

Cambridge Centre for Computational Chemical Engineering

University of Cambridge

Department of Chemical Engineering

Preprint

ISSN 1473 – 4273

A predictor-corrector algorithm for the coupling of stiff ODEs to a particle population balance

Matthew S Celnik¹ Robert I A Patterson¹

Markus Kraft^{1*}, Wolfgang Wagner²

released: 29 May 2008

¹ Department of Chemical Engineering
University of Cambridge
New Museums Site, Pembroke Street,
Cambridge, CB2 3RA, UK
*E-mail: mk306@cam.ac.uk

² Weierstrass Institute for
Applied Analysis and Stochastics
Mohrenstraße 39
D-10117 Berlin, Germany

Preprint No. 58



c4e

Key words and phrases: predictor-corrector, Monte-Carlo, population balance, splitting, soot

Edited by

Cambridge Centre for Computational Chemical Engineering
Department of Chemical Engineering
University of Cambridge
Cambridge CB2 3RA
United Kingdom.

Fax: + 44 (0)1223 334796

E-Mail: c4e@cheng.cam.ac.uk

World Wide Web: <http://www.cheng.cam.ac.uk/c4e/>

Abstract

In this paper a novel predictor-corrector algorithm is presented for the solution of coupled gas-phase—particulate systems. The emphasis of this work is the study of soot formation, but the concepts can be applied to other systems. This algorithm couples a stiff ODE solver to a Monte Carlo population balance solver. Such coupling has been achieved previously for similar systems using a Strang operator splitting algorithm, however, that algorithm demonstrated several numerical issues which resulted in a high computational cost to acquire adequate precision. In particular a source-sink instability was identified whereby a large-magnitude source term present in the ODE system was competing with a similarly sized sink term in the population balance. This instability required that the splitting-step size was very small in order to keep numerical error sufficiently low. A predictor-corrector algorithm has been formulated to negate this instability. An additional efficiency is gained with this algorithm as a principal computational cost of the Strang splitting algorithm is removed: the requirement to re-initialise the ODE solver every splitting-step. The numerical convergence of the new algorithm is demonstrated, and its efficiency is compared to that of the Strang splitting algorithm. Substantial computation time savings are demonstrated, which allow a fixed error in three studied system functionals to be achieved with an order-of-magnitude reduction in computation time.

Contents

1	Introduction	3
1.1	Studied System	4
1.2	Formulation of Equations	4
2	Particle–Chemistry Coupling Strategies	6
2.1	Predictor-Corrector	6
2.1.1	Predictor	7
2.1.2	Corrector	7
2.1.3	Order of the Algorithm	8
2.1.4	Under-relaxation	9
2.2	Methods for the Split Operators	10
3	Numerical Convergence	10
3.1	Error Measures	10
3.2	Test System	11
3.3	Algorithm Parameters	12
3.4	Convergence Results	12
4	Algorithm Efficiency	14
5	Conclusions	18

1 Introduction

Many physical systems involve the formation and growth of a particulate phase within a gas phase, for example soot, titania and silica nanoparticle formation processes. The focus of this work is soot formation, but the principles can be applied to other particulate systems. The formation of soot particles and other carbon products in combustion systems is a highly complex process [1, 5, 14, 15]. In recent years models have become available for the conversion of mass from the gas to the particulate phase [1, 16, 17, 24], which calls for detailed modelling of the resulting soot particle populations.

In order to adequately characterise such systems, numerical models for the gas phase and the particle population must be developed and combined. Generally in these systems the particulate and gas phases are strongly coupled, as gas-phase precursors form and react with particles. Gas-phase kinetics are generally described by a finite set of ordinary differential equations (ODEs), which can be solved using well-known numerical techniques such as Runge-Kutta. Substantial effort has gone into the development of numerical population balance methods. The main numerical methods used are the method of moments with interpolative closure (MoMIC) [13], fixed [21] and moving [26] sectional methods, and Galerkin methods [2]. All of these methods explicitly reduce the evolution of the particle population to a finite set of ODEs, which can be solved simultaneously with the gas-phase kinetics equations. The disadvantage of these methods is that they are restricted to relatively simple particle descriptions.

Stochastic particle methods, which originate from the Bird algorithm for molecular gas dynamics [6], have also been applied to these systems [4], and in particular have been used for the mathematical study of coagulation problems [11]. These stochastic methods allow for arbitrarily complex models for individual particles [8, 10] to be included with little additional impact on computation times [18] when compared to simple models, but even for simple particle models the computational requirements are much greater than for the other population methods mentioned above.

Coupling population balance calculations using the ODE methods mentioned above to chemical reaction calculations is generally possible, but computational cost can be considerable [23]. Direct coupling between stochastic particle population balance methods and the stiff differential equation solvers used to treat chemically reacting systems is a particularly challenging problem, because of the different natures of the two solution methods. Initial work in this area for soot formation avoided a direct coupling by using MoMIC, the simplest of the above methods, to provide approximations to the effects of the soot population during the solution of the gas-phase chemistry [4] and then recalculating the soot population with a stochastic particle method in a post-processing step. However, accuracy requires a close coupling, which can be achieved using an operator splitting method [9]. In that paper [9] the coupling of a stochastic soot population balance method to a standard deterministic stiff chemistry solver for a simple batch reactor was reported. The case of a perfectly-stirred reactor with outflow was also considered, but the results demonstrated that the numerical method was stretched to its limits and possibly beyond.

The **purpose of this paper** is to discuss the limitations of the operator splitting method [9] and to propose a more robust and less computationally demanding alternative. This

introduction is completed by a presentation of the equations describing the coupled systems already mentioned. The numerical convergence of the new method is demonstrated, and the algorithm efficiency is compared to that of the previously used Strang splitting algorithm.

1.1 Studied System

A constant pressure batch reactor model was used to investigate the numerical algorithms. This model shall be briefly described. The gas-phase chemistry requires the solution of each species concentration and the temperature, that is $K + 1$ variables, where K is the number of gas-phase species. The soot particle population evolves according to the discrete Smoluchowski equation.

The governing equations (gas-phase material balance and soot population balance) for a constant pressure batch reactor are:

$$\frac{dc_k}{dt} = \dot{\omega}_k(c, T) + \dot{g}_k(c, n, T) - \gamma c_k \quad (1)$$

$$\frac{d}{dt}n(x) = I(c, x) + K(n, x) + \sum_{l=1}^L S_l(c, x) - \gamma n(x) \quad (2)$$

Where c_k is the molar concentration of species k , $\dot{\omega}_k$ is the molar production rate of species k due to reaction, \dot{g}_k is the molar production rate of species k due to soot processes, and T is the temperature. In the soot balance $n(x)$ is the number density of particles of type x , where the type is a vector of particle volume, composition, and surface area (surface-volume model only). $I(c, x)$, $K(n, x)$ and $S_l(c, x)$ are the inception, coagulation and surface reaction l rates respectively for particles of type x . c denotes the vector of species concentrations. γ is the rate of gas-phase expansion and is given by

$$\gamma = \frac{\dot{\omega}_k}{\rho} + \frac{1}{T} \frac{dT}{dt} \quad (3)$$

where ρ is the molar density of the bulk fluid. Note, the molar production rate due to soot is considered small in magnitude compared to that due to chemical reaction. For this theoretical study, only constant temperature systems were considered, therefore $\frac{dT}{dt}$ is zero.

