DISCLAIMER

Stochastic Algorithms for the Analysis of Numerical Flame Simulations

John B. Bell, Marcus S. Day, Joseph F. Grcar and Michael J. Lijewski
Computing Sciences Directorate
Lawrence Berkeley National Laboratory
Berkeley, California 94720

Stochastic Algorithms for the Analysis of Numerical Flame Simulations

Abstract

Recent progress in simulation methodologies and new, high-performance parallel architectures have made it is possible to perform detailed simulations of multidimensional combustion phenomena using comprehensive kinetics mechanisms. However, as simulation complexity increases, it becomes increasingly difficult to extract detailed quantitative information about the flame from the numerical solution, particularly regarding the details of chemical processes. In this paper we present a new diagnostic tool for analysis of numerical simulations of combustion phenomena. Our approach is based on recasting an Eulerian flow solution in a Lagrangian frame. Unlike a conventional Lagrangian viewpoint in which we follow the evolution of a volume of the fluid, we instead follow specific chemical elements, e.g., carbon, nitrogen, etc., as they move through the system. From this perspective an "atom" is part of some molecule that is transported through the domain by advection and diffusion. Reactions cause the atom to shift from one species to another with the subsequent transport given by the movement of the new species. We represent these processes using a stochastic particle formulation that treats advection deterministically and models diffusion as a suitable random-walk process. Within this probabilistic framework, reactions can be view as a Markov process transforming molecule to molecule with given probabilities. In this paper, we discuss the numerical issues in more detail and demonstrate that an ensemble of stochastic trajectories can accurately capture key features of the continuum solution. We also illustrate how the method can be applied to studying the role of cyano chemistry on NO_x production in a diffusion flame.

Introduction

Recent progress in simulation methodologies and high-performance architectures make feasible detailed simulations of multidimensional combustion phenomena using comprehensive kinetics mechanisms. Smooke and his co-workers [1–5] have performed a number of studies of laminar methane diffusion flames with detailed kinetics. Sullivan et al. [6] have studied nitrogen chemistry in ammonia-enriched methane flames. For premixed flames, Najm and co-workers [7–9], and Bell et al. [10], have studied vortex flame interactions with detailed methane chemistry. Baum et al. [11] have studied two-dimensional turbulent flame interactions for detailed hydrogen chemistry, and Haworth et al. [12] have used a detailed propane chemistry model to simulate

propane–air flames. More recently Tanahashi et al. [13] have performed direct numerical simulations of turbulent, premixed hydrogen flames in three dimension with detailed hydrogen chemistry.

As simulation complexity increases, interpretation of computational results becomes increasingly difficult. Mean and fluctuating statistics can provide basic information about the flame. Probability distribution functions provide a more detailed view of the flame and begin to elucidate flame morphology in greater detail. These approaches provide information about basic flame structure and about the fluid dynamical behavior of the flame. However, understanding chemical activity against the backdrop of multidimensional fluid flow is difficult, even for relatively simple laminar flames.

This paper introduces a new diagnostic procedure for analyzing reacting flow simulation results with an emphasis on following chemical behavior within the system. The methodology is based on tracking individual "atoms" through the system, tabulating their trajectories and the chemical reactions they undergo. We adopt a Lagrangian viewpoint and use the Lagrangian form of the flow equations to derive a stochastic particle algorithm to track atoms through the system. The development and application of Lagrangian methods to reacting flow simulations remains an active area of research. We refer the reader to the survey article of Givi [14] for a discussion of vortex methods for reacting flow. In the present paper, however, we assume that the computation has been performed in an Eulerian frame and use the Lagrangian formulation for diagnostics.

Diagnostic Algorithm

Here we discuss the basic diagnostic model for tracking "atoms" through the system. In essence, we are looking for a numerical techniques for tagging a given atom (or collection of atoms) and monitoring their path through the flow domain and the history of which molecules transport them. The data to be analyzed is obtained by numerical solution of the reacting Navier-Stokes equations. The development is independent of whether the solution is obtained with a compressible or a low Mach number formulation. For clarity of exposition, we will assume that differential diffusion of species is given by a mixture model. Thus, the transport of species is given by the equation

$$\frac{\partial \rho Y_k}{\partial t} + \nabla \cdot u \rho Y_k = \nabla \cdot \rho D_k \nabla Y_k + \rho \omega_k \,, \tag{1}$$

where Y_k is the mass fraction of molecule M_k, ρ is density, u is velocity, and D_k and ω_k are the diffusion coefficient and production rate for M_k. Since we know the numerical solution, all of these fields are specified as functions of space and time. Rewriting this

equation in advective form using the mass conservation equation $\rho_t + \nabla \cdot (\rho U) = 0$ we obtain

$$\frac{D\,Y_k}{D\,t} = \frac{\partial Y_k}{\partial t} + u \cdot \nabla Y_k = \frac{1}{\rho} \nabla \cdot \rho D_k \nabla Y_k + \omega_k \,. \tag{2}$$

