Fast, scalable master equation solution algorithms. IV. Lanczos iteration with diffusion approximation preconditioned iterative inversion

Terry J. Frankcombe and Sean C. Smith

Articles you may be interested in

Statistically testing the validity of analytical and computational approximations to the chemical master equation
J. Chem. Phys. **138**, 204108 (2013); 10.1063/1.4807390

Nucleated polymerization with secondary pathways. II. Determination of self-consistent solutions to growth processes described by non-linear master equations
J. Chem. Phys. **135**, 065106 (2011); 10.1063/1.3608917

The dynamics of conformational isomerization in flexible biomolecules. II. Simulating isomerizations in a supersonic free jet with master equation dynamics
J. Chem. Phys. **120**, 148 (2004); 10.1063/1.1626541

Fast, scalable master equation solution algorithms. III. Direct time propagation accelerated by a diffusion approximation preconditioned iterative solver
J. Chem. Phys. **119**, 12729 (2003); 10.1063/1.1628213

Master equation analysis of intermolecular energy transfer in multiple-well, multiple-channel unimolecular reactions. I. Basic theory
J. Chem. Phys. **107**, 8904 (1997); 10.1063/1.475182

Fast, scalable master equation solution algorithms. IV. Lanczos iteration with diffusion approximation preconditioned iterative inversion

Terry J. Frankcombe[a)] and Sean C. Smith[b)]
Centre for Computational Molecular Science, Chemistry Building 68, University of Queensland, Brisbane, 4072, Australia

In this paper we propose a second linearly scalable method for solving large master equations arising in the context of gas-phase reactive systems. The new method is based on the well-known shift-invert Lanczos iteration using the GMRES iteration preconditioned using the diffusion approximation to the master equation to provide the inverse of the master equation matrix. In this way we avoid the cubic scaling of traditional master equation solution methods while maintaining the speed of a partial spectral decomposition. The method is tested using a master equation modeling the formation of propargyl from the reaction of singlet methylene with acetylene, proceeding through long-lived isomerizing intermediates. © *2003 American Institute of Physics.*
[DOI: 10.1063/1.1628214]

I. INTRODUCTION

The master equation (ME) formulation for solving gas-phase chemical kinetics problems is well-known and commonly employed.[1–5] While the smallest eigenvalue of the energy grained ME matrix is often all that is required in simpler applications, the transient behavior of the system commonly becomes the primary focus in more complex applications. Determining the transient behavior requires far more information than just the smallest eigenvalue and corresponding eigenvector (which only describe the long-time behavior), usually with a disproportionate increase in the amount of computational effort required. Determining transient behavior is particularly important in multi-well systems describing isomerization between a number of isomers, which are increasingly being modeled with ME methods.[6–13]

The matrices arising from multi-well MEs are significantly larger than unimolecular ME matrices, as are the matrices arising from two-dimensional (2-D) MEs resolved in angular momentum as well as energy.[14–19] One can easily construct a 2-D or multi-well ME discretized over tens of thousands of points. The potential also exists to construct 2-D multi-well MEs, with a corresponding further increase in the size of the discretization. For such large problems to be tractable requires an effective method requiring calculation time that is *scalable*. That is, one needs a method that not only works, but also one that does not result in a vast increase in the computational effort required for a modest increase in the size of the problem. Throughout this work, we use the term *scalable method* loosely to mean a method that scales significantly better than the cubic scaling of traditional ME solution methods. While scaling with the square of the size of the system/discretization falls into this loose categorization, we aim for global linear scaling. Specifically, we consider the matrix-vector product with its linear dependence on the number of isomers being modeled and, at worst, qua-

dratic dependence on the number of energy grains to be a scalable operation.

In this series of papers we have made significant progress toward general scalable ME solution methods. The first paper[20] presented one of the first successful scalable methods by generalizing the well-known and highly successful Nesbet method. However this method, denoted the HONE method, has the deficiency that some high-quality initial data must be available. Specifically, the eigenvalues and the relatively large magnitude elements of the eigenvectors must be known. In that regard, multi-well systems are particularly problematic as there may be multiple numerically difficult small eigenvalues. The second paper[21] served largely to set the scene for what was to come by surveying the effectiveness and speed of a number of standard methods used for solving multi-well MEs. The conclusions of the second paper were that the fastest method that was robust over all the temperature and pressure ranges tested was direct time propagation using a stiff ordinary differential equation (ODE) integration algorithm, while in high temperature cases a spectral approach using a shift-invert Lanczos iteration could be orders of magnitude faster. All of the methods tested in Paper II formally scale with n^3 [denoted as $\mathcal{O}(n^3)$ operations, where n is the size of the ME matrix] and are hence inappropriate for solving large problems.

In Paper III[22] we proposed a method that maintains the robustness of standard direct time propagation with a stiff ODE integrator while maintaining scalability for large problems. The component of a stiff ODE integrator which imposes the $\mathcal{O}(n^3)$ scaling on the calculation is solving the nonlinear correction equation by Newton's method. In the first-order linear case of the ME,

$$\frac{d\boldsymbol{\rho}}{dt} = B\boldsymbol{\rho}, \tag{1}$$

applying Newton's method to the correction equation requires solving the linear system,

$$(I - \gamma B)\mathbf{z} = \mathbf{b}, \tag{2}$$

[a)]Electronic mail: T.Frankcombe@chemistry.uq.edu.au
[b)]Electronic mail: S.Smith@uq.edu.au