1.2 Formulation of Equations

At a high level a system comprising soot particles in a reacting gas mixture may be described by a vector $(f_1, f_2)^T$, where f_1 contains gas-phase species concentrations and other gaseous state variables such as temperature and pressure, and f_2 represents the soot

particle population. The differential equation for which an efficient solution method is sought may then be written as:

$$\frac{d}{dt} \begin{pmatrix} f_1 \\ f_2 \end{pmatrix} = \begin{pmatrix} \mathcal{A}_1(f_1) \\ \mathcal{A}_2(f_1, f_2) \end{pmatrix} + \begin{pmatrix} \mathcal{B}_1(f_1, f_2) \\ \mathcal{B}_2(f_1, f_2) \end{pmatrix}. \quad (4)$$

The operator subscript 1 denotes processes which occur to the gas-phase, and the subscript 2 denotes processes which affect soot particles. The physical interpretations of the operators in the context of the above model are:

1. \mathcal{A}_1 is the gas-phase chemical mechanism. This has no soot particle dependence.
2. \mathcal{A}_2 is the change to the particle ensemble due to gas-phase processes. For a constant pressure system this describes the expansion of gas and hence a decrease in particle number density.
3. \mathcal{B}_1 is the change to the gas-phase chemistry due to particle inception and surface growth processes.
4. \mathcal{B}_2 is the change to the particle ensemble due to all particle processes.

The relationships of these operators to the physical model for the batch reactor are illustrated by the equations

$$\frac{dc_k}{dt} = \underbrace{\dot{\omega}_k(c, T) - \gamma(c, T)c_k}_{\mathcal{A}_1} + \underbrace{\dot{g}_k(c, n, T)}_{\mathcal{B}_1} \quad (5)$$

$$\frac{d}{dt}n(x) = \underbrace{-\gamma(c, T)n(x)}_{\mathcal{A}_2} + \underbrace{I(c, x) + \sum_{l=1}^L S_l(c, x) + K(n, x)}_{\mathcal{B}_2} \quad (6)$$

The choice of internal structure for f_1 and f_2 is obviously key. Since f_1 describes a chemically reacting gas-phase system, it is normal to use a vector of species concentrations (or mole fractions) so that $f_1 \in \mathbb{R}^d$ for some $d \in \mathbb{N}$. Exactly which species to include in f_1 is an active area of research, but the present work follows the chemistry developed by Wang and Frenklach [25], and updated by Appel et al. [1]. As discussed in the introduction, there are various models which describe soot particles, each with their own limitations. For the purposes of this analysis a general representation of the soot particle population is used whereby particles are described by a type x , from a type space E . For simpler models, such as the spherical particle [4, 20] and surface-volume [18] models, x is a fixed length array of real numbers. The entire particle population is then represented as a measure on E , following the approach of Eibeck and Wagner [12] and Patterson et al. [19].

2 Particle–Chemistry Coupling Strategies

As discussed above, some method is needed to allow stochastic and deterministic methods to be employed in order to solve for the soot particle population and the chemical species concentrations respectively in zero-dimensional systems (for example, shock tubes and plug-flow reactors).

A direct splitting of \mathcal{A} and \mathcal{B} was successfully used by Celnik et al. [9] to solve (4), but proved computationally expensive (around four minutes per run with 65,536 particles and a time-step size of 100 μs). Execution profiling showed that the high cost arose because a large number of small splitting steps were required to maintain adequate coupling between the \mathcal{A} and \mathcal{B} operators. At the start of each small step it was necessary to restart the stiff ODE solver used to calculate the effects of \mathcal{A} , because the solution had been changed by applying \mathcal{B} , and so a costly step was repeated many times.

The need for very small splitting time steps arose because there was a numerical instability inherent in the splitting, which involved the competition of source and sink terms from the two operators. An example of this problem arose with the gas-phase species pyrene, which is produced from smaller molecules by gas-phase reactions, but is heavily consumed by the formation and growth of soot particles. The magnitudes of the pyrene source term (in \mathcal{A}) and the sink term (in \mathcal{B}) were similar and significant in relation to the size of the pyrene concentration divided by the time step length. For long time steps this resulted in unphysical oscillations in the value of the pyrene concentration, which became very high when operator \mathcal{A} was applied and very low when operator \mathcal{B} was applied. Consequently rates of change of other solution components, which depended on the pyrene concentration, could not be reliably calculated.

The system of chemical reactions represented by \mathcal{A} is inevitably stiff and therefore has to be treated with implicit methods, which involve the inversion of Jacobean matrices. Such inversions are computationally very costly and so efficient solvers rely on reuse of intermediate calculation results [3]. However, the application of the operator \mathcal{B} in a splitting step leads to a discontinuous change in the solution, which makes the re-use of intermediate results impossible. Therefore, after every application of \mathcal{B} for the soot population part of the splitting, it was necessary to completely regenerate the internal data of the stiff ODE solver used for \mathcal{A} resulting in high computational costs.

2.1 Predictor-Corrector

A predictor-corrector strategy was formulated to include approximations for \mathcal{B}_1 in the solution to the \mathcal{A} part of the splitting in order to eliminate the source-sink instability discussed previously. In setting out the algorithm below no mention is made of any internal sub-stepping that may be carried out by the solvers used for the operators \mathcal{A} and \mathcal{B} .

2.1.1 Predictor

Suppose that approximate solutions to (4) have been calculated at times $t_1 < t_2 < t_3, \dots, t_i$ to give values

$$\begin{pmatrix} f_{1,1} \\ f_{2,1} \end{pmatrix}, \begin{pmatrix} f_{1,2} \\ f_{2,2} \end{pmatrix}, \dots, \begin{pmatrix} f_{1,i} \\ f_{2,i} \end{pmatrix}$$

respectively. The method for advancing the solution to a time $t_{i+1} > t_i$, begins with a split-predictor step:

1. Fit a function $B^0(t)$ (for example, an interpolating polynomial) defined on $[t_{i-m}, t_{i+1}]$, for $m \in \mathbb{N}$ previous points, to

$$\mathcal{B}_1(f_{1,i-m}, f_{2,i-m}), \mathcal{B}_1(f_{1,i-m+1}, f_{2,i-m+1}), \dots, \mathcal{B}_1(f_{1,i}, f_{2,i}).$$

Note that time t_{i+1} is outside the interpolation range.

2. Using this approximation to \mathcal{B}_1 , perform the first part of the splitting by solving

$$\frac{d}{dt} \begin{pmatrix} \tilde{f}_1^0 \\ \tilde{f}_2^0 \end{pmatrix} = \begin{pmatrix} \mathcal{A}_1(\tilde{f}_1^0) \\ \mathcal{A}_2(\tilde{f}_1^0, \tilde{f}_2^0) \end{pmatrix} + \begin{pmatrix} B^0(t) \\ 0 \end{pmatrix} \quad t \in [t_i, t_{i+1}] \quad (7)$$

with initial condition

$$\begin{pmatrix} \tilde{f}_1^0(t_i) \\ \tilde{f}_2^0(t_i) \end{pmatrix} = \begin{pmatrix} f_{1,i} \\ f_{2,i} \end{pmatrix}.$$

3. Fit a function $F^0(t)$ (for example, an interpolating polynomial) defined on $[t_{i-m+1}, t_{i+1}]$ to

$$f_{1,i-m+1}, f_{1,i-m+2}, \dots, f_{1,i}, \tilde{f}_1^0(t_{i+1})$$

4. Using this approximation to f_1 , perform the second part of the splitting by solving

$$\frac{d}{dt} \hat{f}_2^0 = \mathcal{B}_2(F^0(t), \hat{f}_2^0) \quad t \in [t_i, t_{i+1}] \quad (8)$$

with initial condition

$$\hat{f}_2^0(t_i) = \tilde{f}_2^0(t_{i+1})$$

to get a predicted solution $\left(\tilde{f}_1^0(t_{i+1}), \hat{f}_2^0(t_{i+1}) \right)^\top$ for t_{i+1} .

