2011-8

On the Acceleration of Explicit Finite Difference Methods for Option Pricing

Stephen O'Sullivan
Technological University Dublin, stephen.osullivan@tudublin.ie

Conall O'Sullivan
University College Dublin

Quantitative Finance

On the acceleration of explicit finite difference methods for option pricing

Stephen O'Sullivan [a] [b] & Conall O'Sullivan [c]

[a] School of Mathematical Sciences, Dublin City University, Glasnevin, Dublin 9, Ireland

[b] UCD School of Mathematical Sciences, University College Dublin, Belfield, Dublin 4,
Ireland

[c] Department of Banking and Finance, The Michael Smurfit School of Business, University
College Dublin, Blackrock, Co. Dublin, Ireland

PLEASE SCROLL DOWN FOR ARTICLE

On the acceleration of explicit finite difference methods for option pricing

STEPHEN O'SULLIVAN*†‡ and CONALL O'SULLIVAN§

†School of Mathematical Sciences, Dublin City University, Glasnevin, Dublin 9, Ireland
‡UCD School of Mathematical Sciences, University College Dublin, Belfield, Dublin 4, Ireland
§Department of Banking and Finance, The Michael Smurfit School of Business,
University College Dublin, Blackrock, Co. Dublin, Ireland

Implicit finite difference methods are conventionally preferred over their explicit counterparts for the numerical valuation of options. In large part the reason for this is a severe stability constraint known as the Courant–Friedrichs–Lewy (CFL) condition which limits the latter class's efficiency. Implicit methods, however, are difficult to implement for all but the most simple of pricing models, whereas explicit techniques are easily adapted to complex problems. For the first time in a financial context, we present an acceleration technique, applicable to explicit finite difference schemes describing diffusive processes with symmetric evolution operators, called Super-Time-Stepping. We show that this method can be implemented as part of a more general approach for non-symmetric operators. Formal stability is thereby deduced for the exemplar cases of European and American put options priced under the Black–Scholes equation. Furthermore, we introduce a novel approach to describing the efficiencies of finite difference schemes as semi-empirical power laws relating the minimal real time required to carry out the numerical integration to a solution with a specified accuracy. Tests are described in which the method is shown to significantly ameliorate the severity of the CFL constraint whilst retaining the simplicity of the underlying explicit method. Degrees of acceleration are achieved yielding comparable, or superior, efficiencies to a set of benchmark implicit schemes. We infer that the described method is a powerful tool, the explicit nature of which makes it ideally suited to the treatment of symmetric and non-symmetric diffusion operators describing complex financial instruments including multi-dimensional systems requiring representation on decomposed and/or adaptive meshes.

Keywords: Numerical methods for option pricing; Black–Scholes model; Computational finance; Equity options; American options; Exotic options

1. Introduction

The Black–Scholes partial differential equation (PDE) (Black and Scholes 1973, Merton 1973) has become a cornerstone of modern derivatives pricing. In general, however, it is rare to find a closed-form solution to the PDE except for the well-known classical cases, such as European call and put options. When a closed-form solution does not exist, for example in the case of American put options, one popular way to proceed is to solve the PDE numerically using finite difference methods (see, for example, Wilmott *et al.* 1995 and Tavella and Randall 2000). Alternative approaches to American option pricing include binomial tree methods or quasi-analytical approximations. The reader is referred to Broadie and Detemple (1996) for a review and comparison of numerical techniques as applied to American put option pricing and Zhu (2006) for a recent semi-closed-form exact solution to the American put option price which consists of an infinite series expansion.

Numerical valuations of options using finite difference methods based on implicit discretizations are usually superior in terms of efficiency to approaches based on conventional explicit discretizations. The principal reason for this is the famous Courant–Friedrichs–Lewy (CFL) stability constraint on explicit schemes which limits the size of the time-step relative to the *square* of the

---

*Corresponding author. Email: steve.osullivan@dcu.ie

spatial step. In this paper the restriction that the CFL constraint imposes is reduced significantly using an acceleration technique for explicit algorithms, known as Super-Time-Stepping (STS), which is completely novel to computational finance. The technique is applied to the problem of pricing European and American put option prices and compared with a number of standard finite difference methods used frequently in the literature. We demonstrate formal stability by appealing to an extended method developed by the authors which does not have the same dependency on having a highly symmetric evolution operator. Furthermore, we propose that this may be the method of choice for more complex pricing models for which high degrees of symmetry in the evolution operator may not be guaranteed.

It is demonstrated that the efficiencies attained are comparable, and often superior, to those of common implicit differencing techniques. Crucially, this acceleration is achieved without any significant increase in implementation complexity relative to the underlying standard explicit scheme.

The paper is divided into six sections. Section 2 reviews a number of standard finite difference methods used frequently in the literature. Section 3 introduces the Super-Time-Stepping technique whose application to options pricing is the main contribution of this paper. Section 4 describes the novel methodology used to compare the different algorithms. Section 5 contains the results and finally in section 6 we offer concluding remarks.

2. Review

This section contains a brief review of the Black–Scholes PDE and the standard finite difference methods, popular in the literature, that are used to solve this PDE for the case of European and American put options.

2.1. Black–Scholes partial differential equation

Let S be the asset price underlying the option at time t whose dynamics are described by a geometric Brownian motion. Let K be the exercise price, T the time to expiry, r the risk-free interest rate, and σ the volatility of the asset. Denoting the option price as $V(S, t)$, the Black–Scholes PDE for European call and put options is then given by

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + rS\frac{\partial V}{\partial S} = rV, \quad 0 \leq t < T, \quad (1)$$

with the payoff function

$$V(S, T) = \begin{cases} \max(S - K, 0), & \text{for a call,} \\ \max(K - S, 0), & \text{for a put,} \end{cases} \quad (2)$$

and the boundary conditions

$$V(0, t) \approx \begin{cases} 0, & \text{for a call,} \\ Ke^{-r(T-t)}, & \text{for a put,} \end{cases} \quad (3)$$

$$V(S_\infty, t) \approx \begin{cases} S, & \text{for a call,} \\ 0, & \text{for a put.} \end{cases} \quad (4)$$

American options can be exercised before expiry and it is never optimal to exercise an American call option before expiry on a stock with no dividend (Merton 1973). However, it may be optimal to exercise an American put option before expiry. The early exercise constraint means that, in the continuation region, the value of an American put option, $V_A(S, t)$, satisfies equation (1). However, the exercise region (where it is optimal to exercise early) induces the following payoff and boundary conditions for the American put option on a stock with no dividend:

$$V(S, T) = \max(K - S, 0), \quad S \geq 0, \quad (5)$$

$$\frac{\partial V}{\partial S}(\overline{S}, t) = -1, \quad (6)$$

$$V(\overline{S}(t), t) = K - \overline{S}(t), \quad (7)$$

$$\lim_{S \to \infty} V(S, t) = 0, \quad (8)$$

$$\overline{S}(T) = K, \quad (9)$$

$$V(S, t) = K - S, \quad 0 \leq S < \overline{S}(t), \quad (10)$$

where $\overline{S}(t)$ represents the free and moving early exercise boundary that separates the continuation region from the early exercise region (Wilmott et al. 1995, Duffie 1996).