12741

with arbitrary γ and $\mathbf{b}$. The standard approach is to factorize $I - \gamma B$ with a dense factorization before solving for $\mathbf{z}$—an $\mathcal{O}(n^3)$ operation and the lion's share of the computational effort in the numerical integration. In the work of Paper III, the direct solution of Eq. (2) was replaced by using a preconditioned iterative solver, with the preconditioning being provided by direct factorization of the matrix arising from the diffusion approximation to the ME. Thus the $\mathcal{O}(n^3)$ direct solution was replaced by a method that scales with the same scaling as the matrix-vector product: better than $\mathcal{O}(n^2)$ and approximately linearly with the number of isomers in the system being modeled. The direct factorization of the diffusion approximation matrix required for the preconditioning can be achieved with $\mathcal{O}(n)$ computational effort. Although derived independently, it turns out that using an iterative solution method within Newton's method for solving a nonlinear system has been used previously in the context of solving multicomponent solidification and fluid flow problems.[23,24]

While Paper III aimed for a robust and scalable method by building on the best robust method identified in Paper II, the current work starts with the fastest available method at the expense of robustness. To that end, the starting point will be partial spectral decomposition using a shift-invert Lanczos iteration, the fastest of the methods tested in Paper II. Once more, the most time-consuming part of the calculation is finding the solution to a linear system. In the shift-invert Lanczos case the system to be solved is similar to that for the Newton's method case of Paper III [Eq. (2)], taking the form

$$(B - \sigma I)\mathbf{z} = \mathbf{b}, \tag{3}$$

where σ is the spectral shift and again $\mathbf{b}$ is arbitrary. The obvious approach to solving Eq. (3) is by direct factorization and triangular solves, similar to the standard approach to solving Eq. (2) used in a stiff ODE integrator. This is the approach taken in the work of Paper II. As in Paper III, this linear system solve can be replaced by a preconditioned iterative method.

In this work and that presented in Paper III, we utilize the diffusion approximation to the ME to speed the solution of the full ME by using it as a preconditioner. Unlike previous applications of the diffusion approximation to the ME, the problem ultimately being solved is the full ME, not the diffusion approximation to the ME. The diffusion approximation matrix is a very good candidate for applying as a preconditioner: the approximation is quite good, yet the numerical effort required to perform operations such as inversion is much lower than for the full ME matrix. As a preconditioner the diffusion approximation is used selectively. Rather than precondition the whole problem—the solution of the ME—we focus on a specific bottleneck in a well-known solution method and use the diffusion approximation to precondition that step. Without the preconditioning the problematic step (in this case a linear system solve) cannot be completed in a fast and efficient manner.

To test the new method developed here, we model the reaction between singlet methylene and acetylene. This reaction is believed to be an important source of propargyl radicals (C_3H_3) in flames.[25,26] The formation of propargyl is believed to be a significant step in the formation of simple aromatic hydrocarbons and thus polycyclic aromatic hydrocarbons and soot.[25–32] The major route to propargyl proposed by Miller and Melius[25] involves the insertion of singlet methylene into acetylene to form C_3H_4, which isomerizes before decomposing to propargyl:

$$^1CH_2 + C_2H_2 \rightleftharpoons C_3H_4 \rightarrow C_3H_3 + H. \tag{4}$$

The rate constant for the reaction of singlet methylene with acetylene to form C_3H_4 and eventually propargyl has been measured experimentally by several different methods over the past 20 years. A good summary is given by Blitz et al.[33] While it appears that a single time-independent bimolecular rate constant is appropriate for the disappearance of $^1CH_2$ under pulsed conditions, the behavior of the remainder of the system is not yet well established. The modeling of the $^1CH_2 + C_2H_2$ system in this work is similar to that used in previous ME studies[6,7,33,34] and identical to the test problems used in Papers II and III.

The structure of the paper is as follows: In the next section we review the ME generally, including an overview of the diffusion approximation. In Sec. III we develop the new method, replacing the linear system solve in the shift-invert Lanczos algorithm. In Sec. IV we discuss the application of the new method to the $^1CH_2 + C_2H_2$ system. In Sec. V we discuss the new method specifically in terms of the amount of computer time required to solve ME problems. In Sec. VI we sum up, including a concise recap of the method.

II. REVIEW OF THE MASTER EQUATION

The ME is well-known and described in detail elsewhere,[1–5,7,8] so only some details pertinent to the current case shall be pointed out here. The energy grained multi-well ME discretized over a set of energy grains p_i (with each isomer described by a subset of the n grains p_i) can be written as a series of equations of the form

$$\frac{dp_i}{dt} = \omega\, \delta E \sum_j\, P_{ij}p_j - \omega p_i - p_i \sum_r\, k_i^{(L,r)} + \sum_r\, k_i^{(G,r)} p_{i_{ri}}, \tag{5}$$

where ω is the collision frequency, δE is the energy grain size and P_{ij} describes collisional energy transfer within each species. The $k_i^{(L,r)}$ and $k_i^{(G,r)}$ are microscopic rate constants for the interconversion reactions, with $k_i^{(L,r)}$ describing the rate of loss of population from grain i and $k_i^{(G,r)}$ describing the rate of gain of population to grain i from grain i_{ri}. Clearly, $k_i^{(L,r)}$ and $k_i^{(G,r)}$ are related by detailed balance. The sum over j is over all energy grains belonging to the same species as grain i while the sums over r are over all reactive channels. For notational simplicity the explicit time dependence of p_i has not been shown.

Bimolecular reactions are easily incorporated if they are modeled under pseudo-first-order conditions (which makes the reaction linear in p_i).[13] The first two terms on the right of Eq. (5) do not apply in the bimolecular case if the reactant not in excess is assumed to maintain its equilibrium distribution, which is a reasonable assumption. The $k_i^{(L,r)}$ and $k_i^{(G,r)}$ terms for reactions from bimolecular states are then formed by the microscopic rate constant for the reaction multiplied

by the total population of the bimolecular species assumed to be in excess and the normalized Boltzmann population of the reactant not in excess. As usual, detailed balance can be invoked to determine the rate constants for the reverse reaction.