2.1.2 Corrector

Subsequent split-corrector iterations are performed as follows, for $J \in \mathbb{N}$ iterations:

1. $j \leftarrow 1$

2. Fit a function $B^j(t)$ (for example, an interpolating polynomial) defined on $[t_{i-m+1}, t_{i+1}]$, for $m + 1$ points (note inclusion of approximation to the point $i + 1$), to

$$\mathcal{B}_1(f_{1,i-m+1}, f_{2,i-m+1}), \mathcal{B}_1(f_{1,i-m+2}, f_{2,i-m+2}), \dots \\ \dots, \mathcal{B}_1(f_{1,i}, f_{2,i}), \mathcal{B}_1(\tilde{f}_1^{j-1}(t_{i+1}), \hat{f}_2^{j-1}(t_{i+1}))$$

3. Solve

$$\frac{d}{dt} \begin{pmatrix} \tilde{f}_1^j \\ \tilde{f}_2^j \end{pmatrix} = \begin{pmatrix} \mathcal{A}_1(\tilde{f}_1^j) \\ \mathcal{A}_2(\tilde{f}_1^j, \tilde{f}_2^j) \end{pmatrix} + \begin{pmatrix} B^j(t) \\ 0 \end{pmatrix} \quad t \in [t_i, t_{i+1}] \quad (9)$$

with initial condition

$$\begin{pmatrix} \tilde{f}_1^j(t_i) \\ \tilde{f}_2^j(t_i) \end{pmatrix} = \begin{pmatrix} f_{1,i} \\ f_{2,i} \end{pmatrix}$$

4. Fit a function $F^j(t)$ (for example, an interpolating polynomial) defined on $[t_{i-m+1}, t_{i+1}]$ to

$$f_{1,i-m+1}, f_{1,i-m+2}, \dots, f_{1,i}, \tilde{f}_1^j(t_{i+1})$$

5. Solve

$$\frac{d}{dt} \hat{f}_2^j = \mathcal{B}_2(F^j(t), \hat{f}_2^j) \quad t \in [t_i, t_{i+1}] \quad (10)$$

with initial condition

$$\hat{f}_2^j(t_i) = \tilde{f}_2^j(t_{i+1})$$

6. $j \leftarrow j + 1$

7. If $j < J$ go to 2, else stop.

8. Set solution at t_{i+1} to be last calculated values, that is,

$$\begin{pmatrix} f_{1,i+1} \\ f_{2,i+1} \end{pmatrix} \leftarrow \begin{pmatrix} \tilde{f}_1^j(t_{i+1}) \\ \hat{f}_2^j(t_{i+1}) \end{pmatrix}.$$

2.1.3 Order of the Algorithm

The first order splitting of \mathcal{A}_2 and \mathcal{B}_2 limits the method set out above to first order convergence with respect to the time step length, $\Delta t_i := t_{i+1} - t_i$. This limitation might be addressed by completely removing \tilde{f}_2^j from step 3 of the corrector iterations and introducing a polynomial for its effects into step 5, thus making the treatment of \tilde{f}_1^j and \tilde{f}_2^j symmetric, but this is not necessary in order to obtain adequate results.

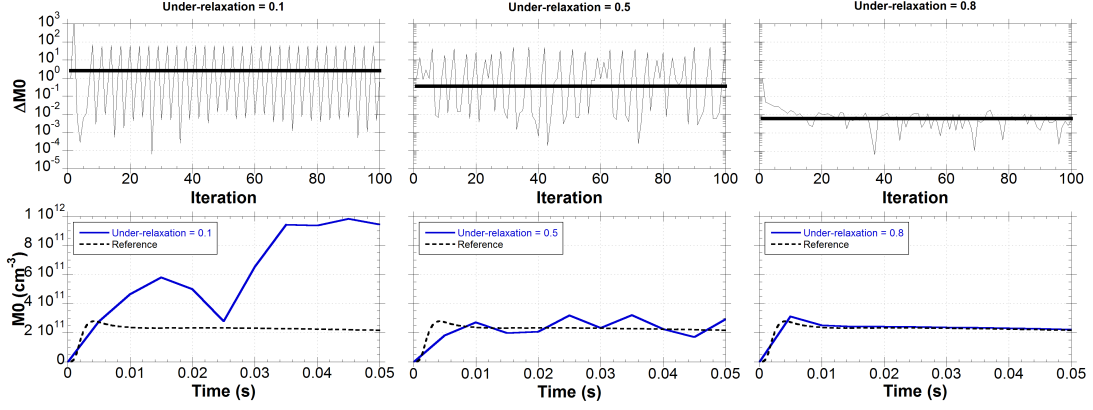


Figure 1: *Convergence with under-relaxation parameter. Simulations were performed with $N_{max} = 65536$ and $\Delta t = 2500 \mu s$. The top plots shows the deviation from the previous iteration for 100 iterations over the final time-step. Bold horizontal lines show the statistical uncertainty calculated for that set of parameters. The bottom plots show the transient behaviour of $M0$ (zeroth moment) compared to the reference solution.*

2.1.4 Under-relaxation

It was found that the above algorithm did not converge reliably for large step sizes. This problem was typically exhibited by oscillations between iterations. A standard under-relaxation technique was used to enable smoothing between iterations by constructing $B^{j+1}(t)$ as explained above, but then replacing it by

$$\tilde{B}^{j+1}(t) = rB^j(t) + (1 - r)B^{j+1}(t) \quad (11)$$

in equation (10). Suitable choice of the under-relaxation parameter r allows this unstable oscillation to be negated.

Figure 1 shows results from the test case (described below) for particle number density $M0$ (zeroth moment) at different values of r . The top row of plots shows the deviation from the previous iteration over 100 iterations for the final time-step. Two points are illustrated by these plots. Firstly the nature of the unstable oscillation is clearly seen in the leftmost plot for $r = 0.1$. Secondly it is clear that the iterations very quickly converge, after 3-5 iterations there is no discernable change in the iteration deviations. The bold line gives the statistical uncertainty calculated for that simulation. The oscillation dominates over the statistical uncertainty, therefore meaningful results cannot be obtained. This finding is mirrored in the bottom row of plots, which show the transient behaviour of $M0$ compared to the reference solution. At $r = 0.1$ the profile is drastically different from the reference. Considering the middle plots at $r = 0.5$, the iteration deviations still swamp the statistical uncertainty, though the profile (bottom plot) is a better match for the reference. At $r = 0.8$ (rightmost plots) the iteration deviations are now of a similar magnitude to the statistical uncertainty and the profile is in good agreement with the reference. It was

found that a value of $r = 0.9$ was suitable to allow all cases tested to converge in this manner.

2.2 Methods for the Split Operators

Equations (7) and (9), which can be thought of as the chemical reaction part of the problem, were solved using DDASSL [7]. DDASSL is a differential/algebraic equation solver which employs backward differentiation formulae (BDF) to solve implicit systems. It is known as a very stable solver, but it suffers from a slow initialisation procedure.

Equations (8) and (10), which contain the soot particle part of the problem, were solved using a stochastic particle method [4, 12, 19].

3 Numerical Convergence

In the following sections the effects of the parameters of the stochastic particle algorithm and the coupling methods are reported. The main stochastic parameters were the number of computational particles, N_{max} , and the number of realisations of the stochastic process, L . The number of computational particles, which is effectively the resolution in particle type space, E , was controlled with a particle doubling method [19, 22] so that the actual particle count lay in the range $[\frac{1}{2}N_{max}, N_{max}]$, except during initial transients. One realisation of the stochastic process was used during each corrector iteration, however, for reporting results below, values are averaged over all the iterations performed for that time step, which implies $J = L$ (see next paragraph). The possibility of giving later iterations, which should be more accurate, a higher weighting in this average remains to be explored.