2.2. Numerical methods

Discretizing derivatives in the stock price S in equation (1) via three-point central differencing results in a semi-discrete representation

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S_j^2 \left[\frac{\delta_S^2 V_j}{\Delta S^2}\right] + rS_j\left[\frac{\Delta_S V_j}{2\Delta S}\right] - r[V_j] = 0, \quad (11)$$

on the domain $[0, S_\infty] \times [0, T]$, where

$$\delta_S^2 V_j = V_{j+1} - 2V_j + V_{j-1},$$

$$\Delta_S V_j = V_{j+1} - V_{j-1},$$

$$S_j = j\Delta S, \quad \text{for } j = \{0, 1, \dots, J\}. \quad (12)$$

All finite difference methods in the present work will be derived from this equation and will therefore converge as ΔS^2 to the same exact solution of equation (11). It must be noted that the Black–Scholes PDE is not transformed in any way so the following analysis is as general as possible. Let $\Delta t = T/N$ and $t_n = n\Delta t$ for $n = \{0, 1, \dots, N\}$. This notation means $V_j^N = V(S_j, T)$ is the payoff of the option at maturity and $V_j^0 = V(S_j, 0)$ is the option price at $t = 0$ (corresponding to current time). More generally, working backwards recursively through the computational mesh it means that, at time t_{n+1}, the values for V^{n+1} are known and the values for V^n must be found. Admixing equation (11) at time levels

t_n and t_{n+1} results in the semi-discrete form of the well known θ-method,

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S_j^2 \left[\theta \frac{\delta_S^2 V_j^{n+1}}{\Delta S^2} + (1-\theta)\frac{\delta_S^2 V_j^n}{\Delta S^2} \right]$$
$$+ rS_j \left[\theta \frac{\Delta_S V_j^{n+1}}{2\Delta S} + (1-\theta)\frac{\Delta_S V_j^n}{2\Delta S} \right]$$
$$- r[\theta V_j^{n+1} + (1-\theta)V_j^n] = 0. \tag{13}$$

This expression may be written compactly as

$$\frac{\partial \tilde{\mathbf{V}}}{\partial t} - \tilde{\mathbf{P}}[\theta\tilde{\mathbf{V}}^{n+1} + (1-\theta)\tilde{\mathbf{V}}^n] = 0, \tag{14}$$

where $\tilde{\mathbf{P}} \in \mathbb{R}^{J-1} \times \mathbb{R}^{J+1}$ is a tri-diagonal matrix given by

$$\tilde{\mathbf{P}} = \begin{bmatrix} a_1 & b_1 & c_1 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & a_{J-1} & b_{J-1} & c_{J-1} \end{bmatrix}, \tag{15}$$

with

$$a_j = \frac{1}{2}r\frac{S_j}{\Delta S} - \frac{1}{2}\sigma^2 \frac{S_j^2}{\Delta S^2}, \tag{16}$$

$$b_j = r + \sigma^2 \frac{S_j^2}{\Delta S^2}, \tag{17}$$

$$c_j = -\frac{1}{2}r\frac{S_j}{\Delta S} - \frac{1}{2}\sigma^2 \frac{S_j^2}{\Delta S^2}, \tag{18}$$

and where $\tilde{\mathbf{V}}^n \in \mathbb{R}^{J+1}$ is the column vector of option prices at time t_n given by

$$\tilde{\mathbf{V}}^n = \begin{bmatrix} V_0^n \\ \vdots \\ V_J^n \end{bmatrix}. \tag{19}$$

Discretizing the temporal derivative $\partial V/\partial t$ to first order as $(V_j^{n+1} - V_j^n)/\Delta t$ yields the fully discretized form of the θ-method

$$(\mathbf{I} + \Delta t(1-\theta)\tilde{\mathbf{P}})\tilde{\mathbf{V}}^n = (\mathbf{I} - \Delta t\theta\tilde{\mathbf{P}})\tilde{\mathbf{V}}^{n+1}. \tag{20}$$

Applying the boundary conditions at S_0 and S_J means we know the option values V_0^n and V_J^n for all n. Hence we can recast equations (14) and (20) as

$$\frac{\partial \mathbf{V}}{\partial t} - \mathbf{P}[\theta\mathbf{V}^{n+1} + (1-\theta)\mathbf{V}^n] - [\theta\mathbf{b}^{n+1} + (1-\theta)\mathbf{b}^n] = 0, \tag{21}$$

and

$$(\mathbf{I} + \Delta t(1-\theta)\mathbf{P})\mathbf{V}^n + \Delta t(1-\theta)\mathbf{b}^n$$
$$= (\mathbf{I} - \Delta t\theta\mathbf{P})\mathbf{V}^{n+1} - \Delta t\theta\mathbf{b}^{n+1}, \tag{22}$$

respectively, where $\mathbf{P} \in \mathbb{R}^{J-1} \times \mathbb{R}^{J-1}$ is a square non-symmetric matrix consisting of the $J-1$ innermost columns of $\mathbf{P}$

$$\mathbf{P} = \begin{bmatrix} b_1 & c_1 & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & a_{J-1} & b_{J-1} \end{bmatrix}, \tag{23}$$

and where $\mathbf{V}^n$, $\mathbf{b}^n \in \mathbb{R}^{J-1}$ are column vectors given by

$$\mathbf{V}^n = \begin{bmatrix} V_1^n \\ V_2^n \\ \vdots \\ V_{J-2}^n \\ V_{J-1}^n \end{bmatrix}, \tag{24}$$

and

$$\mathbf{b}^n = \begin{bmatrix} a_1 V_0^n \\ 0 \\ \vdots \\ 0 \\ c_{J-1} V_J^n \end{bmatrix}. \tag{25}$$

When $\theta = 1$ this is an explicit scheme that is accurate to $O(\Delta t, \Delta S^2)$. The explicit method is very simple to implement, however the stability of the method depends on the size of the time-step, the spatial step and the coefficients in the PDE (see Wilmott *et al.* 1995 and Tavella and Randall 2000 for further details on these conditions). In particular, it is required that

$$\Delta t \leq \frac{\Delta S^2}{\sigma^2 S_\infty^2}, \tag{26}$$

where $S_\infty = S_J$ is the maximum price on the computational mesh. This is known as the Courant–Friedrichs–Lewy (CFL) stability constraint and may be severely restrictive. By way of illustration, if we want to improve accuracy by halving the spatial step we must reduce the time-step by a factor of 4 and computation time goes up by a factor of 8.

When $\theta = 0$ the scheme is fully implicit and is accurate to $O(\Delta t, \Delta S^2)$. The fully implicit scheme has no limitations on the size of the time-step for the method to converge.

When $\theta = \frac{1}{2}$ the resultant scheme is known as the Crank–Nicolson (CN) method and is accurate to $O(\Delta t^2, \Delta S^2)$. Similarly to the fully implicit method, CN has no limitations on the size of the time-step for stability. CN schemes are therefore a frequently favoured method in the literature.

In this paper we shall additionally employ Richardson Extrapolation (RE) to render explicit and the fully implicit schemes second-order accurate in time. RE is carried out on a step-wise basis as follows.

We assume a smoothly convergent first-order accurate method for the temporal integration of the semi-discrete equation (21) with exact solution $V_{\Delta S}(S, t)$. Given a second-order accurate solution at time level $n+1$ such that $V_j^{n+1} = V_{\Delta S}(j\Delta S, (n+1)\Delta t) + (N-n-1)\mathcal{O}(\Delta t^3)$ we may take a single step of size Δt to approximate the solution at time level n using $V_j^n(\Delta t) = V_{\Delta S}(j\Delta S, n\Delta t) + C\Delta t^2 + \mathcal{O}(\Delta t^3)$ for some constant C. Similarly, taking two steps of size $\Delta t/2$, we have $V_j^n(\Delta t/2) = V_{\Delta S}(j\Delta S, n\Delta t) + (C/2)\Delta t^2 + \mathcal{O}(\Delta t^3)$. Subtracting the expression for $V_j^n(\Delta t)$ from twice the expression for

$V_j^n(\Delta t/2)$ yields a second-order advancement in the solution from time level $n+1$ to level n according to

$$V_j^n = 2V_j^n(\Delta t/2) - V_j^n(\Delta t). \tag{27}$$

This is the prescription of the RE we employ as opposed to the more usual *post-processed* form which requires two independently derived solutions for use in the extrapolation $V_j^0 = 2V_j^0(\Delta t/2) - V_j^0(\Delta t)$ (e.g., Geske and Johnson 1984). The significant difference is that, in the former case, a properly second-order integration method is obtained in the sense that a second-order solution is available at all intermediate times.