This equation provides a Lagrangian view of the dynamics of molecules of species k being transported along particles paths with their relative proportion being modulated by diffusion and reaction. To track the progress of a specific atom through the system, we need to interpret (2) from the perspective of atoms in the system. To make this notion more precise, suppose we are interested in the behavior of atoms A of some specific element in a subset of the molecules $M_1, M_2, \ldots, M_m$ that describe the chemical system. At any given time, A is part of one of these molecules and its motion is specified by the movement of that molecule. If we define x_A to be the location of A and assume that A is currently part of molecule M_k (which we indicate by $A \in M_k$), then in the absence of reactions, the movement of A can be described by the stochastic differential equation

$$dx_A = u(x_A, t)dt + dW_k(x_A, t) \,, \tag{3}$$

where dW_k denotes a generalized Wiener measure that determines a suitable Brownian motion with properties chosen to represent the diffusion of M_k. (Characterizing diffusion as a random walk in Lagrangian numerical methods is a common approach, cf. [15].)

Rather than attempt to construct an analytic form of the random walk to model species diffusion, we introduce a spatial scale Δx and a temporal scale Δt and use a lattice model for the random walk. To this end, we consider a discretization of the diffusion part of equation (2) in one dimension of the form

$$Y_k^{n+1} = Y_k^n + \frac{\Delta t}{\Delta x^2 \rho_k^n} \left[(\rho D)_{k+\frac{1}{2}}^n (Y_{k+1}^n - Y_k^n) + (\rho D)_{k-\frac{1}{2}}^n (Y_{k-1}^n - Y_k^n) \right] \,, \tag{4}$$

where the subscript ± 1 represents a right or left shift by Δx, with half-Δx shifts defined analogously. For Δt sufficiently small, we can collect terms and rewrite this equation as

$$Y_k^{n+1} = (1 - p_R - p_L)Y_k^n + p_R Y_{k,+1}^n + p_L Y_{k,-1}^n \tag{5}$$

where $p_R = \Delta t(\rho D)_{k+\frac{1}{2}}^n / \Delta x^2 \rho_k^n$ with p_L defined similarly. In this form we can interpret p_R as the probability that a molecule of species k shifts right Δx in time Δt, p_L as probability that it shift left and $(1 - p_R - p_L)$ as the probability that it does not shift.

We can then define a method to solve equation (3) by defining $x_p^* = x_p^n + \Delta t\, u$. We then choose a random number $\alpha \in [0, 1]$ and define

$$x_p^{n+1} = \begin{cases} x_p^* + \Delta x & \text{if } 0 \leq \alpha \leq p_R, \\ x_p^* - \Delta x & \text{if } p_R < \alpha \leq p_R + p_L, \\ x_p^* & \text{if } p_R + p_L < \alpha \leq 1. \end{cases}$$

For $\Delta t\, D_{max}/\Delta x^2 \ll 1$ the lattice approximation provides sufficient accuracy that statistical sampling error dominates errors arising from the lattice approximation. The generalization of this approach to two and three space dimension is straightforward.

We now need to augment this procedure to include chemical reactions. In keeping with the probabilistic framework, we seek a stochastic view of the chemical kinetics. This is most easily done by introducing a time interval Δt_c over which we wish to model the chemistry. If $A \in M_k$ at (x_A, t), there are a collection of reactions, $R_1, \ldots, R_N$, that transform M_k along with other reaction participants into a new collection of molecules. As a result of this transformation M_k is destroyed and A is transferred from M_k to M_{k_n} by reaction R_n. This destruction of M_k is expressed at the continuum level as

$$\frac{d[M_k]}{dt} = -\sum_{n=1}^{N} R_n \,. \tag{6}$$

where $[M_k]$ represents the molar concentration of M_k. (Reactions representing creation of M_k are not considered because these reactions do not affect A.) Assuming Δt_c is sufficiently small, (6) can be approximated by