The predictor-corrector algorithm introduces two additional parameters: the iteration count J and under-relaxation constant r , as described previously. Both the predictor-corrector and the Strang splitting coupling methods share a dependence on the splitting time-step size, Δt , over which coupling takes place. Multiple runs were performed to reduce the statistical uncertainty by averaging the output variables, and to give an estimation of the confidence interval. For this algorithm an additional speed increase was identified, whereby the iterations over a time step could be used as the independent runs ($L = J$), if a copy of the solution was stored after each iteration. The solutions after each iteration were given equal weighting in this study, though it is envisaged that later iterations could be given a higher weighting, for example by ignoring the first few iterations. As demonstrated by figure 1 after the first 3-5 iterations they are essentially converged, therefore it is not expected for iteration weights to be important after the first few iterations.

3.1 Error Measures

The error analysis used for this study is identical to that used for the previous study [9], and is briefly reviewed here. The systematic, e_{sys} , and statistical, e_{stat} , errors of a set of sample solutions, $y(t)$, to a problem that has true solution F , are defined as

$$\begin{aligned}
e_{sys}(t) &= |\zeta(t) - F(t)| \\
e_{stat}(t) &= |\eta_y(t) - \zeta(t)|
\end{aligned}
\tag{12}$$

where $\zeta(t)$ is the expectation of the distribution from which $y_l(t)$ is sampled. The sample mean is denoted $\eta_y(t)$. It is computationally impractical to find the true solution F to the system described above, therefore a numerical reference solution for which the systematic error was assumed to be very small, \tilde{F} , was used as an approximation to F in this work. For a large number of realisations, L , the central limit theorem may be used to find confidence intervals for $\eta_y(t)$. The half widths of these are estimated by

$$c_p(t) = a_p \sqrt{\frac{\mu_y(t)}{L}} \tag{13}$$

where $\mu_y(t)$ is the approximation of the variance. Representative total relative errors, c_{tot} , and relative statistical uncertainties, c_{stat} , over M time intervals are calculated as follows, where t_m is the time at the end of interval m , $m \in \{0, 1, \dots, M - 1, M\}$.

$$c_{tot} = \frac{1}{M} \sum_{m=1}^M \frac{|\tilde{F}(t_m) - \eta_y(t)|}{\tilde{F}(t_m)} \tag{14}$$

$$c_{stat} = \frac{1}{M} \sum_{m=1}^M \frac{c_p(t_m)}{\tilde{F}(t_m)} \tag{15}$$

For this study total relative errors at 10 fixed time points were averaged. This is the maximum possible number of time points given the poor resolution of the simulation with the longest time-step size (see next sections).

3.2 Test System

The simple test case reactor using in the previous study [9] was used again for this study to allow direct comparisons of the results. The test case was a batch reactor with an initial $C_2H_2/O_2/N_2/Ar$ mixture of equivalence ratio $\Phi = 2.5$. The temperature was held constant at 1650 K and the pressure was 1 atm. No particles were present in the reactor initially and the run time was 50 ms. The soot model was taken from Appel et al. [1], which is based on the chemical mechanism development work of Wang and Frenklach [25].

While this test system is very simple it never-the-less retains some physical basis as the conditions are similar to experimental set-ups. Therefore it is reasonable to suppose that the convergence and efficiency results obtained for this test system would be similar to those for a real combustion system.

In order to adequately discuss the numerical convergence, three functionals were chosen: particle number density, $M0$, equivalent to the zeroth moment of the particle population;

soot volume fraction, Fv , equivalent to the first moment of the particle population; and the gas-phase pyrene concentration ($A4$). These functionals are identical to those used in the previous study and allow the effect of the algorithm parameters on both gas-phase and particulate properties to be discussed.

3.3 Algorithm Parameters

The numerical convergence of the predictor-corrector algorithm was tested by varying the maximum particle count $N_{max} \in \{2^n, n = 6, 7, \dots, 17, 18\}$ and by varying the time-step size Δt in the range $5 - 5000 \mu s$. A similar previous study [9] demonstrated numerical convergence of the Strang splitting algorithm, so this has not been repeated here.

The parameters of the DDASSL solver used for the stiff chemistry were those found to produced good results in earlier chemistry-only calculations, in particular the relative and absolute error tolerances were 1×10^{-3} and 1×10^{-6} respectively. Simulations were performed with more stringent error tolerances over several orders of magnitude and no change in the convergence behaviour was observed, therefore these tolerances were deemed sufficient small to not affect the convergence results. All predictor-corrector simulations were performed using a relaxation constant $r = 0.9$, which was found to give stable results for all parameters. $L = J = 100$ runs/iterations were performed for all simulations.

The reference solution \tilde{F} was generated by using parameters $N_{max} = 524288$, $\Delta t = 5 \mu s$ and $L = J = 50$. For comparison this simulation took approximately 12.5 days ($> 10^7$ s) to complete on a typical computer of today.

3.4 Convergence Results

Figure 2(left) shows the convergence of total relative error of the three functionals $M0$, Fv and $A4$ with respect to time-step size for $N_{max} = 65536$. The trend lines in that figure show ideal first-order convergence. All functionals demonstrate almost first-order convergence until about $\Delta t = 100 \mu s$, at which point no further reduction in error appears possible. The predictor-corrector algorithm is a first-order method with respect to step size, therefore it is unsurprising to observe first-order convergence. Initially it was thought that the error tolerances in the ODE solver were limiting convergence, however, the same results were obtained when the relative and absolute tolerances were changed to 1×10^{-7} and 1×10^{-10} respectively. This suggests that the problem is unrelated to ODE solver tolerance. The study was repeated with a larger $N_{max} = 262144$, which is plotted in figure 2(right). This figure shows similar behaviour to the 65536 study, but limiting error is smaller. This suggests that there are two errors in competition: a splitting error associated with the time-step size, and a stochastic approximation error controlled by the maximum particle count. The error appears to reach a limit once the time-step is made sufficiently short in figure 2 because it is then dominated by the stochastic approximation error.

Figure 3 shows the convergence of total relative error with respect to maximum particle

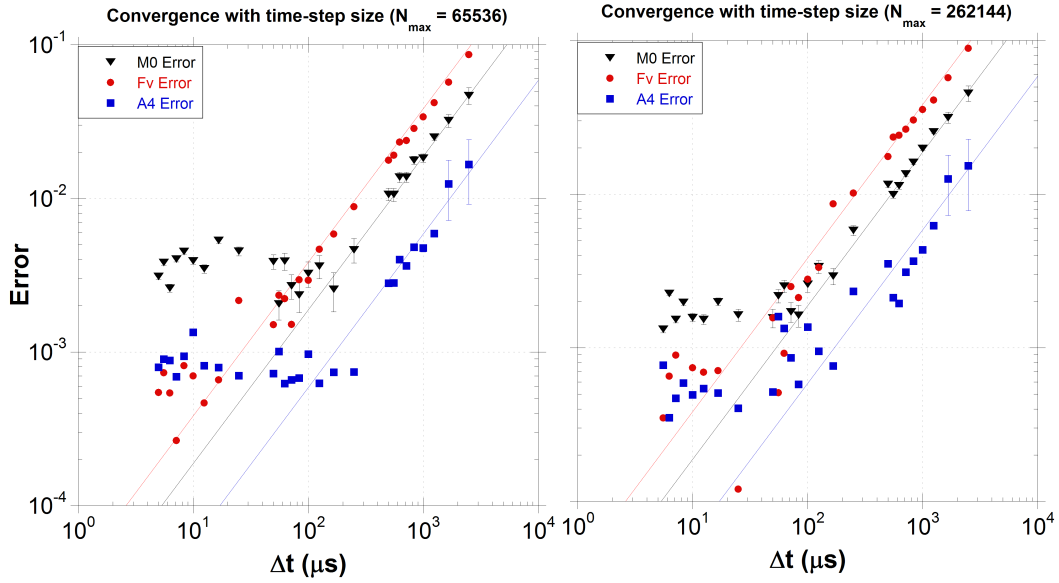


Figure 2: *Relative error convergence with step size of the predictor-corrector algorithm. The error bars show statistical uncertainty. Held constant were $N_{max} = 65536$ (left), $N_{max} = 262144$ (right) and $L = 100$ (both). Trend lines show ideal first order convergence.*