The use of RE comes at the expense of an increase in the computational workload. However, it is simple to implement and of greater applicability than CN, as we shall see.

For implicit schemes, the system of simultaneous equations (22) may be solved exactly via direct matrix inversion. For the vanilla option pricing problems under consideration in this work, the Brennan and Schwartz (1977) algorithm may be reformulated to employ LU-decomposition (Ikonen and Toivanen 2007a). While this is an $\mathcal{O}(N)$ method, it does not generalize well. For example, in the relatively simple case of pricing American put options under Heston's stochastic volatility model, the early exercise region must take a specific form for the Brennan and Schwartz algorithm to work. Since we are interested in comparison of methods of general applicability and examine one-dimensional vanilla options for the purposes of bench-testing computational efficiency only, we do not consider the Brennan and Schwartz class of methods any further.

Iterative approaches may also be taken to obtain solutions to within some prescribed accuracy. The most popular of these is Successive Over-Relaxation (SOR); see Crank (1984) for more detail on SOR. In the case of American options, the early exercise constraint requires the use of a variation known as Projected SOR (PSOR) (Wilmott et al. 1995, Ikonen and Toivanen 2007b). Other iterative methods such as the Gauss–Seidel, originally used by Brennan and Schwartz (1977) in the context of finite difference methods applied to American options pricing, and the Jacobi method are also discussed by Wilmott et al. (1995). Using iterative methods can speed up implicit finite difference schemes relative to direct matrix inversion, particularly when fast inversion techniques such as LU-decomposition are unavailable.

When pricing American put options we have to consider the possibility of early exercise. In explicit schemes this is easily handled recursively as follows. Assume the American constraint has been applied at time-step t_{n+1}. The unknown value V_j^n is calculated from the known values V_{j-1}^{n+1}, V_j^{n+1} and V_{j+1}^{n+1} and then we replace V_j^n with $\max(V_j^n, K - S_j)$.

On the other hand, when solving implicitly for the unknowns V_j^n, we must take into account that the value at any point j may be equal to the corresponding continuation values if early exercise is sub-optimal, or the early exercise value if early exercise is optimal.

Therefore, it is clearly inappropriate to first find the continuation value V_j^n and replace it with $\max(V_j^n, K - S_j)$ for all j since all of these values are coupled and any such replacement must be made simultaneously for all j. Because of this, direct matrix inversion methods can be at best first-order accurate in time and iterative algorithms, such as PSOR, are necessary. During each PSOR iteration, the projection is carried out by replacing V_j^n with $\max(V_j^n, K - S_j)$. Over several iterations this converges to simultaneous early exercise enabling convergence at rates above first order in time.

In the next section we introduce an accelerated explicit finite difference scheme, known as Super-Time-Stepping, that is completely novel to computational finance. This scheme reduces significantly the restriction that the CFL constraint imposes on the size of the time-step relative to the spatial step in conventional explicit finite difference schemes. We go on to describe how STS may be implemented as a component of a composite method with stability properties appropriate for more general problems. We invoke RE for second-order accuracy in time. In the subsequent section the accelerated explicit scheme is applied to European and American put options and is demonstrated to be of comparable or superior efficiency to a number of implicit differencing schemes with no significant increase in implementation complexity relative to standard explicit schemes.

3. Acceleration methods

3.1. Super-Time-Stepping

Super-Time-Stepping (STS) is a technique that can be used to accelerate explicit schemes for parabolic problems. In the following, we shall use the description of Alexiades et al. (1996), itself a variant of a method presented by Gentzsch (1979) and essentially a pared-down Runge–Kutta–Chebyshev (RKC) method (van der Houwen 1977, van der Houwen and Sommeijer 1980, Verwer et al. 1990, Verwer 1996, Sommeijer et al. 1997).

Despite the fact that the STS method is approximately 30 years old, it is extraordinary that it has been reported in use by few researchers. The very limited number of numerical investigations we are aware of employing STS are in engineering and physical disciplines and include: nonlinear degenerate convection–diffusion (Evje et al. 2001); electromagnetic wave scattering (Shi et al. 2006); isotropic and anisotropic diffusion on biological membranes (Sbalzarini et al. 2006); and magnetic field diffusion in astrophysics (Mignone et al. 2007, O'Sullivan and Downes 2007). To our knowledge, it has received no attention to date in the area of finance.

The essence of STS is that rather than requiring stability at each step of the integration, N_{STS} sub-steps of varying size Δt_j are rolled together into a single super-step Δt_{STS} according to

$$\Delta t_{\mathrm{STS}} = \sum_{j=1}^{N_{\mathrm{STS}}} \Delta t_j, \tag{28}$$

and stability is only demanded at the end of the super-step.† To proceed, we assume a *linear* scheme on $\mathbf{V} \in \mathbb{R}^M$ of the form

$$\mathbf{V}^n = (\mathbf{I} - \Delta t_{\mathrm{STS}}\mathbf{A})\mathbf{V}^{n+1} = \left[\prod_{j=1}^{N_{\mathrm{STS}}}(\mathbf{I} - \Delta t_j\mathbf{A})\right]\mathbf{V}^{n+1}, \quad (29)$$

where the solutions at time levels $n+1$ and n are known and unknown, respectively, $\mathbf{I}$ is the identity matrix and $\mathbf{A} \in \mathbb{R}^M \times \mathbb{R}^M$ is a symmetric positive definite matrix. It is well known that, for stability, we must have

$$\left|\prod_{j=1}^{N_{\mathrm{STS}}}(\mathbf{I} - \Delta t_j\lambda)\right| < 1, \quad (30)$$

for all eigenvectors λ of $\mathbf{A}$.

The properties of Chebyshev polynomials of degree N_{STS} (Markoff 1916) then allow us to explicitly enforce stability *while maximizing* Δt_{STS} to provide a set of optimal values for the sub-steps given by

$$\Delta t_j = \Delta t_{\mathrm{STD}}\left[(-1 + \nu)\cos\left(\frac{2j-1}{N_{\mathrm{STS}}}\frac{\pi}{2}\right) + 1 + \nu\right]^{-1}, \quad (31)$$

where Δt_{STD} is the normal explicit time-step limit and ν is a damping factor. Note, in particular, that

$$\Delta t_{\mathrm{STS}} \rightarrow N_{\mathrm{STS}}^2\Delta t_{\mathrm{STD}}, \quad \text{as } \nu \rightarrow 0. \quad (32)$$

While this scheme is stable for any choice of N_{STS} given a large enough value of ν, in practice a balance can be struck between the two parameters to optimize the performance of the scheme. The method is unstable in the limit $\nu = 0$. We illustrate the efficacy of the acceleration process for $N_{\mathrm{STS}} = 30$ in figure 1. It can be seen that the first substep may be up to 25 times the stable limit for a standard explicit integration as $\nu \rightarrow 0$, but subsequent substeps become increasingly small. The effect of this is a cumulative error cancellation that recovers stability over the composite superstep. Crucially, there is a net payoff in terms of the size of the superstep with respect to N_{STS} steps of size Δt_{STD} as described by equation (32).