$$[M_k]^{n+1} = [M_k]^n - \Delta t_c \sum_{n=1}^{N} R_n \,.$$

In this form we note that $\Delta t_c \sum_{n=1}^{N} R_n$ represents the amount of M_k that is destroyed by reaction in the time interval Δt_c. Motivated by this observation, we can rewrite this equation as

$$[M_k]^{n+1} = [M_k]^n (1 - \Delta t_c \sum_{n=1}^{N} R_n/[M_k]) \,.$$

and then define $p_n = \Delta t_c R_n/[M_k]$ to be the probability that reaction R_n transforms A from M_k to M_{k_n} during the time interval Δt_c. We also define $p_0 = (1 - \Delta t_c \sum_{n=1}^{N} R_n/[M_k]) > 0$ which is the probability that the molecule M_k containing A does not react during the time interval. Thus, from a finite time perspective we can represent the transfer of A from molecule to molecule as a result of chemical reactions as a Markov process $\mathcal{M}$.

There are several subtleties associated with constructing the Markov process $\mathcal{M}$. First, although most of the data required to construct $\mathcal{M}$ can be obtained from a standard CHEMKIN reaction file, some additional detail may be required. Issues arise when molecules in the reaction contain more than one atom of element A. When the molecule is symmetric with respect to A, we can assign probabilities based on simple counting arguments. The issue is more difficult when the structure of M is asymmetric with respect to the A's. In this situation, when M reacts, we need to know which position (up to symmetry) A occupies and we must augment the reaction probabilities to reflect the location of A in M. If the target molecule is asymmetric

in A, we must assign probabilities for the location of A in the new molecule and view different possible locations in the new molecule as distinct events.

Two other issues that arise are treating multiple reactions going from $M_k \overset{\mathcal{M}}{\to} M_{k'}$, and dealing with reversible reactions. In defining the algorithm, one can amalgamate all of the reactions that give the same transformation of A into a single transition probability. Similarly, one can lump forward and backward rates for reversible reactions into a single net rate for each cell. In the present implementation we have done neither. Each reaction has its own probability so we can identify, for each molecular transformation, the reaction that produced the change. For reversible reactions, we allow both forward and backward reactions so that it is possible for an atom to jump back and forth between to molecules several times within a cell. These choices represent the most general form of the algorithm; however, in some cases, another choice may be appropriate. This issue will be explored in future work.

Combining the two stochastic process, we can now define a probabilistic algorithm for tracing the path of A through the domain. First, we identify Δt for transport and Δt_c for chemistry. Here, we have chosen Δt_c and Δt so that the probabilities of reacting or taking a diffusion step are less than 10%. For stiff chemistry a suitable Δt_c is potentially very small so we subcycle the Markov process, $\mathcal{M}$, with respect to the other processes. Thus, the basic stochastic particle algorithm is, given x_A with $A \in M_k$, to first update

$$x_A^{n+1} = x_A^n + \Delta t \, u(x_A, t) + dW_k^{\Delta t, Lat}(x_A^n + \Delta t u, t)$$

where $dW_k^{\Delta t, Lat}$ is the lattice random walk algorithm described above. Then we apply the Markov process $\mathcal{M}^{\Delta t_C}$ r times until $r\Delta t_c = \Delta t$.

To analyze a numerical simulation, we apply this algorithm to an ensemble of trajectories with initial starting location and initial molecule based on the question being considered. Before illustrating the performance of the method we first make a couple of observations about the method. First, although at first examination it would appear that the computations are quite costly, in fact, since most of the required data can be precomputed, the algorithm can be implemented with minimal computational cost. Also, since each trajectory is independent, the method parallelizes very well. Exploiting these characteristics of the method, we have been able to follow a million trajectories in just a few hours.

A second observation about this approach is a word of caution. With the stochastic description of the algorithm it is tempting to try to relate the stochastic particles to a Boltzmann description of the fluid. However, unlike a Boltzmann description we do not maintain velocity distribution for the particles. Thus, our particles represent, at best, "pseudo-particles" whose velocities are averages of the Boltzmann velocity description. Perhaps the most accurate description of an ensemble of stochastic trajectories is as a path integral representation of an approximation to the continuum solution.