count for a fixed time-step size of $10 \mu\text{s}$. The trend lines show order 0.5 convergence. All three functionals exhibit approximately order 0.5 convergence with respect to particle count. The very small statistical uncertainties, exhibited by small error bars, suggest that the apparent scatter of the points is due to variations in splitting error rather than the stochastic approximation. A time-step size of $10 \mu\text{s}$ was chosen because the results in figure 2 suggest that this is sufficient short for total error to be controlled by N_{max} . This is backed up by figure 3 which demonstrates no limit of convergence in the studied range.

In the previous study [9] for Strang splitting it was found that there was approximately no dependence of total error on maximum particle count, however, this new study suggests that this was probably due to the relatively large splitting time-step of approximately $167 \mu\text{s}$ used. With that time-step size the splitting error would be dominant at all investigated particle counts. That time-step size was used because it was difficult to reduce the statistical uncertainty sufficiently to get meaningful results using the Strang splitting algorithm, which problem was not encountered using the predictor-corrector algorithm.

Figure 4 shows the transient behaviour of particle number density in the test system for different time-step sizes. This figure illustrates that the predictor-corrector algorithm manages to approximate the reference solution far better than the Strang splitting algorithm for larger time-steps. At $\Delta t = 5000 \mu\text{s}$, for which only 10 steps are performed, the predictor-corrector algorithm manages to oscillate around the reference solution, whereas the Strang splitting algorithm hopelessly under-predicts. At $\Delta t = 500 \mu\text{s}$ the predictor-corrector algorithm already produces an excellent fit of the reference case, while the Strang splitting algorithm has not yet reached the reference. Similar behaviour was observed for the other

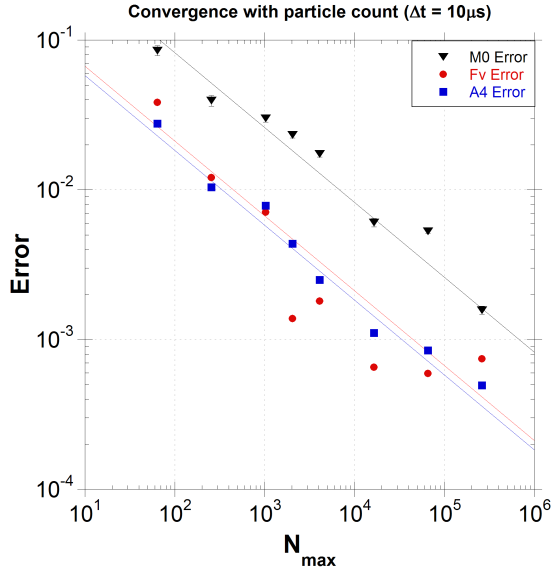


Figure 3: *Relative error convergence with maximum particle count of the predictor-corrector algorithm. The error bars show statistical uncertainty. Held constant were $\Delta t = 10 \mu s$ and $L = 100$. Trend lines show ideal order 0.5 convergence.*

two functionals studied; soot volume fraction and pyrene concentration.

4 Algorithm Efficiency

The discussion of computation times in this section is conducted with an understanding that such times are implementation dependent, in particular the source code compiler and computer architecture used are important. It is sufficient for this discussion to state that both algorithms were implemented in the same FORTRAN90 code using the same compiler, and that all simulations were performed on computers of the same specification. The computers were comparable to standard desktop machines available at the time.

Figure 5 shows the dependence of computation time on the time-step size. The total computation time is further subdivided into the time spent solving the gas-phase ODE system and that spent solving the Monte Carlo stochastic particle algorithm. These plots show an approximately first order relationship between computation time and step size, which is expected. Results for both Strang splitting and predictor-corrector are plotted. For $N_{max} = 64$ (left plot) the computation time is dominated by ODE solution time, while for $N_{max} = 65536$ (right plot) the computation time is dominated by the Monte Carlo solver. This is as expected because Monte Carlo computation times generally scale approximately linearly with the number of computation particles. This is an inherent time overhead when performing a Monte Carlo step as internal variables must be set and the LPDA algorithm [18] loops over all particles at least once per step. This explains why the Monte Carlo computation time has a dependence on time-step size.

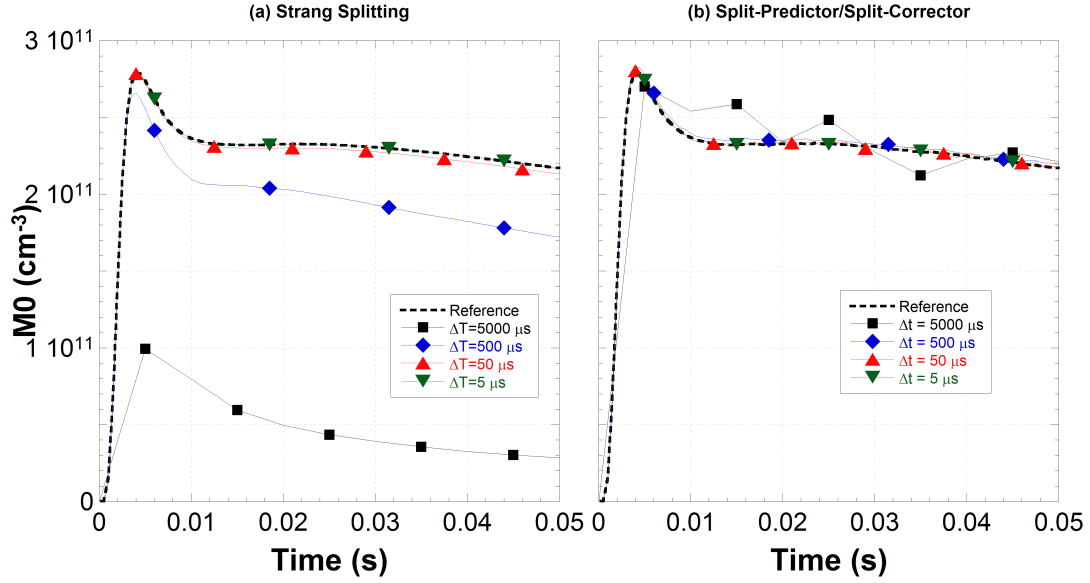


Figure 4: Test system particle number density vs. time for Strang splitting algorithm (left) and predictor-corrector algorithm (right). Simulations were performed with $N_{max} = 65536$ and $L = 100$.

Figure 6 shows the dependence of computation time on the maximum particle count in a similar manner to figure 5. The approximately first order dependence of Monte Carlo computation time on computational particle count is clearly seen. The ODE solution time is unsurprisingly not a function of computational particle count for either algorithm. These plots clearly demonstrate the computational advantage of not re-initialising the ODE solver after each time-step; the ODE solution time for the predictor-corrector algorithm (open symbols) is clearly much lower than that for the Strang splitting algorithm (solid symbols). The difference is far greater at shorter time-step sizes, approximately two orders of magnitude of computation time at $\Delta t = 5 \mu\text{s}$. Additionally it can be seen from figure 6 that this ODE step optimisation is only active for lower values of N_{max} ; when higher N_{max} values are used, the Monte Carlo solution time dominates and the total computation times for both algorithms are similar.