Equation (5) can be written as a simple first-order linear matrix ODE to facilitate matrix analysis methods. Particularly when employing spectral solution methods, the efficiency and stability of the solution can be improved by symmetrizing the coefficient matrix before solving the ME (though the symmetric form of the ME is not automatically the best choice[20,35]). If $\mathbf{f}$ is the vector describing the Boltzmann population of the system and the diagonal matrix S is given by

$$S_{ii} = f_i^{-1/2}, \tag{6}$$

then transforming the original population vector $\mathbf{p}$ and ME ODE coefficient matrix A according to

$$\boldsymbol{\rho} = S\mathbf{p} \tag{7}$$

and

$$B = SAS^{-1} \tag{8}$$

yields the ODE already given as Eq. (1), with the matrix B symmetric (such that $B = B^T$). When the symmetrized ME is solved, it yields transformed "populations" $\boldsymbol{\rho}(t)$, which are transformed back to the real, observable population distributions via

$$\mathbf{p}(t) = S^{-1}\boldsymbol{\rho}(t). \tag{9}$$

In this work, the symmetrized representation of the matrix is used to improve the stability of the spectral decomposition. (Symmetric matrices are guaranteed to have purely real eigenvalues and eigenvectors.)

The first term on the right of Eq. (5) describes collisional energy transfer (CET) within each isomer and manifests itself within the ME matrix as a dense block. Invoking the diffusion approximation describes CET as derivatives of particular energy-dependent functions. Using finite differences to approximate the derivatives gives CET modeled by a purely local process,

$$\sum_j P_{ij}p_j \approx a_i p_{i-1} + b_i p_i + c_i p_{i+1}, \tag{10}$$

where the constants a_i, b_i and c_i depend on the particulars of the diffusion approximation and finite difference scheme used. Green, Robertson and Pilling[36] concluded that of the various diffusion approximation formulations available, drift-determined diffusion gave the best results for unimolecular MEs. Substituting Eq. (10) into Eq. (5) yields tridiagonal blocks replacing the dense blocks in the ME matrix. The well-defined sparse structure of the diffusion approximation matrix allows matrix-vector products with both the matrix and its inverse to be calculated very quickly. While solving the diffusion version of the ME is clearly an approximation to the solution of the full ME, this approximation has proved useful.[6,7,16,36–39]

One of the keys to the scalability of the methods presented in the current work is the inversion of the diffusion approximation matrix. As described elsewhere,[6,7] rearranging the ordering of the energy grains within the state space $\mathbf{p}$ to bring grains of the same energy together results in a banded matrix with the bandwidth equal to the number of isomers being modeled, p. A banded matrix of this type can be factorized with computational effort scaling at $\mathcal{O}(np^2)$, a vast improvement over the $\mathcal{O}(n^3)$ standard solve as usually $p \ll n$. The increased well-structured sparsity of the matrix also leads to a significant reduction in the memory needed to store the matrix, with the required storage scaling at $\mathcal{O}(np)$.

Including linearized bimolecular reactions changes the banded matrix arising from the diffusion approximation to a banded "arrowhead" matrix, with nonzero rows and columns added to the bottom and right of the banded matrix. Critically, the factorization of such an arrowhead matrix does not lead to fill-in outside of the arrowhead structure so that bimolecular channels do not alter the basic scaling of solving linear systems involving the diffusion approximation ME.[6,7]

III. DEVELOPING THE SOLUTION METHOD

Rather than develop an entirely new methodology, in this work we aim to adapt and combine existing techniques to yield a fast and scalable method. A key component of our approach is the utilization of the diffusion approximation to the ME.[36] While using the diffusion model to approximate the full ME has shown promise[6,7,16,36–39] it has not previously been used to facilitate the fast and accurate solution of the full, dense ME.

A. The diffusion approximation matrix

As one would expect, the matrix derived from the diffusion approximation to the full ME is an approximation to the matrix derived from the full ME. While one could consider this approximation from the point of view of the individual elements of the matrices, it is more useful to consider the action of the matrix. If B is the full ME matrix, D is the symmetrized diffusion approximation ME matrix and $\mathbf{v}$ is some vector, then D is an approximation to B in the sense that

$$D\mathbf{v} \approx B\mathbf{v}. \tag{11}$$

Similarly, the inverse D^{-1} is an approximation to B^{-1}. The value in using the diffusion approximation to the full ME comes from the fact that both $D\mathbf{v}$ and $D^{-1}\mathbf{v}$ can be calculated much faster than $B\mathbf{v}$ and $B^{-1}\mathbf{v}$, and with much better scaling as the size of the discretization of the ME increases.

Generally speaking, there are two subtly different approaches to using such an approximation to speed up the solution of a problem. While you could use a solution to the approximation to the original problem to build a solution to the original problem, in this work we use the widely known technique of preconditioning. The philosophy behind preconditioning is simple: use an available approximation to transform a problem that is hard to solve into a problem that is easier to solve.

B. Iterative inversion of *B* for shift-invert Lanczos

We showed in Paper II that a very fast way of solving multi-well ME problems—provided the temperature being modeled is high enough—is to generate a partial spectral decomposition of the population evolution operator using the Lanczos iteration. The full spectral expansion is given by

$$\boldsymbol{\rho}(t) = \sum_{i=1}^{n} \alpha_i \exp(\lambda_i t) \mathbf{y}_i, \tag{12}$$

where λ_i and $\mathbf{y}_i$ are the n eigenvalues and eigenvectors of B and $\alpha_i = \langle \mathbf{y}_i, \boldsymbol{\rho}(0) \rangle$. While Eq. (12) is exact for all times t, for medium- to long-time behavior only a limited number of the eigenpairs of B are required, being those with the smallest eigenvalues.