Computational Results

In this section we demonstrate the behavior of the method and illustrate the relations between the particle trajectories and the continuum solution. The examples consider a laminar, steady, diffusion flame in cylindrical geometry. This flame was modeled using a reaction mechanism due to Glarborg [16] with a low Mach number adaptive mesh refinement algorithm [17]. This mechanism, which includes detailed nitrogen chemistry, contains 65 species and 447 reactions. The case we consider here was part of a combined experimental and numerical study of the effect of fuel-bound nitrogen in the form of NH_3 on NO_x formation [6, 18]. Here we consider only the case with no added NH_3. Temperature and X_{NO} obtained from the continuum solution are shown in Fig. 1.

For all of the examples we interpolate the solution to the finest level of the adaptive mesh hierarchy and evaluate the reaction rates and the diffusion coefficients on the resulting uniform grid. We then use the mesh spacing of this grid to define the lattice spacing for the random walk. Neither of these choices represent inherent limitations of the approach.

Dominant, Carbon Chemistry

Our first example traces the carbon chemistry of the system. For this problem, all of the carbon enters the system as CH_4 in the inflow fuel stream. From the inflow velocity profile and composition, we construct a probability measure on the inflow radial location that uniformly samples the molar flux of CH_4 into the system. We create an ensemble of 10^6 trajectories and tabulate the molecular transformations given by $\mathcal{M}$. From these transformations, we compute the net flow of carbon through each edge of the carbon reaction network. Suitably scaled, this collection of net transformations corresponds to a reaction path diagram for the carbon chemistry, which is depicted graphically in Fig. 2. For steady flow we can also perform a reaction path analysis by integrating the chemical rates over the domain. In Table 1 we compare the strengths of the reaction pathways from the particle simulation to the comparable data computed directly from the continuum data. For most of the edges the stochastic simulation closely matches the result from the continuum solution. The only major discrepancy is the strength of the $CO \rightarrow CO_2$ edge. The stochastic algorithm predicts more net reaction than is indicated by the continuum integration. Preliminary investigations suggest that this discrepancy reflects underlying splitting errors in the methodology used for the simulation. For the remaining edges, four have errors in the 1–3% range with the remaining errors less than 1%, even for relatively weak edges.

Trace, Nitrogen Chemistry

An important scientific objective in studying these types of flames is to study the formation of NO_x. For the case we are considering the only source of nitrogen is N_2 in both the fuel and the oxidizer streams. Because N_2 is a relatively stable molecule, the number of trajectories in which nontrivial reactions occur is very small; we are essentially looking for rare events. Although we can examine NO_x chemistry as before by sampling N in N_2 molecules entering the domain, we would require a large number of trajectories to obtain a statistically significant set of interesting trajectories. Alternatively, we can *a priori* decide to look only at "interesting" trajectories; i.e., trajectories where N_2 reacts. This is done by using the continuum reaction rates for N_2 over the domain to construct a probability distribution that reflects where N_2 will first react. We use one random variable to sample this distribution for points in space at which to begin the trajectories, and then we use a second random variable to sample the distribution of N_2 reactions at such points for the initiating reactions. This type of procedure is similar to stochastic models for studying rare reaction phenomena in biological models originally developed by Gillespie [19, 20]. We again simulate 10^6 trajectories, in this case for the N atom, beginning with the an initial breakup of N_2. We again use the particle trajectories to compute a net reaction path diagram for the nitrogen chemistry which is presented in Fig. 3. As before, we can compute the analogous net reaction graph from the continuum data. A comparison of the resulting edge weights is given in Table 2. The stochastic algorithm captures most of the reactions to within a few percent. The only exceptions are that the stochastic algorithm predicts an increased level of the $NO \rightarrow HONO \rightarrow NO_2$ loop and a decrease in $N_2 \rightarrow NNH$ as compared to the continuum solution. These cases correspond to situations in which forward and reverse reaction rates are quite large and suggest that we have insufficient statistics to accurately capture the net.

Species Concentrations

The above tests illustrate the ability of the stochastic particles to capture the chemical kinetics in the system. We now test the ability of the stochastic particles to represent the spatial structure of the continuum solution. Assume that we are given a lattice that covers the computational domain. For each cell in the lattice and for each trajectory that crosses that cell, we determine the residence time of N while it is part of an NO molecule in that cell. If we sum these residence times over an ensemble of trajectories the result is proportional to the molar concentration of NO. In Fig. 4, we show the continuum $r \cdot [NO]$ compared to the residence time for an increasing number of particles in the ensemble. (The continuum data is scaled by radius because the residence time is in units proportional to moles/lattice cell not moles/volume.) Even for a modest number of particles the residence time provides a

reasonable, if somewhat ephemeral, view of the concentration profile. As the number of samples is increased, however, the agreement becomes increasingly good. We also note that better agreement can be obtained with fewer trajectories if we use a coarser lattice.