Note that although formal results only exist for linear schemes, there is ample evidence, as described above, that nonlinear target systems are equally amenable to the STS method.

It can be shown (Alexiades *et al.* 1996) that STS is essentially first order in time. It is not possible to introduce additional temporal structure to an STS step since intermediate values obtained during a STS cycle are physically meaningless and may not be used as approximations to the solution in any sense. Therefore, predictor–corrector-style methods are not applicable should higher-order convergence be required. On the other hand, we have found that RE works perfectly well. By this method all the advantages of the first-order STS method are easily transferred to second- (or higher-) order schemes.
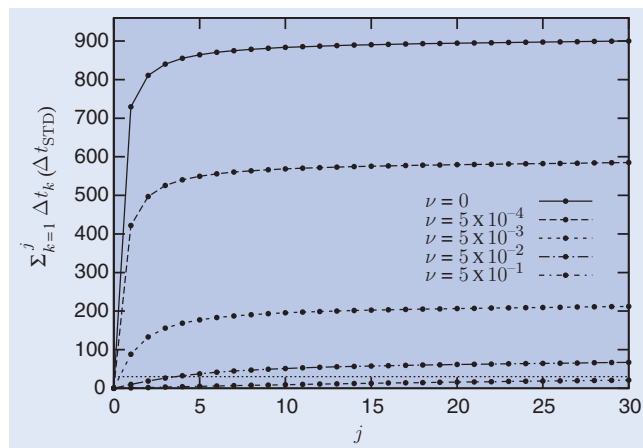

Figure 1. Illustration of acceleration via STS. Accumulated time $\sum_{k=1}^{j}\Delta t_k$ is shown in units of standard explicit time-step Δt_{STD} over single superstep Δt_{STS} with $N_{\mathrm{STS}} = 30$ for a range of damping factors ν. A reference line at $N_{\mathrm{STS}}\Delta t_{\mathrm{STD}}$ indicates the time attained over N_{STS} unaccelerated (standard explicit) steps. Note that acceleration approaches N_{STS} times this value as $\nu \rightarrow 0$, in agreement with equation (32). Note also that deceleration occurs for the highest considered damping factor of $\nu = 0.5$. In this work, $\nu = 5 \times 10^{-4}$ and $N_{\mathrm{STS}} = 30$ for all tests.

The principal advantage of the STS method is not efficiency, however, but simplicity. Explicit discretizations of even the most complex systems of parabolic equations are very straightforward. In particular, implementation of adaptive mesh refinement (AMR) technologies and/or parallelization via domain decomposition techniques present no great challenges from within an explicit framework. On the contrary, when implicit methods are involved, tackling problems of even a moderate level of complexity can be an exceedingly intricate task.

In the next section the performance of the STS method applied via equation (22) is compared with finite difference schemes described in section 2.

3.2. Composite

Since stability is assured by the above analysis for symmetric positive definite $\mathbf{A}$ in equation (29), we provide a formal stability analysis for an alternative discretization of equation (1) in appendix A. The scheme presented therein is formally stable under application of STS to a split symmetric positive definite operator. In practice, we find that, as suggested by other authors previously (e.g., Alexiades *et al.* 1996 and O'Sullivan and Downes 2007), this alternative scheme is not strictly necessary when the evolution operator has a dominant symmetric component. For the cases under consideration here, in fact, we find it has negligible impact and therefore it is not used.

We wish to emphasize that, while the composite scheme is employed in this work for the purposes of stability analysis, its greater generality may prove it to be the appropriate choice for the numerical integration of

---

†It has been claimed by Verwer (1996) that factorized RKC methods are impractical as they suffer from severe internal instability. We find no evidence of this for $N_{\mathrm{STS}} \lesssim 30$.

systems of equations for which the evolution operator does not display the high degree of symmetry observed for pricing vanilla options under Black–Scholes.

4. Methodology

In terms of the real-world usefulness of a numerical scheme X for pricing financial instruments, it is important to recognize that the wall-time (the real time taken to carry out the numerical integration) W_X required for a scheme to perform a calculation will depend on the maximum permitted error E of the solution. Since the greatest error between the numerical solution V_j and the exact solution $V(S_j)$ arises when S_j is the equal to the exercise price K, where there is a discontinuity in the first derivative of the payoff function, we use the definition $E \equiv |V_k^0 - V(K, 0)|$ for the error in the numerical solution (where the index k is the index corresponding to the spot price $S_k = K$). For consistency with this measure, the SOR schemes are required to converge to within a user-prescribed tolerance on V_k^n rather than over a range of values as is more usual. From this point on, we shall generally omit the X subscript from quantities associated with a scheme X for clarity of notation. Dependent variables will be implicitly assumed to be associated with a given scheme.

The schemes that are used as benchmarks for comparison with the STS explicit scheme with Richardson extrapolation (STS_RE) are as follows: the standard explicit scheme with RE (STD_RE), the fully implicit scheme that uses matrix inversion and Richardson extrapolation (INV_RE), the fully implicit scheme that uses SOR and RE (SOR_RE), the CN scheme that uses matrix inversion (INV_CN) and the CN scheme that uses SOR (SOR_CN). Scheme comparison is traditionally performed by presenting timings for solutions obtained to non-uniform accuracies, although Broadie and Detemple (1996) go further by considering simultaneously timings and accuracy of the pricing methods they considered. Presenting timings for solutions obtained to non-uniform accuracies is not a fair comparison of the efficiencies as clearly a scheme which has run to a higher accuracy has expended greater computational resources than necessary to achieve a lower value. The efficiency of finite difference schemes which are derived from a single semi-discrete representation (in this case equation (11), the semi-discrete representation of equation (1)) may be represented by three parameters for a given problem, with an additional three parameters necessary to impose the CFL condition on explicit schemes. We now proceed to elucidate this idea by deriving semi-empirical functional forms for the wall-time W as a function of the error E.

All schemes under consideration here are second order accurate in price by construction from equation (11). Assuming order p in time the total error is given by

$$E = E_t + E_S, \qquad (33)$$

where we say

$$E_t = O(\Delta t^p) \approx \zeta \Delta t^p, \qquad E_S = O(\Delta S^2) \approx \eta \Delta S^2 \quad (34)$$

are the errors arising from the temporal and price discretization, respectively. Assuming W scales inversely with Δt we have

$$W \propto E_t^{-1/p}, \qquad (35)$$

and hence

$$W = \beta \Delta t^{-1}, \qquad (36)$$

for some β. In fact, we know $p = 2$ in all cases bar INV_CN applied to the American put option where $p = 1$. Furthermore, if the scaling of the wall-time with ΔS is given by the power law

$$W \propto \Delta S^{-2\psi}, \qquad (37)$$

we have

$$W \propto E_S^{-\psi}. \qquad (38)$$

As we shall see later, for the schemes considered in this paper, implicit schemes have $\psi \gtrsim 1.1$ and explicit schemes have $\psi \lesssim 0.4$. Note that since all schemes are approximating the solution to the same semi-discrete equation (11), η depends only on the problem and not the scheme.