Krylov subspace based methods, such as the Lanczos method, tend to converge the extremes of the eigenvalue spectrum first, particularly for large magnitude eigenvalues. It has been found that despite the desired smallest eigenvalues being extremal, the fact that these eigenvalues are so many orders of magnitude smaller than the eigenvalues at the other extreme of the spectrum, combined with the spread of magnitudes in the eigenvectors, means that full convergence of these eigenpairs is nearly impossible to achieve.[21] To combat this, we turn to a shift-invert spectral transformation.[40,41] As an eigenvalue λ of B becomes an eigenvalue $1/(\lambda - \sigma)$ of $(B - \sigma)^{-1}$ with an unchanged eigenvector, the shift-invert transformation transforms eigenvalues of B near σ to be those of largest magnitude of $(B - \sigma I)^{-1}$, and hence those expected to converge the fastest. A shift-invert strategy is often used to find eigenpairs on the interior of the spectrum, as these normally converge slowly with Krylov subspace based methods. In this case the smallest eigenvalues are problematic. Hence we wish to transform the matrix so that the eigenpairs with the smallest eigenvalues become those with the largest eigenvalues, making a selection of $\sigma = 0$ appropriate. In difficult cases when convergence of required interior eigenpairs is problematic, σ can be set appropriately to focus on these slowly converging eigenpairs. We demonstrated in Paper II that the zero-shift shift-invert transformation makes the application of the Lanczos method to this ME possible, and in fact leads to fast convergence of the Lanczos iteration, at the cost of solving Eq. (3) at each iteration. Despite the lack of a spectral shift, we continue to refer to the zero-shift case as a shift-invert method and generally present the method retaining the $B - \sigma I$ and $D - \sigma I$ notation.

In Paper II the ARPACK[40] package was used as the implementation of the Lanczos method as this package has several useful features such as implicit restarting and eigenpair locking. The zero-shift case of Eq. (3) was solved directly via a general triangular factorization. Similarly to Paper III, here we replace this direct linear system solve with an iterative method. Generally, iterative solution methods require little more than the matrix-vector product, giving them similar scaling. In this case the matrix-vector product involving the full ME matrix scales well: approximately linearly with the number of isomers being modeled. In this work we use the

GMRES method,[42,43] as implemented in the Sparse Linear Algebra Package (SLAP) of SLATEC.[44]

In its pure form, the GMRES algorithm requires only the matrix-vector product and some relatively minor operations to build and solve the system in a small Krylov subspace. A ME solution methodology can then be implemented as code to first construct the ME matrix (using whatever storage scheme or decompositions are required to fit the matrix into the available memory and give a fast matrix-vector multiply routine) before commencing the Lanczos iteration. At each iteration of the Lanczos method, the GMRES routine is called to solve Eq. (3) using the matrix-vector product specific to the implementation of the ME matrix. This approach will be denoted the Lanczos/GMRES method. As the size of the ME being solved changes, the scaling of such a scheme is dominated by the scaling of the specific matrix-vector product routine, or approximately linearly with the number of isomers being modeled.

When an easily inverted approximation is available, it is standard practice to use that approximation to precondition the iterative solution.[43] This is particularly common in quantum scattering calculations,[45,46] where the shift-invert spectral transformation is known as applying Green's operator and many different forms of preconditioning are used, both physically based (similar to the current case) and chosen on purely mathematical or numerical grounds.

To precondition a linear system solve of the form of $A\mathbf{z} = \mathbf{b}$ one wants a preconditioning matrix M that is an approximation to A so that

$$M^{-1}A \approx I. \tag{13}$$

If the available approximation is not easy to invert then it is not useful—forming M^{-1} must be significantly faster than forming A^{-1}, as the latter is the problem we are trying to accelerate. In the current case an easily inverted approximation is readily available in the banded arrowhead matrix derived from the diffusion approximation. D is a sparse approximation to B, so that

$$(D - \sigma I)^{-1}(B - \sigma I) \approx I, \tag{14}$$

to give an appropriate preconditioning matrix as $D - \sigma I$. Specifically, the preconditioned GMRES method requires the solution of

$$(D - \sigma I)\mathbf{c} = \mathbf{d}, \tag{15}$$

for $\mathbf{c}$. In the general case, the presence of a nonzero σ does not alter the structure of the diffusion approximation matrix so that $D - \sigma I$ can still be permuted to banded arrowhead form, factorized and solved quickly and with good scalability. In this work the asymmetric LU factorization was used to factorize $D - \sigma I$ rather than the faster Cholesky factorization. While the Cholesky factorization is only appropriate for *definite* matrices (matrices whose eigenvalues are all of the same sign and nonzero) which is not guaranteed for general σ, recall that we take $\sigma = 0$ here, exposing the underlying definiteness of D. The LU factorization is much more tolerant of near singularity than the Cholesky factorization and

was found to be much more stable. The extra work required in not exploiting symmetry was not significant to the overall calculation time.

Each GMRES iteration requires one matrix-vector multiply with $B - \sigma I$ and one solution of Eq. (15), the latter effectively two band triangular substitutions. Provided the number of isomers is smaller than the number of energy grains describing each isomer, the $B - \sigma I$ matrix-vector multiply is the more computationally demanding of the two and dominates the calculation time required. Thus the computational effort required for each GMRES iteration scales in a similar manner to the $B - \sigma I$ matrix-vector multiply: approximately linearly with the number of isomers being modeled. This scaling dominates the solution of Eq. (3), and hence the Lanczos iteration and overall master equation solution.

