phenomenon of choked flow has also been accurately predicted by the particle method.

The convergence of the method has been successfully demonstrated through a comprehensive study of two test cases. Four types of numerical errors are present on the method. Statistical error scales as $\overline{N}^{-1/2}$, while the primary contribution to the bias scales as $\overline{N}^{-1}$. The others are spatial error, which scales as $\overline{h}^2$, and temporal error which scales as Δt^2. All four errors have been shown to scale at the expected rates.

The results of this convergence study also show that it is possible to reduce all errors below 1% using a very reasonable amount of CPU time. For the two test cases considered, the required CPU time to achieve an error level of approximately 1% is on the order of 3 min on a workstation. These results show that this particle method provides a feasible way to obtain accurate PDF solutions to compressible turbulent flow problems.

ACKNOWLEDGMENTS

Financial support for this work was provided by NASA Lewis Research Center through Grant NGT-50842 as part of the NASA Graduate Student Researchers Program. The authors also wish to thank NASA Lewis Research Center for allowing generous use of its LACE computing cluster.

REFERENCES

1. S. B. Pope, Computations of turbulent combustion: Progress and challenges, in *23rd International Symposium on Combustion* (The Combustion Institute, Pittsburgh, 1990), p. 591.
2. S. B. Pope, PDF methods for turbulent reactive flows, *Prog. Energy Combust. Sci.* **11,** 119 (1985).
3. S. B. Pope, Lagrangian PDF methods for turbulent flows, *Annu. Rev. Fluid Mech.* **26,** 23 (1994).
4. D. C. Haworth and S. B. Pope, A generalized Langevin model for turbulent flows, *Phys. Fluids* **29,** 387 (1986).
5. S. B. Pope and Y. L. Chen, The velocity-dissipation probability density function model for turbulent flows, *Phys. Fluids A* **2,** 1437 (1990).
6. S. B. Pope, Application of the velocity-dissipation probability density function model to inhomogeneous turbulent flows, *Phys. Fluids A* **3,** 1947 (1991).
7. C. Dopazo, Recent developments in PDF methods, in *Turbulent Reacting Flows* (Academic Press, New York, 1994), p. 375.
8. A. T. Hsu, Y.-L. P. Tsai, and M. S. Raju, Probability density function approach for compressible turbulent reacting flows, *AIAA J.* **7,** 1407 (1994).
9. M. S. Anand, S. B. Pope, and H. C. Mongia, PDF calculations for swirling Flows, AIAA Paper 93-0106 (unpublished).
10. M. S. Anand, S. B. Pope, and H. C. Mongia, Pressure algorithm for elliptic flow calculations with the PDF method, in *CFD Symposium on Aeropropulsion* (NASA Lewis Research Center, Cleveland, 1990).
11. S. B. Pope, Position, *Velocity and Pressure Correction Algorithm for Particle Method Solution of the PDF Transport Equations,* FDA Report 95-06, Cornell University, 1995.
12. D. C. Haworth and S. H. El Tahry, Probability density function approach for multidimensional turbulent flow calculations with appli-

cation to in-cylinder flows in reciprocating engines, *AIAA J.* **29,** 208 (1991).

13. S. M. Correa and S. B. Pope, Comparison of a Monte Carlo PDF/finite-volume mean flow model with bluff-body Raman data, in *24th International Symposium on Combustion* (The Combustion Institute, Pittsburgh, 1992), p. 279.
14. G. C. Chang, A Monte Carlo PDF/all-speed finite-volume study of turbulent flames, Ph.D. thesis, Cornell University, 1995 (unpublished).
15. J. J. Monaghan, Smoothed particle hydrodynamics, *Annu. Rev. Astronom. Astrophys.* **30,** 543 (1992).
16. J. J. Monaghan, Why particle methods work, *SIAM J. Sci. Stat. Comput.* **3,** 422 (1982).
17. S. B. Pope, Particle method for turbulent flows: Integration of stochastic differential equations, *J. Comput. Phys.* **117,** 332 (1995).
18. D. C. Haworth and S. B. Pope, A pdf modeling study of self-similar turbulent free shear flows, *Phys. Fluids* **30,** 1026 (1987).
19. S. B. Pope, On the relationship between stochastic Lagrangian models of turbulence and second-moment closures, *Phys. Fluids* **6,** 973 (1994).
20. M. S. Anand and S. B. Pope, Diffusion behind a line source in grid turbulence, in *Turbulent Shear Flows,* Vol. 4 (Springer-Verlag, Berlin, 1985), p. 46.
21. B. J. Delarue and S. B. Pope, Stochastic model for compressible turbulence, *Bull. Am. Phys. Soc.* **39,** 1860 (1994).
22. G. N. Mil'shtein, A method of second-order accuracy integration of stochastic differential equations, *Theory Probab. Appl.* **23,** 396 (1978).
23. G. N. Mil'shtein, Weak approximation of solutions of systems of stochastic differential equations, *Theory Probabl. Appl.* **30,** 750 (1985).
24. C. Hirsch, *Numerical Computation of Internal and External Flows,* Vol. 2 (Wiley, New York, 1990).
25. A. Jameson, Steady-state solution of the Euler equations for transonic flow, in *Transonic, Shock, and Multidimensional Flows: Advances in Scientific Computing* (Academic Press, New York, 1982), p. 37.
26. P. Campbell, Some new algorithms for boundary value problems in smooth particle hydrodynamics (unpublished).
27. A. Monin and A. Yaglom, *Statistical Fluid Mechanics,* Vol. 2 (MIT Press, Cambridge, 1975).
28. P. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations* (Springer-Verlag, Berlin, 1992).