Figures 7 to 9 illustrate the efficiency of both algorithms by highlighting the regions in the error-computation time space occupied by the studied range of the N_{max} (64, 256, 1024, 4096, 16384, 65536) and Δt (5-100 μs) parameters. Only point for which the error is greater than the statistical uncertainty have been shown. Approximate threshold lines of lowest achievable error vs. computation time are drawn to emphasise the differences between the algorithms. It is expected that these lines are the best possible cases because any further increase in N_{max} or decrease in Δt to reduce the error can only result in a higher computation time, in accordance with the results in figures 5 and 6. Several features are common to all three plots. There is a general trend towards smaller errors at larger computation times, which is expected. There also appears to be a limit on the lowest possible computation time, which will be a result of fixed computational costs in

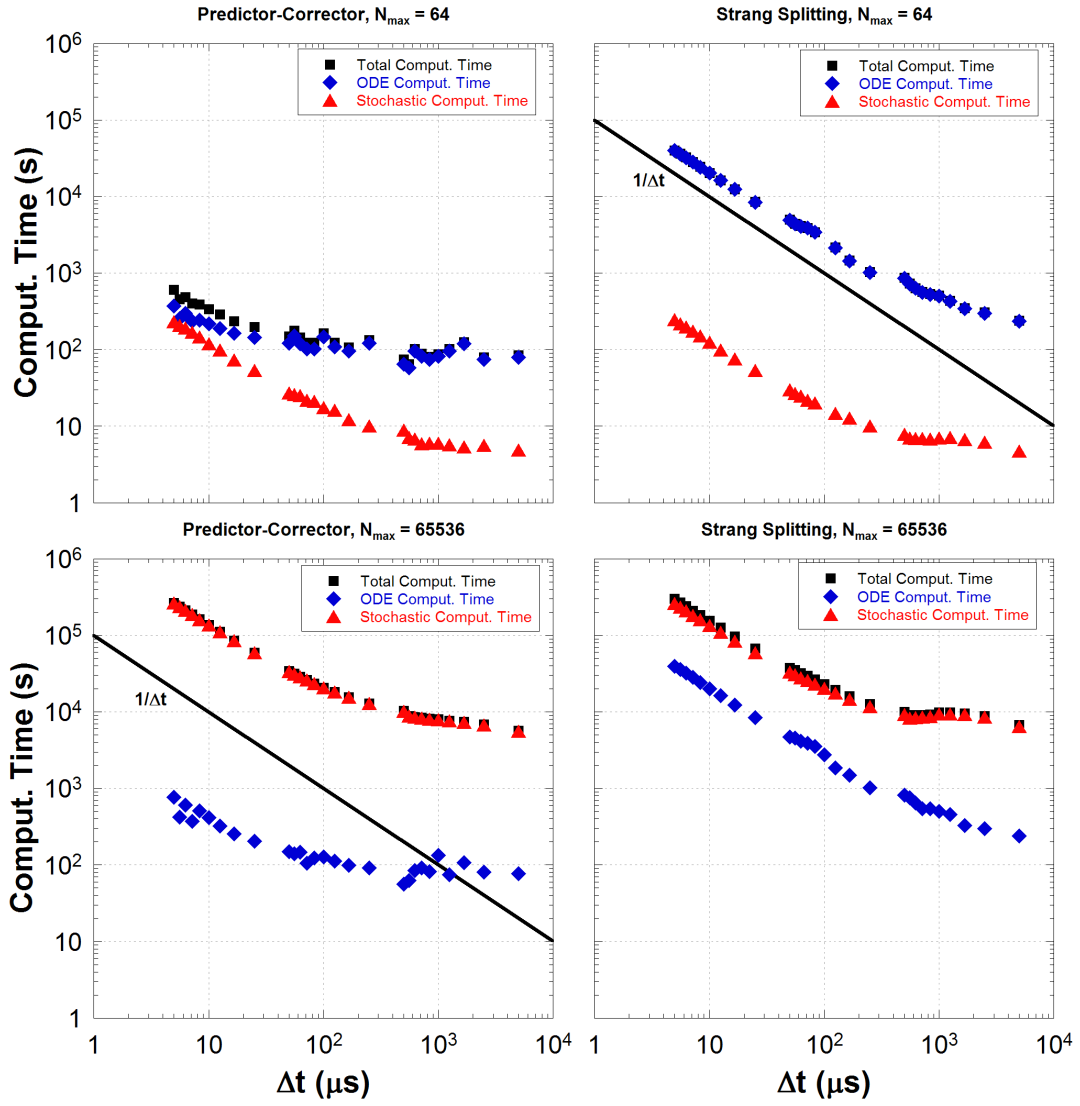


Figure 5: Computation time as a function of time-step size for $N_{max} = 64$ (top) and $N_{max} = 65536$ (bottom). The left-hand plots show results for the predictor-corrector algorithm, and the right-hand plots show results for the Strang splitting algorithm.