While using a preconditioned GMRES iteration to replace the inversions required for shift-invert Lanczos and those required for stiff ODE integration (as per Paper III) are similar in principle, there is at least one significant difference. In the stiff integration case, the matrix being inverted changes for each inversion, as the γ of Eq. (2) generally changes for each successive right hand side $\mathbf{b}$. On the other hand, in the current shift-invert Lanczos case the shift remains fixed for the duration of the calculation. Hence when using direct factorization in the shift-invert Lanczos case the matrix $B - \sigma I$ need be factorized once only, whereas for the stiff integration the matrix $I - \gamma B$ must be factorized for each time step. While the factorization will dominate for large systems in the shift-invert Lanczos case, a much greater proportion of the calculation will be spent doing $\mathcal{O}(n^2)$ triangular system solves than for the stiff integration case. This means that while the scaling of the method will be improved by replacing direct factorization with preconditioned GMRES in the present case, the resulting method is less likely to be significantly faster than the direct factorization version than was the case for the stiff ODE integrator results of Paper III.

IV. RESULTS FOR THE $^1CH_2 + C_2H_2$ REACTION

A. The model

In previous work[7,8,21,22,34] we have modeled the propargyl formation reaction [Eq. (4)] using various multi-well ME methods. The reaction proceeds through a multi-well collision complex. The C_3H_4 species exists as three interconverting isomers:

$$\text{propyne} \rightleftharpoons \text{cyclopropene} \rightleftharpoons \text{allene}. \tag{16}$$

The $^1CH_2 + C_2H_2$ reaction produces the cyclopropene isomer, which must isomerize to allene or propyne before irreversibly decomposing to the propargyl product. This reaction scheme is summarized in Fig. 1.

An energy grain size of 200 cm^{-1} was used throughout, giving a matrix of order 714. The collision frequency was taken as the Lennard-Jones value. The rotational constants and vibrational frequencies were taken from Karni et al.[47] The $^1CH_2 + C_2H_2$ microscopic rate constants were derived from the data of Blitz et al.[33] The propargyl formation from
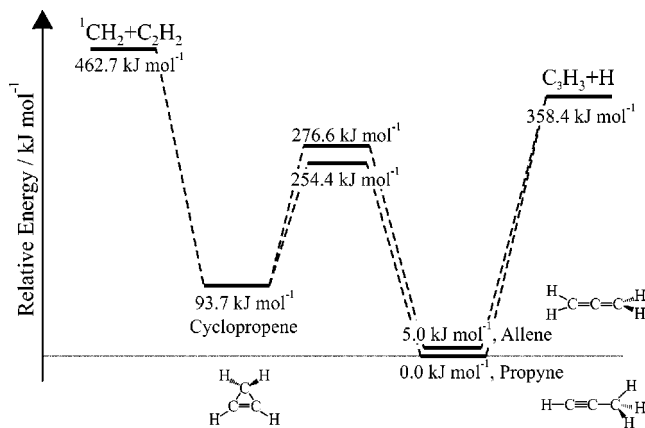

FIG. 1. Schematic reaction scheme for the modeled $^1CH_2 + C_2H_2$ reaction.

allene and propyne was modeled according to the expressions of Harding and Klippenstein.[48] In all cases, the initial population represented the dissociated state of singlet methylene with no C_3H_4 present. Under all of the conditions tested in this work (300–2000 K and 1–1000 Torr), the population profiles of the five species involved can readily be calculated by other means, as demonstrated in Paper II.

While the methylene plus acetylene channel was linearized and treated reversibly under pseudo-first-order conditions, our previous work shows that at low temperatures treating the propargyl formation reaction in a similar manner significantly alters the dynamics through the reformation of C_3H_4.[8] Explicitly including the products of irreversible reactions prevents the symmetrization of the ME matrix and excludes spectral approaches to solving the ME.[7] Thus the propargyl population was calculated by consideration of conservation of the total population.

B. Lanczos/GMRES solution in double precision

Fresh modeling based on the full eigendecomposition of the ME calculated in high precision indicates that 10 eigenpairs is sufficient to yield accurate population profiles over a reasonable range of times at all temperatures and pressures modeled in this work. Throughout this work, the ARPACK routines were called with the arguments set to calculate the ten largest magnitude eigenvalues of $(B - \sigma I)^{-1}$ with the corresponding eigenvectors, from a Krylov subspace of dimension 20. This is different to the 5 or 25 eigenpairs used in Paper II and was chosen for consistency across the pressure regime. As previously stated, $\sigma = 0$ was used throughout this work.

One would expect that the capabilities of the Lanczos/GMRES method be similar to those reported in Paper II for the zero-shift shift-invert Lanczos method with the standard solution by triangular factorization. That is, in double precision arithmetic the method should be accurate at high temperatures and pressures, but less reliable at lower temperatures and pressures unless one resorts to increasing the precision of the entire calculation.

For the double precision calculations, the GMRES solution was deemed converged when the error norm was less than 10^{-13}. Unless otherwise stated, all GMRES calculations
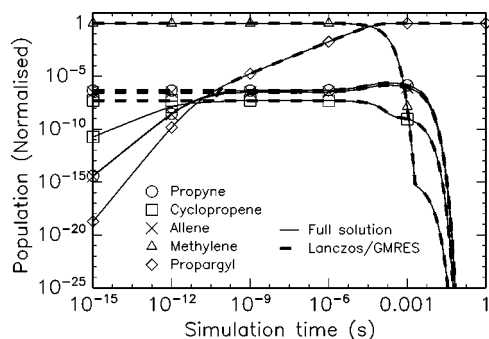
FIG. 2. Population profiles for the five species involved in the $^1CH_2$ $+C_2H_2$ reaction modeled at 1600 K and 1000 Torr. Accurate populations from the full spectral solution and populations calculated with the truncated expansion from the Lanczos/GMRES method.
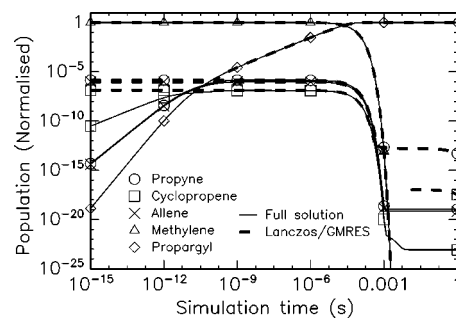

FIG. 3. As per Fig. 2 but modeling 1200 K and 1 Torr using double precision arithmetic. (Sections of the population profiles that were calculated as being negative are not shown.)

were restarted after 20 iterations. This should be considered a very frequently restarted GMRES calculation, which in general can lead to convergence difficulties.