Carbon-Nitrogen Chemistry

The examples presented thus far evaluate diagnostics that are easily obtained from the continuum solution; we have examined them to provide a measure of validation for the stochastic particle method. As a final illustration of the use to stochastic particles, we examine an issue for these types of flames that is not easily determined from the continuum solution. To pose the question, we consider the nitrogen reaction path analysis for this flame presented in Fig. 3. The N in nitrous oxides leaving the system enters the system as N_2 which is broken into N, NNH, HCN, etc. and eventually exits the domain as either NO, or NO_2. The path diagram shows a loop in which nitrogen atoms reside for a time in carbon species. Indeed, the flow through some carbon species is greater than that out of N_2. This indicates that carbon chemistry plays an important role in the formation of NO_x with an N atom possibly recycling through the carbon species multiple times before exiting the domain in NO_x.

To understand the role of this carbon recycling on NO_x chemistry we examine the trajectories used to compute Fig. 3's nitrogen reaction path diagram and extract the subset of those paths that exit the domain as NO or NO_2. For each of these trajectories, we calculate the number of times the N atom we are tracking changes from a non-carbon containing species to a carbon species, and refer to this quantity as the number of cycles for that trajectory. We can then compute a probability distribution for the number of carbon recycling cycles undergone by NO and NO_x molecules leaving the system. The resulting data, presented in Fig. 5, are well approximated by the discrete geometric probability distribution

$$P(n) = \lambda(1 - \lambda)^n \quad \text{for } \lambda = 0.33$$

indicating that entering the recycling loop can be modeled as the arrival time for a Bernoulli process.

We can also use the particle trajectories to quantify the spatial structure of the recycling behavior. In Fig. 6 we plot the initial reaction location for each particle that exits the domain as NO or NO_2. We can see that particles that initially react on the outer edge of the flame are not affected by carbon recycling. However, considerable carbon recycling occurs for trajectories initiating on the rich side of the flame sheet, and it becomes increasingly important as we approach the base of the flame. This provides some quantification of the overall behavior of the system and allows us to obtain a spatial picture that indicates where carbon chemistry plays an important role in NO_x formation.

Summary and Conclusions

We have developed a new diagnostic methodology for analyzing combustion simulations. This approach is based on tracking atoms through the system using a stochastic particle formulation that models advective transport, differential diffusion and reactions using the results of a pre-existing solution to the reacting flow equations. We have demonstrated that the method can recover key properties of the continuum solution and provides a mechanism for diagnosing the behavior of complex reacting flows. Perhaps, one of the most intriguing properties of this approach is that is allows questions about the chemistry and the flow to be posed in a natural and straightforward manner. We are currently working to extend our implementation to directly access the hierarchical grid system generated by our block-structured adaptive reacting flow algorithm with the goal of using stochastic particles to study chemical behavior in three-dimensional, time-dependent turbulent combustion.

References

[1] M. D. Smooke, A. Ern, M. A. Tanoff, B. A. Valdati, R. K. Mohammed, D. F. Marran, and M. B. Long. Computational and experimental study of NO in an axisymmetric laminar diffusion flame. *Proc. Comb. Inst.*, 26:2161–2170, 1996.

[2] B. A. V. Bennett and M. D. Smooke. A comparison of the structures of lean and rich axisymmetric laminar bunsen flames: application of local rectangular refinement solution-adaptive gridding. *Combust. Theory Modelling*, 3:657–687, 1999.

[3] R. H. Mohammed, M. A. Tanoff, M. D. Smooke, and A. M. Schaffer. Computational and experimental study of a forced, time-varying, axisymmetric, laminar diffusion flame. In *Proc. Comb. Inst.*, volume 27, pages 693–702, 1998.

[4] B. A. V. Bennett, C. S. McEnally, L. D. Pfefferle, and M. D. Smooke. Computational and experimental study of axisymmetric coflow partially premixed methane/air flames. *Combust. Flame*, 123:522–546, 2000.