We can write

$$W = \alpha E_t^{-1/p} E_S^{-\psi}, \qquad (39)$$

where the parameters (α, ψ) are fit by experiment for a given test case as follows. To determine ψ for each scheme, we run a series of tests for a fixed value of Δt at different price resolutions ΔS. Once ψ is known, a series of tests is run for a fixed value of ΔS at different time resolutions Δt. The exact solution to equation (11), $V_{\Delta S}$, is estimated via an independent second-order explicit code (Ødegaarde 2007) run to very high accuracy. This code was minimally modified by the authors to attain second-order accuracy in time via RE. The value obtained is then used to estimate $E_t = |V_k(K, 0) - V_{\Delta S}(K, 0)|$ from the experimental data which may then be fit to

$$W = \beta \zeta^{1/p} E_t^{-1/p}. \qquad (40)$$

From equation (39),

$$\beta = \alpha \zeta^{-1/p} E_S^{-\psi}, \qquad (41)$$

and hence α follows directly given $E_S = |V(K, 0) - V_{\Delta S}(K, 0)|$. The exact solution to equation (1), $V(K, 0)$, is derived to high accuracy using an analytical model for European options or, for American options, an independent binomial tree code (Ødegaarde 2007) run to extremely high accuracy.

Minimizing the wall-time W for a fixed total error E determines the dependencies of E_t and E_S on E according to the relations

$$E_t = \frac{1}{1 + \psi p} E, \qquad E_S = \frac{\psi p}{1 + \psi p} E. \qquad (42)$$

Equation (39) may then be recast in the form

$$W = \alpha \frac{(1 + \psi p)^{\psi+1/p}}{(\psi p)^{\psi}} E^{-\psi-1/p}. \qquad (43)$$

From this efficiency equation, an approximation to the wall-time required for a scheme X to evaluate the solution for a given problem to a maximum specified accuracy E is characterized by three parameters (α, ψ, p). In a subsequent section we shall derive these parameters for a set of schemes applied to given option pricing problems. Firstly, however, it is necessary to consider an additional constraint on the viable choices for the pair (E_t, E_S) in optimizing the efficiency of explicit schemes.

4.1. Stability constraints on explicit schemes

For STD_RE, the CFL condition (equation (26)) requires $\Delta t \leq \Delta S^2/(\sigma^2 S_{\infty}^2)$ where S_{∞} is the maximum price on the computational mesh. For a given number of uniformly spaced price mesh points J, the maximum stable time-step is $1/(\sigma^2 J^2)$. Therefore, in an integration over a time T, a minimum number of time-steps $T\sigma^2 J^2$ is required. More generally, for an explicit scheme X this constraint will determine the scheme efficiency for $E \leq E_{\text{crit}}$. The critical point is the lower bound of the regime described by equations (42) where the CFL condition is marginally satisfied, i.e.

$$\Delta t = \epsilon \Delta S^2. \qquad (44)$$

Using equations (34) we may then write

$$E_{\text{crit}} = \left[\frac{1}{\zeta} \left(\frac{\eta}{\epsilon} \right)^p \right]^{1/(p-1)} \frac{1 + \psi p}{(\psi p)^{p/(p-1)}}. \qquad (45)$$

Note that if $p = 1$ then as long as $\eta/4\zeta \leq \epsilon$, equations (42) trivially satisfy the CFL condition for all E.

We have theoretical values for ϵ given by

$$\epsilon_{\text{STD\_RE}}^{\text{theory}} = 1/\sigma^2 S_{\infty}^2, \qquad (46)$$

and

$$\epsilon_{\text{STS\_RE}}^{\text{theory}} = \frac{N_{\text{STS}}}{2\sqrt{\nu}} \left(\frac{(1 + \sqrt{\nu})^{2N_{\text{STS}}} - (1 - \sqrt{\nu})^{2N_{\text{STS}}}}{(1 + \sqrt{\nu})^{2N_{\text{STS}}} + (1 - \sqrt{\nu})^{2N_{\text{STS}}}} \right) \epsilon_{\text{STD\_RE}}. \qquad (47)$$

We shall see later that, in practice, $\epsilon_{\text{STD\_RE}} \approx \epsilon_{\text{STD\_RE}}^{\text{theory}}$ but $\epsilon_{\text{STS\_RE}} > \epsilon_{\text{STS\_RE}}^{\text{theory}}$. Thus the theoretical value is conservative in the case of STS_RE.

For $E < E_{\text{crit}}$, the relationship between E_t and E_S is no longer freely tunable for optimal efficiency but set by the marginal CFL condition $\Delta t = \epsilon \Delta S^2$ giving $(E_t/\zeta)^{1/p} = \epsilon E_S/\eta$. For $p = 2$, $E = E_t + E_S$ may be written as a quadratic equation in E_S with one admissible root,

$$\frac{\zeta \epsilon^p}{\eta^p} E_S^p + E_S - E = 0. \qquad (48)$$

Equation (39) may then be written as a function of E_S,

$$W = \frac{\alpha \eta}{\zeta^{1/p} \epsilon} E_S^{-1-\psi}. \qquad (49)$$

5. Efficiency tests

In the following we shall consider European and American put options with $T = 1$, $\sigma = 0.2$, $r = 0.05$ and $K = 100$. A uniform mesh is assumed in S over the range $[0, 5K]$. The authors have confirmed similar results for other choices of T, σ, and r and for a uniform mesh in $\log(S)$. All tests are carried out using MATLAB R2007a on a 2.2 GHz Intel Core 2 Duo processor under Fedora 8 linux.

In all tests, SOR_CN and SOR_RE are tuned with an over-relaxation parameter $w = 1.1$ and tolerance set to 10^{-12}. No maximum iteration count is assigned so that this tolerance is always achieved. For STS_RE, in all tests $\nu = 5 \times 10^{-4}$ and $N_{\text{STS}} = 30$. These values formally imply a limiting stable time-step 380.2961 times longer than that for STD_RE. Convergence to the solution of equation (1) is measured against an analytic solution for European put options and, for American put options, against a high accuracy solution ($E \approx 10^{-6}$) obtained from a binomial tree method. To measure convergence rates to the exact solution of equation (11) for a given ΔS, very high accuracy solutions from an independent second-order explicit scheme (Ødegaarde 2007) are used ($\Delta t = T/10^7$).

5.1. Convergence studies: determining ζ, η, p

Convergence studies of all schemes for European and American put option pricing are presented in figures 2 and 3, respectively.

The scaling of E_S is obtained for each scheme for fixed Δt and $\Delta S = \{5K/(100 \times 2^n) \mid n = 0, 1, 2, 3, 4, 5, 6\}$. Given a sufficiently small Δt, the temporal error E_t will be negligible compared with the price error E_S and $E = |V(K, 0) - V_j(K, 0)| \approx E_S$. We find $\Delta t = T/1000$ for the implicit solvers fulfills this condition. Stability over the full range of spatial resolutions demands lower values for the explicit schemes: $\Delta t = T/3600$ for STS_RE and $\Delta t = T/1\,660\,000$ for STD_RE. In the upper panel of figure 2, second-order convergence with ΔS is confirmed for all schemes pricing the European option. Note that, at the very lowest values of ΔS, the temporal error begins to dominate. While second-order convergence is also clear from the upper panel of figure 3, the early exercise constraint results in temporal errors that dominate the spatial errors at larger values of ΔS than in the European case. In particular, the temporal error for INV_CN is dominant at even moderate values of ΔS because the early exercise constraint reduces this scheme to first-order accuracy in time when an iterative solver is not used (Wilmott *et al.* 1995).

The temporal order of convergence is determined for a fixed price spacing $\Delta S = 5K/100$ and a range of temporal resolutions $\Delta t = \{T/(20 \times 2^n) \mid n = 0, 1, \ldots, 12, 13\}$. In the case of STD_RE, the stable temporal resolutions are restricted and we consider $\Delta t = \{T/(20 \times 2^n) \mid n = 5, 6, \ldots, 14, 15\}$. The temporal error E_T is then determined from $V_{\Delta S = 5K/100}$.

