A New Approach to the Solution of Large Thermonuclear Burning Networks

Mike Guidry∗**Univ. Tennessee/Oak Ridge National Lab**

E-mail: `guidry@utk.edu`

A new algorithm is proposed for solving large networks of stiff coupled differential equations in various scientific applications. The algorithm replaces differencing of abundance variables with evolution of discrete populations. It reproduces quantitatively the results of standard methods, but with important advantages. (1) The algorithm is explicit, yet it decouples accuracy from stability for stiff systems, permitting explicit integration with a timestep set by the former rather than the latter, thereby avoiding implicit solves. (2) It exploits sparseness perfectly, computing only those reaction links that the physical system traverses. (3) It scales linearly with the number of couplings for large sparse networks, in contrast to the quadratic to cubic scaling of standard methods. (4) Unlike Monte Carlo, for large physical particle number, execution time is independent of the number of test particles, allowing even weaker populations to be tracked efficiently. (5) The decoupling of stability from accuracy allows stable tuning of large networks to optimize accuracy versus computational time. We propose that this new approach can be used to solve large, stiff networks for many complex systems, such as the coupling of realistic networks to multidimensional hydrodynamics, that tax the capability of standard methods.

∗Speaker.

1. Introduction

Phenomena in many scientific disciplines may be modeled by fluxes transferring population between sources and sinks for various species. Let us refer generically to these sources and sinks as *boxes,* and term the resulting systems of boxes connected by fluxes *reaction networks.* Of particular interest in the present context are thermonuclear reaction networks in astrophysics, but the method described here should be applicable to a much broader class of problems. For such systems one usually solves numerically a coupled set of N ordinary differential equations

$$\frac{dY_i}{dt} = \sum_j F_{ij},$$ (1.1)

where $Y_i (i = 1 \ldots N)$ describe the abundances, t is the the time, the fluxes between species i and j are given by F_{ij}, and the sum for each i is over all j coupled to i by a non-zero F_{ij}. Thus, Eqs. (1.1) describe a continuous and deterministic evolution of populations between the boxes.

2. The Explicit Stochastic Method

Nature does not know about differential equations like (1.1). Instead, new species are produced when initial species interact according to the physical theory governing the system and undergo reactions that alter populations. For microscopic systems these reactions and decays are inherently stochastic: in a thermonuclear reaction network, an individual nucleus of ^{13}N does not evolve continuously to other isotopes. Instead, in any time interval it has a certain probability to remain ^{13}N, a certain probability to beta decay, a certain probability to capture a proton and emit a gamma ray, and so on, with these (stochastic) probabilities governed by the rules of quantum mechanics.

The description implied by Eq. (1.1) may be converted into a stochastic description more in keeping with actual physical processes through the following steps. Assume the boxes i to have populations Y_i. In a small time Δt, the total number of particles transferred from box i to box j is given by $F_{ij}\Delta t$ and the total number of particles transferred to all boxes from box i is $\Delta Y_i = \sum_j F_{ij}\Delta t$, where we assume on physical grounds that $\sum_j F_{ij}\Delta t \leq N_i$. Therefore, the probability P_{ij} for a given entity in box i to make a transition to box j is

$$P_{ij} = \frac{F_{ij}\Delta t}{\sum_j F_{ij}\Delta t} = \frac{F_{ij}}{\sum_j F_{ij}},$$ (2.1)

and the sum of all such probabilities (including the probability of no transition) is unity, $\sum_k P_{ik} = 1$. This suggests the algorithm indicated schematically in Fig. 1 to evolve network populations.

3. The Large Particle Number Limit of the Stochastic Algorithm

The preceding discussion has introduced the stochastic method in a basic form that closely approximates the actual physical processes modeled by Eqs. (1.1). As a numerical method, the algorithm in this form works best for evolution of strong populations, since resolution of weak populations requires tracking large total numbers of test particles. In the (usually fulfilled) case
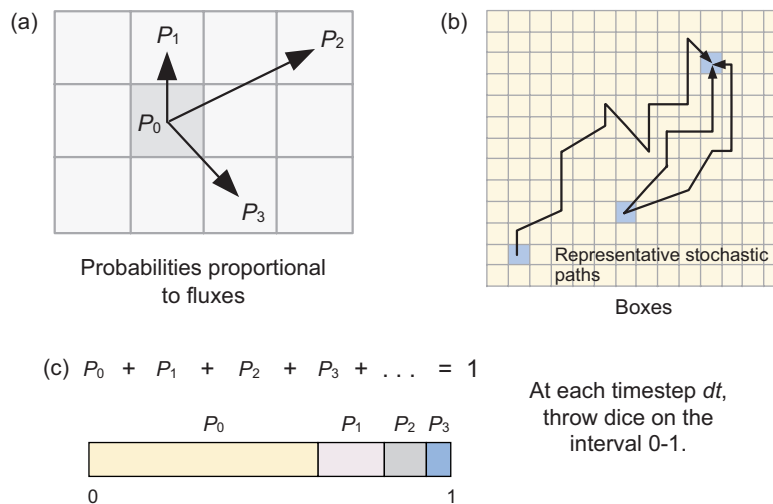
Figure 1: The stochastic algorithm.