The behavior of the GMRES solution of Eq. (3) was similar to that reported in Paper III for the solution of Eq. (2). That is, without preconditioning tens of thousands of iterations were required to converge the solution, even with very loose convergence criteria and irrespective of the restarting strategy. Preconditioning the iterative solution with the arrowhead diffusion approximation matrix lead to rapid convergence, taking between 20 and 500 iterations to converge in all cases.

Figure 2 shows the population profiles for the system at 1000 Torr and 1600 K modeled using the new Lanczos/ GMRES method in double precision. The results shown in Fig. 2 are typical of the results of modeling at 100 Torr or 1000 Torr and 1200, 1600, or 2000 K. In all of these cases the population profiles calculated with the Lanczos/GMRES method are indistinguishable from those calculated via standard shift-invert Lanczos with the inversion being performed through a dense factorization. That is, the deviation from the accurate model population profiles at short times (at times shorter than around 10^{-10} s for the 1000 Torr and 1600 K case shown in Fig. 2) is due to the truncation of Eq. (12), not any effect of using GMRES iteration to solve Eq. (3).

With some caveats, the double precision implementation of the Lanczos/GMRES method behaves similarly when modeling pressures of 10 or 1 Torr to when modeling higher pressures. When modeling 1200 K or above, the method performs reasonably well. Typical results are shown in Fig. 3 for 1200 K and 1 Torr.

The long-time behavior is not being accurately modeled. This is a common mode of failure of this type of multi-well ME including a bimolecular reactant source, where the expansion coefficients α_i in Eq. (12) are determined solely by the element of the eigenvector corresponding to the bimolecular species. When this particular element is determined inaccurately, as in this case, the calculated population profiles exhibit systematic errors, sometimes globally, sometimes in particular regimes. In this case, the long-time populations are inaccurate, giving too large and sometimes negative populations. As indicated in Fig. 3, while the long-

time populations of the C_3H_4 isomers are significantly in error, the calculated long-time populations are very low and are likely to be insignificant in any practical application.

C. Using higher precision to model lower temperatures

When the temperature being modeled was reduced to 900 K, the double precision Lanczos/GMRES calculated solution was slightly less stable than the standard full inversion shift-invert Lanczos approach. In the 1000 Torr case, while the calculated population profiles of propyne, allene, the propargyl product and the methylene reactant were accurate from times around 10^{-10} s, the calculated population of cyclopropene, significantly lower than the other isomers, was generally in error by around an order of magnitude. When the cyclopropene population was calculated to be larger than 5×10^{-7} at simulation times in the region of $10^{-6} - 10^{-3}$ s the error was smaller. At 100 Torr, the 900 K results were less reliable still, with all three C_3H_4 isomer population profiles being significantly in error for at least some significant portion of the chemically interesting timescales. At temperatures lower than 900 K the population profiles calculated in double precision were erroneous.

At both 10 and 1 Torr and temperatures of 900 K or less the calculated populations were largely spurious. At 900 K only the C_3H_4 isomer populations were affected, with the methylene reactant disappearance and the propargyl product production not significantly in error. At lower temperatures no accurate population profiles were produced.

Implementing the Lanczos/GMRES method in quadruple precision gives the same range of applicability as the standard shift-invert Lanczos indicated in Paper II. That is, the method produces accurate population profiles (within the truncated spectral expansion) at temperatures down to 600 K. For these calculations the error norm convergence tolerance was maintained at 10^{-13}. The population profiles were not significantly different to those calculated using the triangular factorization inversion reported in Paper II, meaning that even at 1 Torr the last 15 of the 25 eigenpairs included in the truncated spectral expansion of the work in Paper II had little effect.

Achieving accurate modeling of the population profiles at a temperature of 300 K required a very substantial investment in computational effort. Increasing the numerical pre-
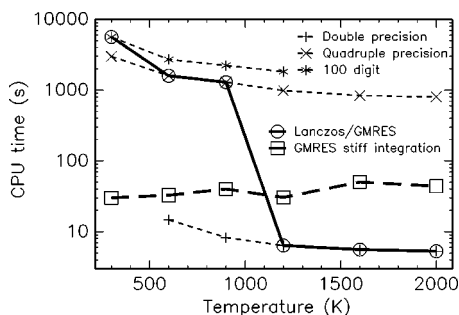
FIG. 4. CPU time required to solve the $^1CH_2 + C_2H_2$ ME as a function of the modeled temperature at 1000 Torr. The fainter dashed lines indicate the CPU times for the method at various precisions. The solid line indicates the times required for accurate modeling, taking into consideration the level of numerical precision required as the temperature changes. The time required for modeling using the stiff ODE integrator with the GMRES solve of Paper III shown for reference. Intel Pentium 4 1.9 GHz CPU.

cision well beyond quadruple precision was required. The MPFUN package of Bailey[49] was used to implement the Lanczos/GMRES method with arbitrary floating point precision. We have used MPFUN in the past to circumvent extreme ill-conditioning in spectral solutions to MEs, with 50 decimal digits of precision usually being sufficient to solve room temperature problems, including the current $^1CH_2 + C_2H_2$ multi-well system.[8,20,21,34] In the current case at 300 K the GMRES iteration would not converge to a solution with the specified error norm of 10^{-13} when 50 digit arithmetic was used, even when 200 GMRES iterations were allowed before restarting. The GMRES method did achieve convergence to the 10^{-13} error norm level when 100 digit arithmetic was used and the restart frequency was increased to 200 iterations. Shifting the calculated spectrum away from zero (through a nonzero σ parameter) did not lead to an easier calculation. Despite the relatively low GMRES convergence threshold of 10^{-13} (compared to the numerical model maintaining precision to approximately one part in 10^{100}), the population profiles calculated at 300 K were accurate.