[5] B. A. V. Bennett, C. S. McEnally, L. D. Pfefferle, M. D. Smooke, and M. B. Colket. Computational and experimental study of axisymmetric coflow partially premixed ethylene/air flames. *Combust. Flame*, 127:2004–2022, 2001.

[6] N. Sullivan, A. Jensen, P. Glarborg, M. S. Day, J. F. Grcar, J. B. Bell, C. Pope, and R. J. Kee. Ammonia conversion and NO_x formation in laminar coflowing nonpremixed methane-air flames. *Combust. Flame*, 2001. Submitted.

[7] H. N. Najm and P. S. Wyckoff. Premixed flame response to unsteady strain rate and curvature. *Combust. Flame*, 110(1–2):92–112, 1997.

[8] H. N. Najm, O. M.. Knio, P. P. Paul, and P. S. Wyckoff. A study of flame observables in premixed methane-air flames. *Combust. Sci. Technol.*, 140:369–403, 1998.

[9] H. N. Najm, P. P. Paul, C. J. Mueller, and P. S. Wyckoff. On the adequacy of certain experimental observables as measurements of flame burning rate. *Combust. Flame*, 113(3):312–332, 1998.

[10] J. B. Bell, N. J. Brown, M. S. Day, M. Frenklach, J. F. Grcar, and S. R. Tonse. The effect of stoichiometry on vortex flame inter actions. *Proc. Comb. Inst.*, 28:1933–1939, 2000.

[11] M. Baum, T. J. Poinsot, D. C. Haworth, and N. Darabiha. Direct numerical simulation of $H_2/O_2/N_2$ flames with complex chemistry in two-dimensional turbulent flows. *J. Fluid Mech.*, 281:1–32, 1994.

[12] D. C. Haworth, R. J. Blint, B. Cuenot, and T. J. Poinsot. Numerical simulation of turbulent propane-air combustion with nonhomogeneous reactants. *Combust. Flame*, 121:395–417, 2000.

[13] M. Tanahashi, M. Fujimura, and T. Miyauchi. Coherent fine scale eddies in turbulent premixed flames. In *Proc. Comb. Inst.*, volume 28, pages 529–535, 1998.

[14] P. Givi. Spectral and random vortex methods in turbulent reacting flow. In P. A. Libby and F. A. Williams, editors, *Turbulent Reacting Flows*, pages 475–572. Academic Press, 1994.

[15] A. J. Chorin. Numerical study of slightly viscous flow. *J. Fluid Mech.*, 57:785–796, 1973.

[16] P. Glarborg, P. G. Kristensen, K. Dam-Johansen, M. U. Alzueta, A. Millera, and R. Bilbao. Nitric oxide reduction by non-hydrocarbon fuels. Implications for reburning with gasification gases. *Energy Fuels*, 14:828–838, 2000.

[17] M. S. Day and J. B. Bell. Numerical simulation of laminar reacting flows with complex chemistry. *Combust. Theory Modelling*, 4:535–556, 2000.

[18] J. B. Bell, M. S. Day, J. F. Grcar, W. G. Bessler, C. Schultz, P. Glarborg, and A. D. Jensen. Detailed modeling and laser-induced fluorescence imaging of nitric oxide in a NH_3-seeded non-premixed methane/air flame. *Proc. Comb. Inst.*, 29, 2002. Submitted.

[19] D. T. Gillespie. Stochastic simulation of chemical processes. *J. Comput. Phys.*, 22:403, 1976.

[20] D. T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.*, 81:2340–2361, 1977.

Figure 1: *Temperature and NO mole fraction for the laminar nonpremixed flame. The dotted white line is the stoichiometric boundary between the fuel and air.*


Figure 2: *Reaction path diagram for carbon chemistry. Only edges at least 3% of the strongest are shown.*

Table 1: *Comparison of edge weights for a reaction path diagram based on carbon (C) as determined from a continuum simulation and as reproduced by the stochastic particle method. Edges are ordered by their stochastic-particle weight; only those at least 1% of the largest are shown.*