The lower panel of figure 2 illustrates the case of the European option pricing test case. All schemes show
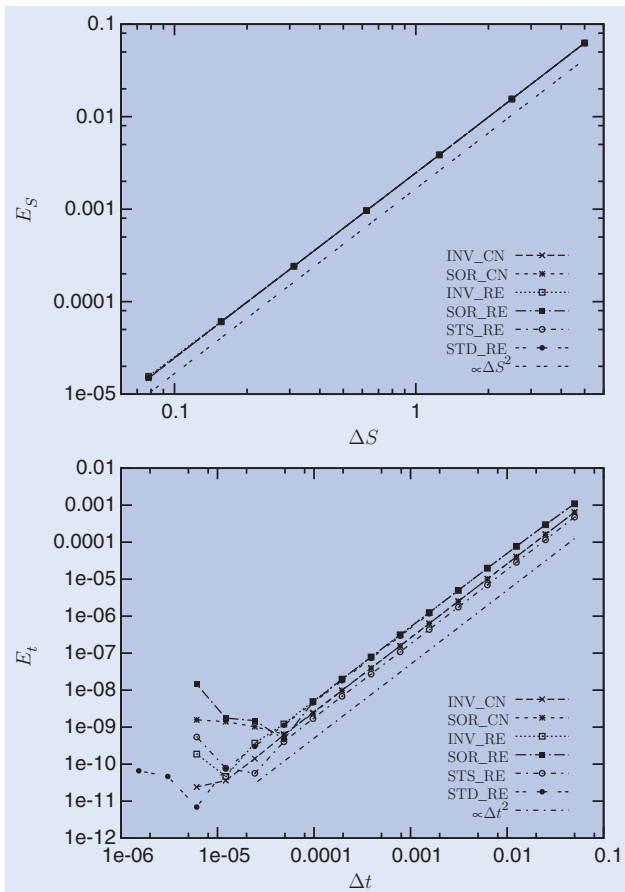
Figure 2. Scheme convergence tests for European put option pricing. Upper panel: Spatial error at the strike price, E_S, plotted as a function of ΔS. For this test, errors are measured against a value obtained from an exact solution to the Black–Scholes equation. The influence of temporal errors in the numerical integration is minimized by using a sufficiently small value for Δt in all cases. Points are almost exactly coincident in all cases. Clearly, all tested schemes converge to the exact solution at second order according to $E_s = \eta \Delta S^2$ for some η. Values for η are obtained from this data as necessary. A reference line scaling with ΔS^2 is shown to guide the eye. Lower panel: Temporal error at the strike price, E_t, plotted as a function of Δt for $\Delta S = 5K/100$. Temporal errors are calculated using a high-accuracy approximation to the exact solution of equation (11) obtained from an independent code modified to be second order in time (Ødegaarde 2007). It is clear that all schemes converge to this value with second-order accuracy according to $E_t = \zeta \Delta t^2$, for some ζ, down to levels where machine accuracy becomes important. Values for ζ are obtained from this data as required. A reference line scaling with Δt^2 is also shown to guide the eye.

Figure 3. Scheme convergence tests for American put option pricing. Upper panel: Spatial error at the strike price, E_S, plotted as a function of ΔS. For this test, errors are measured against a high-accuracy approximation to the exact solution using a binomial tree method. The influence of temporal errors in the numerical integration is minimized by using a sufficiently small value for Δt in all cases. Down to small ΔS all schemes, except INV_CN, which is first order in time, demonstrate convergence to the estimated exact solution at second order according to $E_s = \eta \Delta S^2$ for some η. Values for η are obtained from this data as necessary. A reference line scaling with ΔS^2 is shown to guide the eye. Lower panel: Temporal error at the strike price, E_t, plotted as a function of Δt for $\Delta S = 5K/100$. Temporal errors are calculated using a high-accuracy approximation to the exact solution of equation (11) obtained from an independent code modified to be second order in time (Ødegaarde 2007). Except for INV_CN, which shows first-order temporal accuracy, all schemes converge to this value with approximate second-order accuracy according to $E_t = \zeta \Delta t^2$, for some ζ. Values for ζ are obtained from this data as required. Reference lines scaling with Δt and Δt^2 are also shown to guide the eye.

second-order convergence with time down to levels where round-off error becomes significant. The results are qualitatively similar for the American case as shown in the lower panel of figure 3 except that, again, the early exercise constraint degrades the smoothness of convergence in all cases except INV_CN, which, as previously commented on, is reduced to first-order accuracy in time. We find good experimental agreement for $p = 2$ in all cases except INV_CN applied to the American option problem where $p = 1$, as expected.

The power laws given by equations (34) are fit to the well-behaved data-points; as previously observed, other sources of error become significant at the lowest values of
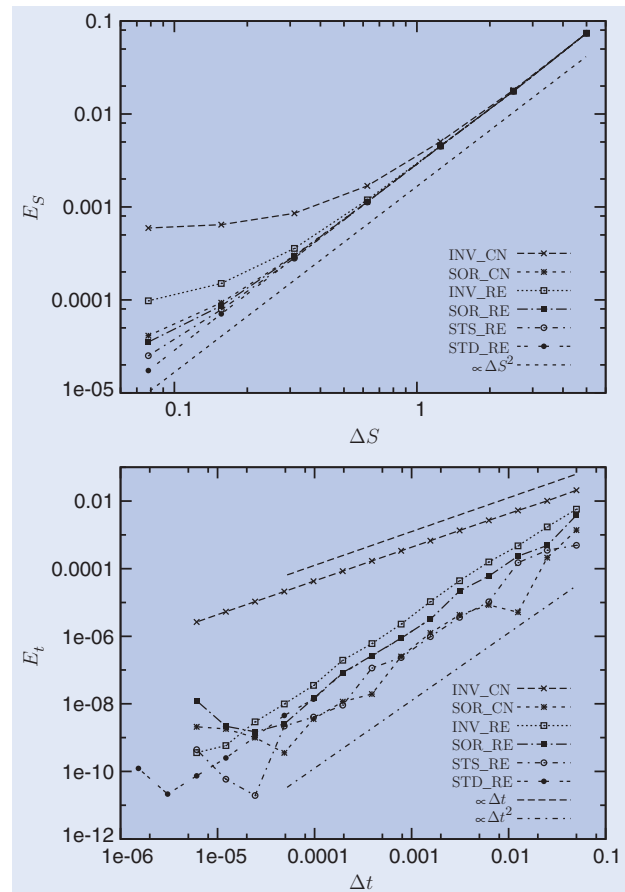
Δt and ΔS. The best-fit values obtained for ζ and η are used in the following sections to evaluate the dependent parameters as required. The values corresponding to the explicit schemes are detailed in table 1 as they are required to apply the CFL constraint. This is discussed later in section 5.4.