V. SPEED OF THE NEW METHOD

When implemented in double precision, the Lanczos/GMRES method is indeed fast. Figure 4 shows sample timings of the Lanczos/GMRES method to solve the model problem when implemented on an Intel Pentium 4 1.9 GHz CPU. Shown for comparison are the times published in paper III for the stiff ODE integrator using the GMRES iteration to solve the correction equation. Note that the vertical CPU time axis of Fig. 4 is a logarithmic axis.

In double precision, the Lanczos/GMRES method is much faster than the integration method of Paper III, taking around 5 s of CPU time compared to the latter's 30–50 s. Both of these methods should have similar scaling, approximately that of the matrix-vector product or approximately linear in the number of isomers being modeled. This 5 s execution time is significantly slower than the standard shift-invert Lanczos method reported in Paper II, which required around 0.5 s to solve the ME. When one considers the scaling of solving Eq. (3) for the two approaches, the current Lanczos/GMRES should be faster than the standard shift-

invert Lanczos approach when solving problems three to nine times larger than the current three isomer, 714 energy grain case.

Once the temperature being modeled is decreased to a point where the double precision implementation of the Lanczos/GMRES method is unable to accurately model the population profiles, a large penalty is paid in terms of execution time. The CPU time required for the quadruple precision solution was over 1000 s, more than 30 times greater than that required to solve the problem with the stiff integrator using GMRES. However, it should be recalled that the results presented in Paper II show the penalty for moving to quadruple precision is not nearly as high on a 64 bit processor (such as the HP Alpha or SGI R14000) as on the 32 bit Pentium processors used in this work.

VI. SUMMARY

To aid in the implementation of the new Lanczos/GMRES method, our procedure is restated here [bearing in mind when referring back to Eqs. (3) and (15) that $\sigma = 0$].

(1) Set up multi-well ME, storing only the dense blocks on the main diagonal of the ME matrix, the off-block-diagonal interconversion terms and the "arrowhead" terms associated with the bimolecular reactions for an efficient and fast specialized matrix-vector product.
(2) Set up the diffusion approximation matrix from the full matrix, storing the matrix permuted to a banded arrowhead matrix in a number of vectors.
(3) Factorize the diffusion approximation using a nonsymmetric LU factorization, overwriting the diffusion approximation storage.
(4) Loop over calls to the ARPACK driver routine, requesting convergence of the largest magnitude eigenvalues (of the inverted matrix) with corresponding eigenvectors.
(5) When the ARPACK routine returns and requests a matrix-vector multiplication, call the SLAP GMRES routine to solve Eq. (3).
(6) The matrix-vector multiply routine passed to the GMRES routine should utilize the structure of the multi-well matrix to calculate $B\mathbf{v}$ as quickly as possible.
(7) The preconditioning routine passed to the GMRES routine to solve Eq. (15) should first permute $\mathbf{d}$ to the order giving the banded arrowhead structure of the diffusion approximation matrix before solving Eq. (15) from the factorized D and permuting the solution back to the original ordering.
(8) After convergence of the Lanczos iteration, call the ARPACK post-processing routine and propagate the initial population vector using a truncated version of Eq. (12).

The importance of an algorithm to solve the ME at or near linear scaling with the system size cannot be overstated in the context of solving large problems. Current $\mathcal{O}(n^3)$ algorithms simply cannot be used effectively when dealing with 2-D, multi-well or even 2-D multi-well MEs discretized over tens of thousands of points. The methods presented here and in Paper III are very significant steps toward solving large ME problems. The key to these new methods is using

the GMRES method to iteratively solve linear systems involving the ME matrix, which is only feasible when using the diffusion approximation to provide fast and effective preconditioning. While not as robust as the method based on direct integration presented in Paper III, the Lanczos/GMRES method presented here is an order of magnitude faster with the same scaling when the conditions being modeled are sufficiently high in temperature and pressure to allow an accurate solution to be calculated in double precision.

ACKNOWLEDGMENTS

[1] R. G. Gilbert and S. C. Smith, *Theory of Unimolecular and Recombination Reactions* (Blackwell Scientific, Oxford, 1990).

[2] I. Oref and D. C. Tardy, Chem. Rev. **90**, 1407 (1990).

[3] S. Nordholm and H. W. Schranz, in *Advances in Chemical Kinetics and Dynamics*, edited by J. R. Barker (JAI, Greenwich, 1995), Vol. 2A, p. 245.

[4] G. D. Billing and K. V. Mikkelsen, *Introduction to Molecular Dynamics and Chemical Kinetics* (Wiley, New York, 1996).

[5] K. A. Holbrook, M. J. Pilling, and S. H. Robertson, *Unimolecular Reactions*, 2nd ed. (Wiley, Chichester, 1996).

[6] K. E. Gates, S. H. Robertson, S. C. Smith, M. J. Pilling, M. S. Beasley, and K. J. Maschhoff, J. Phys. Chem. A **101**, 5765 (1997).