edge	stoch.	cont.
CH4 → CH3	100.0	100.0
CO → CO2	98.9	92.5
HCO → CO	63.4	64.0
CH2O → HCO	60.1	60.5
C2H4 → C2H3	44.7	42.5
C2H3 → C2H2	40.2	37.6
C2H5 → C2H4	37.0	36.5
CH2OH → CH2O	26.6	25.8
C2H2 → HCCO	24.3	22.3
CH3 → CH2OH	24.1	23.1
CH3 → C2H6	19.6	20.0
C2H6 → C2H5	19.6	20.1
CH3 → CH2O	19.3	21.3
CH3 → C2H5	17.4	16.4
HCCO → CO	13.5	12.6
CH2(S) → CH2	11.1	10.7
CH3 → CH2(S)	8.5	8.9
HCCO → CH2(S)	7.8	7.2
HCCO → C2O	7.0	6.5
CH2 → C2H4	6.6	6.1
CH2 → CH2O	6.5	6.1
C2O → CO	6.5	6.0
CH3 → C2H4	5.9	5.3
C2H2 → CO	5.4	4.9
C2H2 → HCCOH	5.1	5.1
HCCOH → HCCO	5.1	5.0
CH2 → CH	5.1	4.6
C2H2 → CH2	4.9	4.4
CH3 → CH2	4.7	4.4
C2H4 → C2H2	3.5	3.6
C2H3 → CH2HCO	3.4	3.7
CH → CH2O	3.4	3.2
CH3 → CH3OH	2.6	2.8
CH2(S) → CO	2.5	3.0
CH2(S) → CH2O	2.2	2.1
CH2HCO → CO	2.2	2.3
CH2HCO → CH3	2.1	2.3
CH3OH → CH2OH	2.1	2.3
CH2 → CO	2.1	2.3
C2H2 → CH2CO	1.9	1.9
CH3O → CH2O	1.3	1.5
CH2CO → CH3	1.2	1.2
CH → HCO	1.2	1.2
CH2CO → CO	1.2	1.2
C2H2 → C2H	1.1	1.0
C2H4 → CH2HCO	1.0	1.1

Table 2: *Comparison of edge weights for a reaction path diagram based on nitrogen (N) as determined from a continuum simulation and as reproduced by the stochastic particle method. Edges are ordered by their stochastic-particle weight; only those at least 1% of the largest are shown.*

| edge	stoch.	cont.
NO → HCN	100.0	100.0
N → NO	97.2	95.9
NH → N	76.9	75.9
NCO → NH	71.9	70.3
HCN → NCO	63.0	60.7
HNCO → NH2	50.1	50.8
NCO → HNCO	46.2	46.6
NH2 → NH	45.9	46.6
HNO → NO	42.4	41.8
NO → HONO	41.0	32.3
HONO → NO2	40.2	31.9
CN → NCO	39.3	38.9
NH → NO	37.1	35.2
NH → HNO	37.0	36.7
HOCN → NCO	31.9	32.7
HCN → HOCN	30.0	30.6
N2 → N	29.7	30.8
N2 → NNH	29.0	40.1
H2CN → HCN	24.8	25.9
HCN → CN	24.1	23.3
NO2 → NO	23.9	23.1
N2 → HCN	23.1	24.3
NNH → N2O	19.2	17.7
N → H2CN	18.4	19.0
HCN → NH	16.2	15.7
NCO → NO	14.8	14.1
N2O → NO	14.7	12.5
HCN → CH3CN	13.3	13.9
N2 → N2O	12.6	8.7
CH3CN → CH2CN	11.4	11.6
CH2CN → CN	11.4	11.3
N2O → NH	11.2	8.8
NO → CN	11.2	10.9
CN → N	8.9	8.3
NO → H2CN	6.4	6.9
NH2 → HNO	5.2	5.0
HCN → HNCO	5.2	5.6
NNH → NO	4.9	5.4
N2 → NO	4.7	4.7
NNH → NH	4.5	5.3
HCNO → HCN	3.9	4.0
NO → HCNO	3.9	4.1
CH3CN → HOCN	1.9	2.1
N2 → CN	1.5	1.7
HNCO → NH	1.3	1.2
NCO → N	1.1	1.2
N → HCN	1.1	1.2
HCN → NH2	1.0	1.0

Figure 3: *Reaction path diagram for nitrogen chemistry. Only edges at least 3% of the strongest are shown.*


Figure 4: *NO concentration in moles/area comparing 32,000, 128,000 and 1,000,000 particles with the continuum solution.*

Figure 5: *Distribution of carbon cycles in* NO_x *chemistry.*

Figure 6: *Location of initial reaction for exiting* NO_x *particles. Black points mark initial reaction locations for particles that do not participate in carbon recycling. The remaining points are color-coded to indicate the number of carbon cycles in the particle history using a rainbow palette ranging from red (= 1) to violet (≥ 12).*