5.2. Wall-time scaling with price resolution: determining ψ

In this section, the scaling of the wall-time required to achieve a given accuracy as a function of the spot price

Figure 4. Temporal error and wall-time scaling with price resolution for European put option pricing. Upper panel: Temporal error at the strike price, E_t, plotted as a function of ΔS for a fixed value of $\Delta t = T/1280$ (except for $\Delta t = T/640\,000$ in the case of STD_RE). Errors are calculated using a high-accuracy approximation to the exact solution of equation (11) obtained from an independent code modified to be second order in time (Ødegaarde 2007). To a fair approximation, the temporal error has no dependence on the spatial resolution. Lower panel: The wall-time W in seconds as a function of price resolution ΔS for the tests described above. This data is used to fit the scaling law $W \propto \Delta S^{-2\psi}$. At moderate to high resolution the law is clearly well obeyed. Reference lines scaling with ΔS^{-1} and ΔS^{-2} are also shown.

resolution is assessed. The range of values used for the price mesh spacing is $\Delta S = \{5K/(250 \times 2^n) \mid n = 0, 1, 2, 3, 4\}$. For all schemes, $\Delta t = T/1280$, except for STD_RE where $\Delta t = T/640\,000$ is required for stability. Error and wall-time data for European and American option pricing are presented in figures 4 and 5, respectively. The upper panel in each case shows the dependence of the temporal error E_t on ΔS and the lower panels illustrate the wall-time W scaling. Notably, the SOR schemes tend to diverge from a simple power law fit at high values of ΔS, and at low values for E_t in the European case. The temporal error E_t is particularly erratic for STS_RE applied to the American option. This does not reflect any inherent difficulty with the scheme, but merely indicates that the power law approximation is not well suited to STS_RE with the early exercise constraint applied in this way. We defer discussion of the scheme's performance under the more modern penalty method approach to
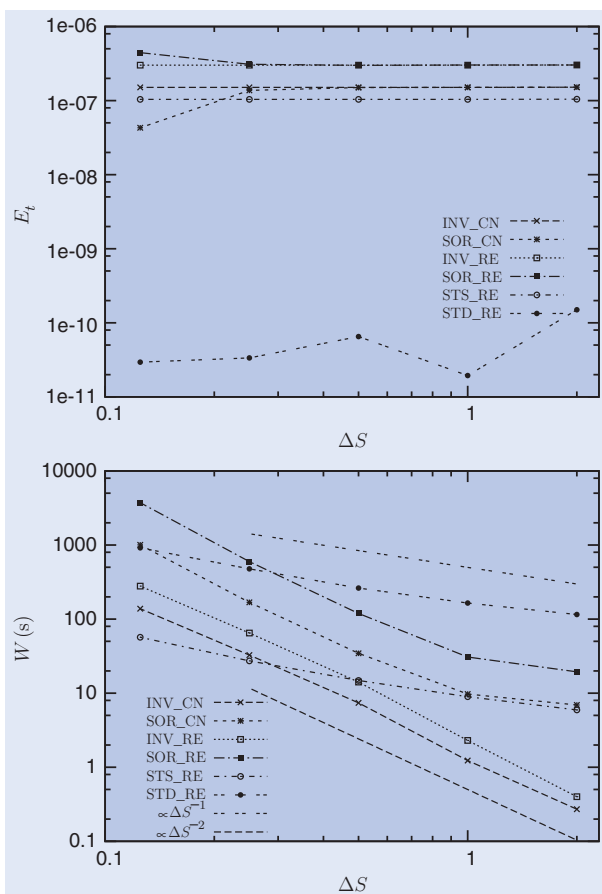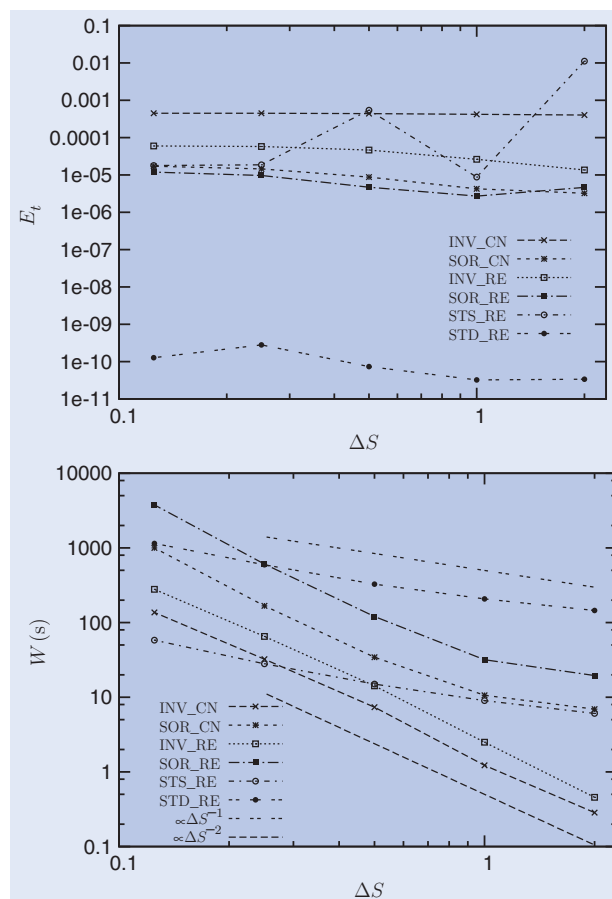

Figure 5. Temporal error and wall-time scaling with price resolution for American put option pricing. Upper panel: Temporal error at the strike price, E_t, plotted as a function of ΔS for a fixed value of $\Delta t = T/1280$ (except for $\Delta t = T/640\,000$ in the case of STD_RE). Errors are calculated using a high-accuracy approximation to the exact solution of equation (11) obtained from an independent code modified to be second order in time (Ødegaarde 2007). Except for STS_RE at low to moderate values of ΔS, which shows substantial volatility in E_t due to the implementation of the early exercise constraint, the temporal error is reasonably approximated as independent of the spatial resolution. Lower panel: The wall-time W in seconds as a function of price resolution ΔS for the tests described above. This data is used to fit the scaling law $W \propto \Delta S^{-2\psi}$ and is almost identical to the data presented for the European case in figure 4. At moderate to high resolution the law is clearly well obeyed. Reference lines scaling with ΔS^{-1} and ΔS^{-2} are also shown.

pricing American options, for example Nielsen *et al.* (2002). Power law indices are fit to the well-behaved datapoints according to $W \propto \Delta S^{\psi}$ and tabulated in table 2. Reference lines are plotted for some values in the corresponding figure panels.

Clearly, from the upper panel of figure 4, the temporal accuracy of all schemes is independent of the mesh spacing in price.

5.3. Wall-time scaling with temporal resolution: determining α

We now proceed to establish the values of α using equations (40) and (41). The temporal error E_t and