[7] T. J. Frankcombe, S. C. Smith, K. E. Gates, and S. H. Robertson, Phys. Chem. Chem. Phys. **2**, 793 (2000).

[8] T. J. Frankcombe and S. C. Smith, Faraday Discuss. **119**, 159 (2002).

[9] N. Snider, J. Chem. Phys. **42**, 548 (1964).

[10] M. Quack, Ber. Bunsenges. Phys. Chem. **88**, 94 (1984).

[11] J. A. Miller, S. J. Klippenstein, and S. H. Robertson, J. Phys. Chem. A **104**, 9806 (2000).

[12] W. Tsang, V. Bedanov, and M. R. Zachariah, Ber. Bunsenges. Phys. Chem. **101**, 491 (1997).

[13] M. A. Hanning-Lee, N. J. B. Green, M. J. Pilling, and S. H. Robertson, J. Phys. Chem. **97**, 860 (1993).

[14] S. C. Smith and R. G. Gilbert, Int. J. Chem. Kinet. **20**, 307 (1988).

[15] S. C. Smith and R. G. Gilbert, Int. J. Chem. Kinet. **20**, 979 (1988).

[16] S. H. Robertson, M. J. Pilling, and N. J. B. Green, Mol. Phys. **89**, 5131 (1996).

[17] S. J. Jeffrey, K. E. Gates, and S. C. Smith, J. Phys. Chem. **100**, 7090 (1996).

[18] J. A. Miller, S. J. Klippenstein, and C. Raffy, J. Phys. Chem. A **106**, 4904 (2002).

[19] P. K. Venkatesh, A. M. Dean, M. H. Cohen, and R. W. Carr, J. Chem. Phys. **107**, 8904 (1997).

[20] T. J. Frankcombe and S. C. Smith, Comput. Phys. Commun. **141**, 159 (2001).

[21] T. J. Frankcombe and S. C. Smith, J. Theor. Comput. Chem. **2**, 179 (2003).

[22] T. J. Frankcombe and S. C. Smith, J. Chem. Phys. **119**, 12729 (2003), preceding paper.

[23] D. A. Knoll, D. B. Kothe, and B. Lally, Numer. Heat Transfer, Part B **35**, 439 (1999).

[24] D. A. Knoll, W. B. Vanderheyden, V. A. Mousseau, and D. B. Kothe, SIAM J. Sci. Comput. (USA) **23**, 381 (2001).

[25] J. A. Miller and C. F. Melius, Combust. Flame **91**, 21 (1992).

[26] J. D. Adamson, C. L. Morter, J. D. DeSain, G. P. Glass, and R. F. Curl, J. Phys. Chem. **100**, 2125 (1996).

[27] H. Wang and M. Frenklach, Combust. Flame **110**, 173 (1997).

[28] H. Böhm, K. Kohse-Höinghaus, F. Lacas, C. Rolon, N. Darabiha, and S. Candel, Combust. Flame **124**, 127 (2001).

[29] C. H. Wu and R. D. Kern, J. Phys. Chem. **91**, 6291 (1987).

[30] R. D. Kern, H. J. Singh, and C. H. Wu, Int. J. Chem. Kinet. **20**, 731 (1988).

[31] U. Alkemade and K. H. Homann, Z. Phys. Chem., Neue Folge **161**, 19 (1989).

[32] R. D. Kern, H. Chen, J. H. Kiefer, and P. S. Mudipalli, Combust. Flame **100**, 177 (1995).

[33] M. Blitz, M. S. Beasley, M. J. Pilling, and S. H. Robertson, Phys. Chem. Chem. Phys. **2**, 805 (2000).

[34] T. J. Frankcombe, Ph.D. thesis, University of Queensland, 2002.

[35] T. J. Frankcombe and S. C. Smith, J. Comput. Chem. **21**, 592 (2000).

[36] N. J. B. Green, S. H. Robertson, and M. J. Pilling, J. Chem. Phys. **100**, 5259 (1994).

[37] S. H. Robertson, M. J. Pilling, D. A. Baulch, and N. J. B. Green, J. Phys. Chem. **99**, 13452 (1995).

[38] S. H. Robertson, A. I. Shushin, and D. M. Wardlaw, J. Chem. Phys. **98**, 8673 (1993).

[39] S. H. Robertson, M. J. Pilling, K. E. Gates, and S. C. Smith, J. Comput. Chem. **18**, 1004 (1997).

[40] R. B. Lehoucq, D. C. Sorensen, and C. Yang, *ARPACK Users' Guide: Solution of Large Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods* (SIAM, Philadelphia, 1998).

[41] Y. Saad, *Numerical Methods for Large Eigenvalue Problems* (Manchester University Press, Manchester, 1992).

[42] Y. Saad and M. H. Schultz, SIAM J. Sci. Comput. (USA) **7**, 856 (1986).

[43] Y. Saad, *Iterative Methods for Sparse Linear Systems* (PWS Publishing Co., Boston, 1996).

[44] M. K. Seager, Tech. Rep. UCRL-100267, Lawrence Livermore National Laboratory, Livermore, 1988.

[45] U. Peskin, W. H. Miller, and Å. Edlund, J. Chem. Phys. **103**, 10030 (1995).

[46] B. Poirier and W. H. Miller, Chem. Phys. Lett. **265**, 77 (1997).

[47] M. Karni, I. Oref, S. Barzilai-Gilboa, and A. Lifshitz, J. Phys. Chem. **92**, 6924 (1988).

[48] L. B. Harding and S. J. Klippenstein, Proc. Combust. Inst. **28**, 1503 (2000).

[49] D. H. Bailey, ACM Trans. Math. Software **21**, 379 (1995).