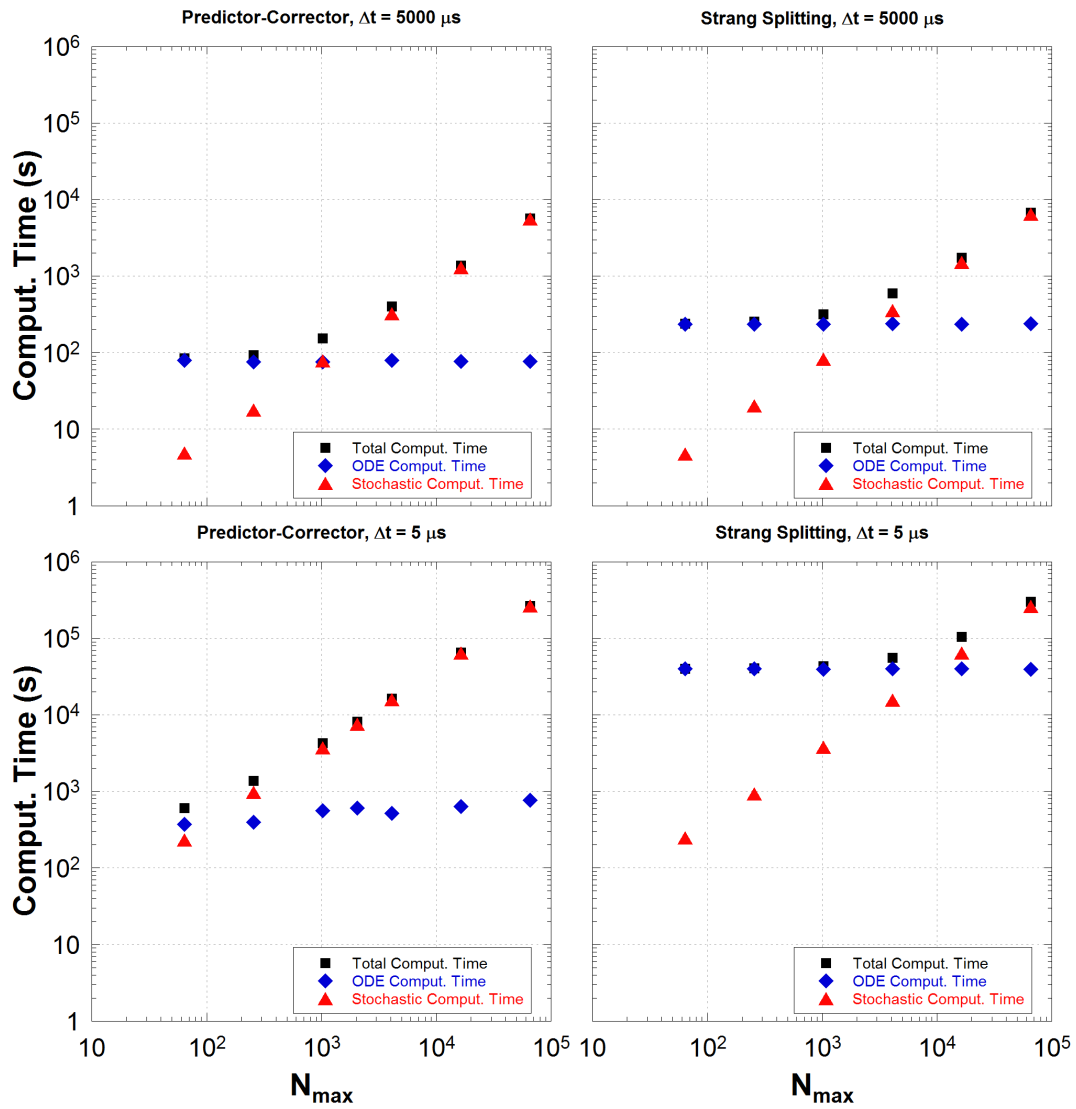


Figure 6: Computation time as a function of maximum particle count for $\Delta t = 5000 \mu\text{s}$ (top) and $\Delta t = 5 \mu\text{s}$ (bottom). The left-hand plots show results for the predictor-corrector algorithm, and the right-hand plots show results for the Strang splitting algorithm.

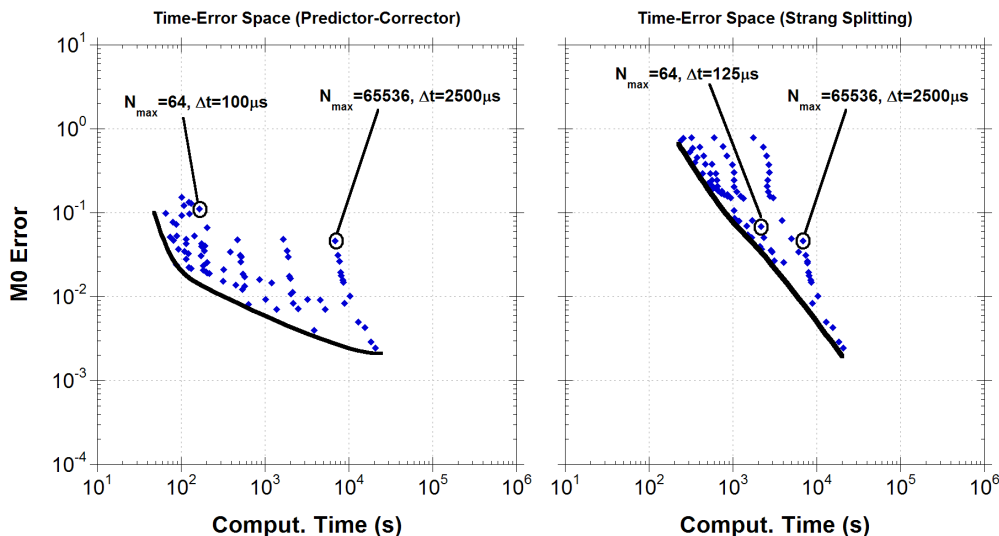


Figure 7: *Computation Time—Error space diagram for particle number density (zeroth moment). The point-cloud shows the error vs. computation time for each parameter set in the studied range (see text). The lines show approximately the lowest achievable error for a given computation time.*

the algorithms not controlled by either parameter.

In all figures it is clear that the predictor-corrector algorithm provides a significant efficiency saving over the Strang splitting algorithm. It is possible to select a fixed error of 10^{-2} (1%) with the predictor-corrector algorithm and achieve an order of magnitude reduction in computation time for functionals $M0$ and Fv , while the Strang splitting algorithm does not achieve that error for $A4$, though an almost two order of magnitude reduction in $A4$ error for a fixed computation time is observed. The predictor-corrector algorithm can achieve $A4$ errors below 10^{-3} , which was the relative error tolerance set in the ODE solver, while the Strang splitting algorithm cannot. The result for $A4$ is particularly encouraging, which must be due to the inclusion of an approximation to the pyrene sink term in the gas-phase ODE system. By removing the source-sink instability a very significant improvement in error is achieved.

5 Conclusions

A predictor-corrector algorithm has been developed and compared to a Strang splitting algorithm for the solution of stiff ODEs coupled to a particle population balance, solved using a stochastic particle algorithm. This new predictor-corrector algorithm incorporates an approximation to the source/sink terms for the ODE variables due to particle processes into the ODE solution. In this way the source-sink instability exhibited by Strang split-

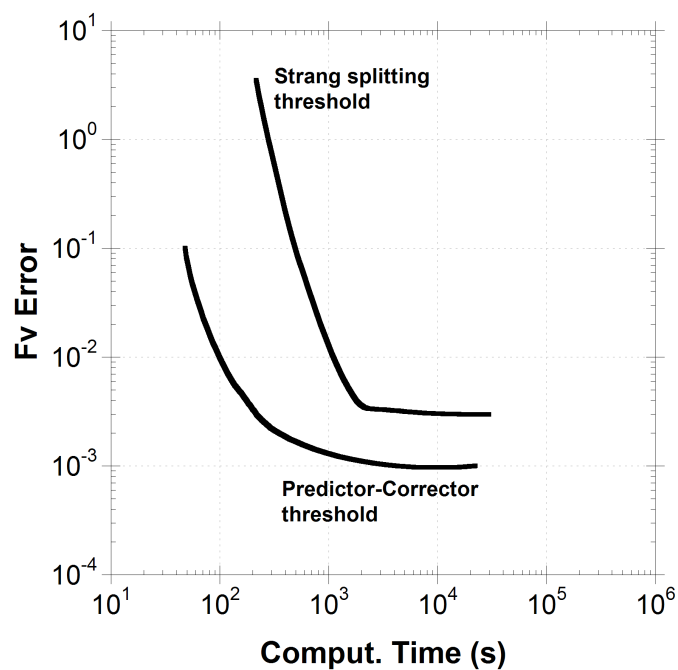


Figure 8: *Computation Time—Error space diagram for soot volume fraction (first moment). The lines show approximately the lowest achievable error for a given computation time. The point-clouds (as in fig. 7) have not been shown for clarity.*

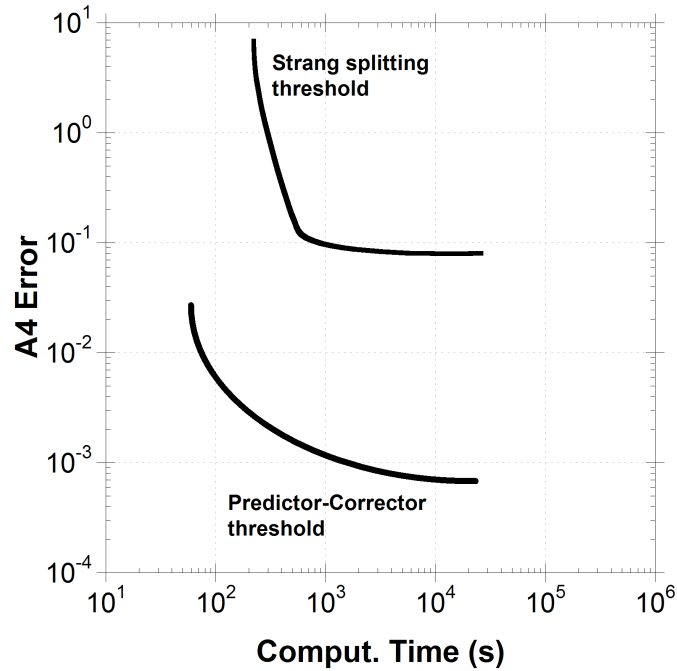


Figure 9: *Computation Time—Error space diagram for gas-phase pyrene concentration (A4). The lines show approximately the lowest achievable error for a given computation time. The point-clouds (as in fig. 7) have not been shown for clarity.*