corresponding to a macroscopic number of physical particles, it is no longer necessary to track the evolution of individual test particles explicitly by throwing dice because we already know the outcome of this exercise with extremely high certainty. In the large particle number limit, we may simply transfer whole blocks of test particles to daughter boxes in proportion to the fluxes connecting the parent and daughter boxes in a given timestep. Then the change in population N_i for box i in the time interval Δt is given by

$$\Delta N_i = \sum_j \Delta N_{ij} = \sum_j F_{ij}\Delta t. \tag{3.1}$$

However, because of the physical heritage of the model (evolution of populations), we cannot transfer particles that don't exist and therefore no negative fluxes are permitted: $F_{ij} \to \max(F_{ij},0) \equiv \tilde{F}_{ij}$. The corresponding algorithm is efficient and reproduces the results of a variety of test networks from various fields of science. Although it is an explicit algorithm (all information needed to advance a timestep is known at the beginning of the timestep, without iteration or matrix inversion), it has proven stable for applications in very stiff systems, as will be demonstrated below.

4. An Example: Hot CNO Burning under Nova Conditions

Figure 2(a) compares some representative stochastic populations under nova (hot CNO) conditions with the results of a standard implicit calculation. Note the almost perfect agreement between implicit and stochastic methods over six orders of magnitude in the mass fractions in Fig. 2(a). Figure 2(b) displays the fastest and slowest rates entering a representative stochastic nova simulation as a function of time. The difference of about 18 orders of magnitude between the fastest and slowest rates at any timestep is an indication that this is an extremely stiff system. For standard explicit algorithms, the largest timestep permitted by stiffness stability criteria generally is of order the inverse of the fastest rate in the network. For the calculations illustrated in Fig. 2(b), the inverse
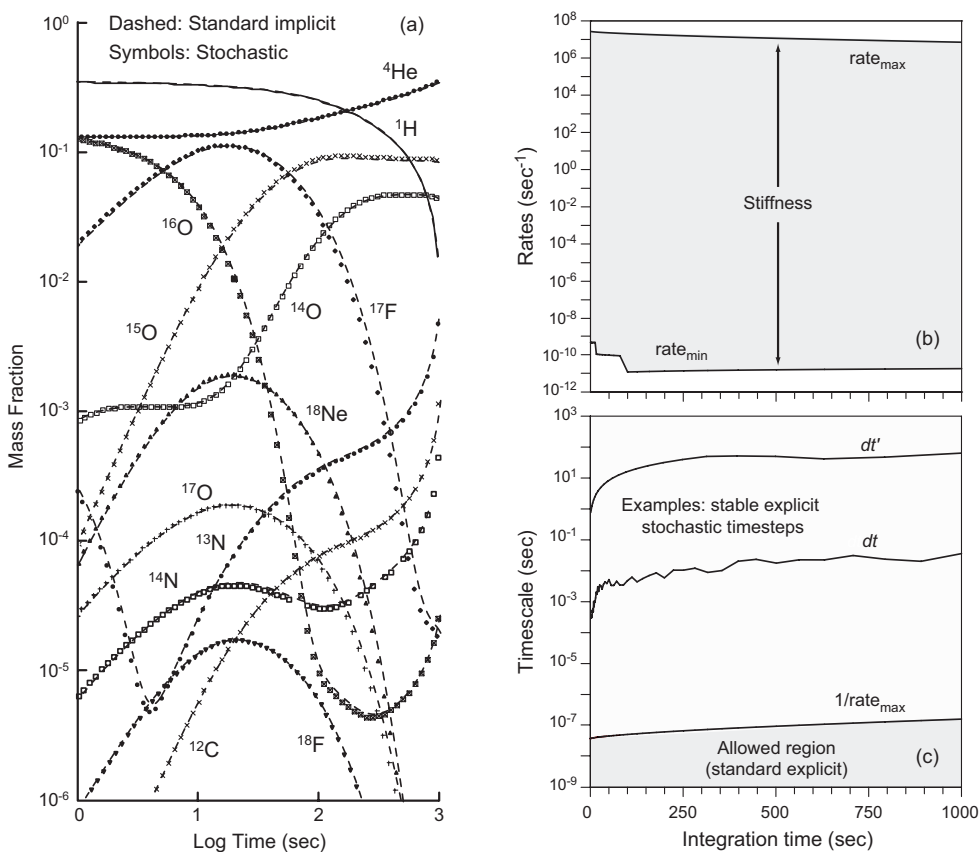
Figure 2: (a) Representative isotopic mass fractions under nova conditions computed with the explicit stochastic algorithm (symbols), compared with an implicit backward Euler calculation (dashed lines). The network contained 145 isotopes, with 924 non-zero couplings. For simplicity in interpretation, constant temperature and density profiles were assumed. (b) Rates and timescales characteristic of a stochastic nova simulation. Conditions as for part (a) but with a larger reaction library: 896 isotopes with 8260 couplings were included but the stochastic algorithm traverses only the non-zero flux links in any timestep. (c) Comparison of maximum stable timestep [$\simeq 1/\text{rate}_{max}$ from part (b)] possible for a standard explicit integration with much larger stable timesteps dt and dt' for some representative explicit stochastic integrations.

of the fastest rate gives the lower curve in Fig. 2(c). Thus a normal explicit algorithm would be restricted by stability requirements to timesteps lying approximately in the shaded region below this curve ($dt \simeq 10^{-7}$ seconds or less).