ting at large splitting steps is eliminated. By removing these source/sink terms from the population balance solver a substantial computation time saving is also achieved. The predictor-corrector algorithm has been shown to converge numerically order 0.5 with maximum computational particle count (N_{max}) and first order with splitting-step size (Δt). The computation times of the predictor-corrector and Strang splitting algorithms have been compared, and it has been shown that the predictor-corrector algorithm can provide substantial computation time savings, greater than one order of magnitude, depending on the choice of parameters. Finally the efficiency of the two algorithms has been compared by plotting the error-computation time space diagram for three functionals. The predictor-corrector algorithm is shown to be substantially more efficient than the Strang splitting algorithm, and can achieve the same error for an order of magnitude reduction in computation time.

References

- [1] J. Appel, H. Bockhorn, and M. Frenklach. Kinetic modeling of soot formation with detailed chemistry and physics: Laminar premixed flames of C₂ hydrocarbons. *Combust. Flame*, 121:122–136, 2000. doi:10.1016/S0010-2180(99)00135-2.
- [2] J. Appel, H. Bockhorn, and M. Wulkow. A detailed numerical study of the evolution of soot particle size distributions in laminar premixed flames. *Chemosphere*, 42: 635–645, 2001. doi:10.1016/S0045-6535(00)00237-X.
- [3] U. M. Ascher and L. R. Petzold. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. SIAM, Philadelphia, PA, 1998. ISBN 0-89871-412-5.
- [4] M. Balthasar and M. Kraft. A stochastic approach to solve the particle size distribution function of soot particles in laminar premixed flames. *Combust. Flame*, 133: 289–298, 2003. doi:10.1016/S0010-2180(03)00003-8.
- [5] A. C. Barone, A. D’Alessio, and A. D’Anna. Morphological characterization of the early process of soot formation by atomic force microscopy. *Combust. Flame*, 132: 181–187, 2003. doi:10.1016/S0010-2180(02)00434-0.
- [6] G. A. Bird. *Molecular gas dynamics*. Clarendon press, 1976.
- [7] K. E. Brenan, S. L. Campbell, and L. R. Petzold. *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*, volume 14 of *Classics in Applied Mathematics*. SIAM, 1995. ISBN 9780898713534.
- [8] M. Celnik, R. West, N. Morgan, M. Kraft, A. Moisala, J. Wen, W. Green, and H. Richter. Modelling gas-phase synthesis of single-walled carbon nanotubes on iron catalyst particles. *Carbon*, 46(3):422–433, 2008. doi:10.1016/j.carbon.2007.12.004.
- [9] M. S. Celnik, R. I. A. Patterson, M. Kraft, and W. Wagner. Coupling a stochastic soot population balance to gas-phase chemistry using operator splitting. *Combust. Flame*, 148(3):158–176, 2007. doi:10.1016/j.combustflame.2006.10.007.
- [10] M. S. Celnik, A. Raj, R. H. West, R. I. A. Patterson, and M. Kraft. An aromatic site description of soot particles. *Combust. Flame*, page in press, 2008.
- [11] A. Eibeck and W. Wagner. An efficient stochastic algorithm for studying coagulation dynamics and gelation phenomena. *SIAM J. Sci. Comput.*, 22(3):802–821, 2000. doi:10.1137/S1064827599353488.
- [12] A. Eibeck and W. Wagner. Stochastic interacting particle systems and nonlinear kinetic equations. *Ann. Appl. Probab.*, 13(3):845–889, 2003. doi:10.1214/aoap/1060202829.
- [13] M. Frenklach. Method of moments with interpolative closure. *Chem. Eng. Sci.*, 57: 2229–2239, 2002. doi:10.1016/S0009-2509(02)00113-6.

- [14] M. Frenklach, C. A. Schuetz, and J. Ping. Migration mechanism of aromatic-edge growth. *Proc. Combust. Inst.*, 30:1389–1396, 2005. doi:10.1016/j.proci.2004.07.048.
- [15] M. J. Height, J. B. Howard, J. W. Tester, and J. B. V. Sande. Flame synthesis of single-walled carbon nanotubes. *Carbon*, 42:2295–2307, 2004. doi:10.1016/j.carbon.2004.05.010.
- [16] D. F. Kronholm and J. B. Howard. Analysis of soot surface growth pathways using published PFR data with new PSD measurements and published premixed flame data. *Proc. Combust. Inst.*, 28:2555–2561, 2000. doi:10.1016/S0082-0784(00)80672-4.
- [17] B. Öktem, M. P. Tolocka, B. Zhao, H. Wang, and M. V. Johnston. Chemical species associated with the early stage of soot growth in a laminar premixed ethylene-oxygen-argon flame. *Combust. Flame*, 142:364–373, 2005. doi:10.1016/j.combustflame.2005.03.016.
- [18] R. I. A. Patterson and M. Kraft. Models for the aggregate structure of soot particles. *Combust. Flame*, 151:160–172, 2007. doi:10.1016/j.combustflame.2007.04.012.
- [19] R. I. A. Patterson, J. Singh, M. Balthasar, M. Kraft, and J. Norris. The Linear Process Deferment Algorithm: A new technique for solving population balance equations. *SIAM J. Sci. Comput.*, (28):303–320, 2006. doi:10.1137/040618953.
- [20] R. I. A. Patterson, J. Singh, M. Balthasar, M. Kraft, and W. Wagner. Extending stochastic soot simulation to higher pressures. *Combust. Flame*, 145(3):638–642, 2006. doi:10.1016/j.combustflame.2006.02.005.
- [21] H. Richter, S. Granata, W. H. Green, and J. B. Howard. Detailed modelling of PAH and soot formation in a laminar premixed benzene/oxygen/argon low-pressure flame. *Proc. Combust. Inst.*, 30:1397–1405, 2005. doi:10.1016/j.proci.2004.08.088.
- [22] K. K. Sabelfeld, S. V. Rogasinsky, A. A. Kolodko, and A. I. Levykin. Stochastic algorithms for solving Smoluchovsky coagulation equation and applications to aerosol growth simulation. *Monte Carlo Methods and Appl.*, 2(1):41–87, 1996.
- [23] M. D. Smooke, C. S. McEnally, L. Pfefferle, R. J. Hall, and M. B. Colket. Computational and experimental study of soot formation in a coflow, laminar diffusion flame. *Combust. Flame*, 117:117–139, 1999.
- [24] A. Violi. Modeling of soot particle inception in aromatic and aliphatic premixed flames. *Combust. Flame*, 139:279–287, 2004. doi:10.1016/j.combustflame.2004.08.013.
- [25] H. Wang and M. Frenklach. A detailed kinetic modeling study of aromatic formation in laminar premixed acetylene and ethylene flames. *Combust. Flame*, 110:173–221, 1997. doi:10.1016/S0010-2180(97)00068-0.

- [26] J. Z. Wen, M. J. Thomson, S. H. Park, S. N. Rogak, and M. F. Lightstone. Study of soot growth in a plug flow reactor using a moving sectional model. *Proc. Combust. Inst.*, 30(1):1477–1484, 2005. doi:[10.1016/j.proci.2004.08.178](https://doi.org/10.1016/j.proci.2004.08.178).