In contrast, we show two curves in Fig. 2(c) for stable stochastic integration (adaptive) timesteps lying far above this region. The curve marked dt is for a timestep small enough to give accuracy comparable to Fig. 2(a). This timestep is some 10^5 times larger than would be stable for a normal explicit integration. The curve marked dt' is for a much larger timestep that compromises accuracy for the weaker transitions but remains stable and calculates stronger transitions correctly. The timestep $dt' \sim 100$ seconds is about 10^9 times larger than would be stable for a standard explicit algorithm. Since this is comparable to the characteristic timescale ($\sim 10^3$ seconds) of the physical event being simulated, this stable explicit integration timestep is effectively arbitrarily large with
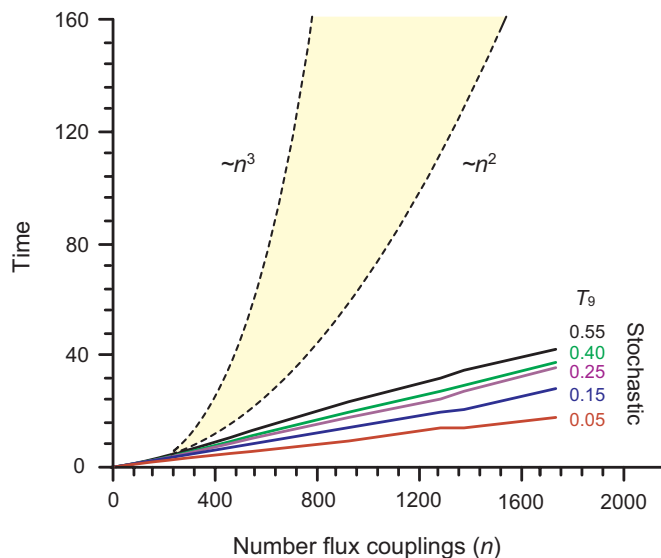
Figure 3: The scaling of execution time with the size of the network for the explicit stochastic algorithm.

respect to the usual upper limit for explicit methods.

5. Scaling with Network Size

Implicit methods scale from quadratically to cubically with network size. The structure of the explicit stochstic algorithm implies linear scaling for large sparse networks. We give an example of actual scaling for explicit stochastic integration in Fig. 3, which shows the variation of time to execute the stochastic algorithm (taking the same constant timesteps for all cases) with the size of the network. These calculations assumed a constant density of 500 g cm^{-3} and constant temperatures ranging from 0.05×10^9 to 0.55×10^9 K. The scaling with respect to network size is indeed seen to be weakly linear. The slope of computing time versus number of couplings depends indirectly on the temperature because the algorithm exploits the sparseness optimally, calculating only those transition links corresponding to fluxes above a user-specified minimum value. For higher temperatures, the number of fluxes exceeding this threshold can increase and more calculations are required in each timestep.

6. Approach to Equilibrium

Perhaps the most difficult issue for stiff systems is the approach to equilibrium. Initial tests of the present explicit method in various astrophysics and geochemical applications indicate that it remains stable in the approach to chemical and nuclear statistical equilibrium. We are presently testing the method extensively for the nuclear statistical equilibrium conditions appropriate for Type Ia supernova explosions. If these tests are successful, the explicit stochastic method could represent a major advance in the solution of large stiff networks for astrophysics and a number of other disciplines.

7. Summary and Conclusions

We conclude by summarizing briefly the most important properties of the stochastic algorithm as applied to the solution of large, sparse reaction networks: (1) The algorithm decouples stability and accuracy issues for the stiff system. (2) The algorithm exploits the sparseness of the equivalent linear algebra system perfectly, since the population naturally follows only the paths of non-zero transition probability. (3) The scaling of the explicit stochastic algorithm is weakly linear with size of the network for sparse systems, in contrast to the quadratic to cubic scaling of implicit solvers. (4) In the large physical particle number limit, execution time is independent of the number of test particles (unlike Monte Carlo), allowing even weak populations to be tracked efficiently. (5) The algorithm allows a user-defined tradeoff of accuracy against computational time that is independent of stability issues. These characteristics suggest that the explicit stochastic algorithm can be used to solve a variety of large network problems in many scientific and technical fields that lie beyond the grasp of currently-used technologies.