Figure 6. Temporal error and wall-time scaling with time resolution for European put option pricing. Temporal error E_t at the strike price (upper panel), and wall-time W in seconds (lower panel) plotted against Δt for a fixed value of $\Delta S = 5K/500$. Reference value are indicated for STD_RE with $\Delta t = 10^{-4}$ as it is unstable over the plotted range of Δt. Errors are calculated using a high-accuracy approximation to the exact solution of equation (11) obtained from an independent code modified to be second order in time (Ødegaarde 2007). This data from the upper panel is used to fit β to the scaling law $\beta = \alpha \zeta^{-1/p} E_S^{-\psi}$. Reference lines illustrating the approximate scaling laws $E_t \propto \Delta t^2$ and $W \propto \Delta t^{-1}$ are presented in the upper and lower panels, respectively. Note that $p = 2$ in all cases.
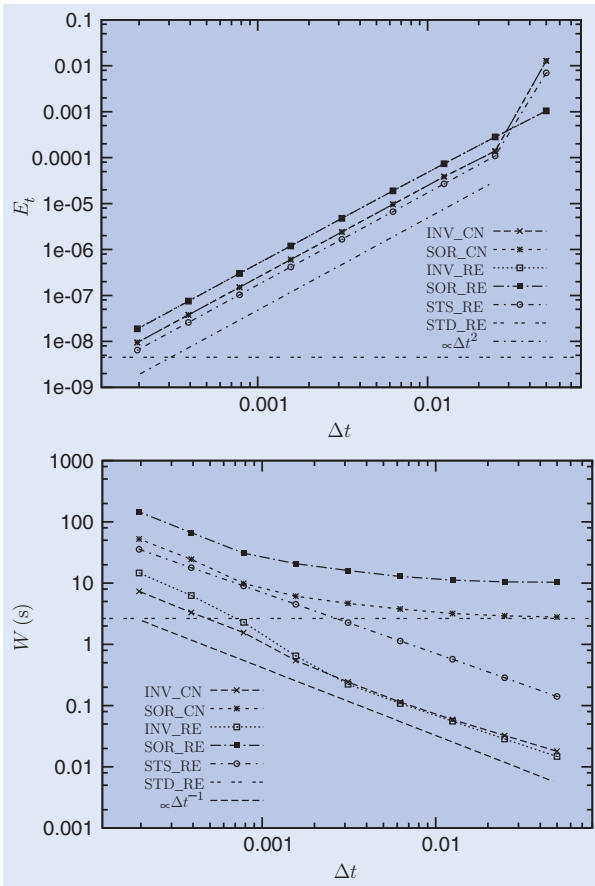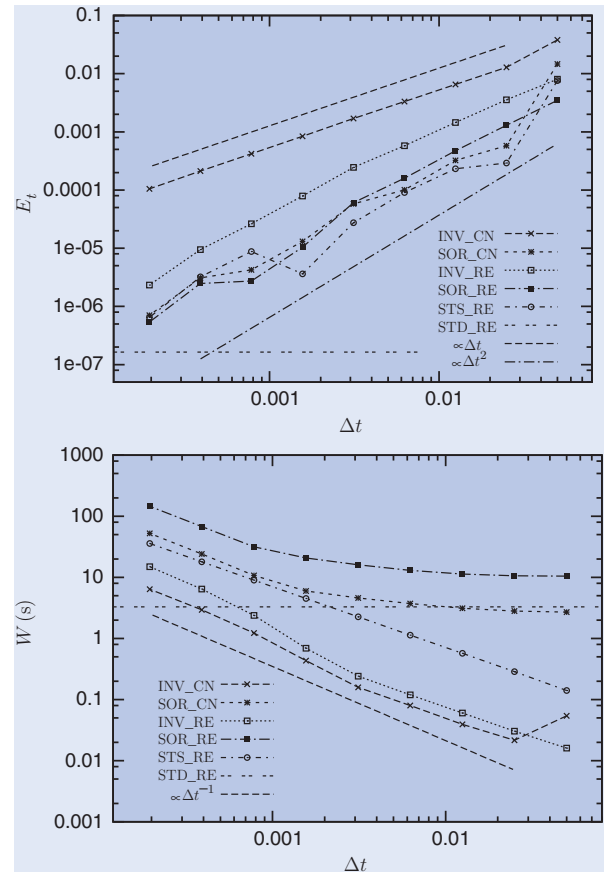
Figure 7. Temporal error and wall-time scaling with time resolution for American put option pricing. Temporal error E_t at the strike price (upper panel), and wall-time W in seconds (lower panel) plotted against Δt for a fixed value of $\Delta S = 5K/500$. Reference value are indicated for STD_RE with $\Delta t = 10^{-4}$ as it is unstable over the plotted range of Δt. Errors are calculated using a high-accuracy approximation to the exact solution of equation (11) obtained from an independent code modified to be second order in time (Ødegaarde 2007). This data from the upper panel is used to fit β to the scaling law $\beta = \alpha \zeta^{-1/p} E_S^{-\psi}$. Reference lines illustrating the approximate scaling laws $E_t \propto \Delta t^p$ and $W \propto \Delta t^{-1}$ are presented in the upper and lower panels, respectively. Note that $p = 2$ in all cases except for INV_CN where $p = 1$.

Table 1. Additional semi-empirical efficiency function parameters for CFL condition compliance of explicit schemes: (ϵ, ζ, η).

Method	European			American		
	ϵ	ζ	η	ϵ	ζ	η
STS_RE	0.0591	0.1791	0.0025	0.0591	0.4368	0.0031
STD_RE	0.0001	0.4787	0.0025	0.0001	1.7397	0.0031

Table 2. Parameters for semi-empirical efficiency functions: (α, ψ, p).

Method	European			American		
	α	ψ	p	α	ψ	p
INV_CN	5.1672e − 07	1.1361	2	4.1479e − 07	1.1266	1
SOR_CN	5.6559e − 06	1.1188	2	7.6095e − 06	1.0978	2
INV_RE	6.3219e − 07	1.1852	2	2.5394e − 06	1.1607	2
SOR_RE	1.9042e − 05	1.1539	2	4.0585e − 05	1.1497	2
STS_RE	2.6033e − 04	0.4063	2	4.2956e − 04	0.4078	2
STD_RE	1.9161e − 05	0.3758	2	4.8678e − 05	0.3740	2

wall-time W are plotted in figures 6 and 7 for the European and American option pricing problems, respectively. A range of temporal resolutions is considered, $\Delta t = \{T/(20 \times 2^n) \mid n = 0, 1, \ldots, 7, 8\}$, for the single spot price spacing value $\Delta S = 1$. The upper panel in each case shows the temporal error E_t while the lower panel illustrates the wall-time W. STD_RE is not stable for any values of Δt in this range and so a reference value is plotted for the minimal stable value of $\Delta t = 10^{-4}$. Table 3 explicitly provides all the errors and timings for this set of tests as well as the benchmark values from which $E_S = |V(K,0) - V_{\Delta S}(K,0)|$ is derived.

From the plots of E_t in figures 6 and 7, the temporal order of accuracy, p, for each scheme is 2, except for INV_CN in the American option pricing case where $p = 1$, as has already been confirmed. Again, the early exercise constraint results in less smooth convergence

Table 3. Numerical value $V(K)$ at strike price $S = K$ with wall-time W in seconds, and temporal error E_t derived from high-accuracy benchmark value (denoted Bench). $\Delta S = 5K/500$ in all cases. Exact solution for European put option obtained from Black–Scholes equation (denoted BS) or high-accuracy solution for American put option derived from binomial tree method (denoted BT) also shown for reference.

N	Method	European put			American put		
		$V(K)$	W (s)	E_t	$V(K)$	W (s)	E_t
	BS/BT	5.5735260222			6.0903702250		
	Bench	5.5710548584			6.0874933186		
20	INV_CN	5.5582337988	0.0180	1.3e − 02	6.0494286083	0.0543	3.8e − 02
	SOR_CN	5.5582337980	2.7834	1.3e − 02	6.0728821885	2.7081	1.5e − 02
	INV_RE	5.5700055968	0.0147	1.0e − 03	6.0794452317	0.0160	8.0e − 03
	SOR_RE	5.5700055955	10.3012	1.0e − 03	6.0840121118	10.4513	3.5e − 03
	STS_RE	5.5640341499	0.1405	7.0e − 03	6.0799065958	0.1407	7.6e − 03
40	INV_CN	5.5711933839	0.0326	1.4e − 04	6.0746091394	0.0216	1.3e − 02
	SOR_CN	5.5711933832	2.9173	1.4e − 04	6.0869159940	2.8169	5.8e − 04
	INV_RE	5.5707713289	0.0284	2.8e − 04	6.0839310254	0.0304	3.6e − 03
	SOR_RE	5.5707713279	10.4519	2.8e − 04	6.0861990167	10.6100	1.3e − 03
	STS_RE	5.5709447761	0.2850	1.1e − 04	6.0872009272	0.2859	2.9e − 04
80	INV_CN	5.5710935500	0.0593	3.9e − 05	6.0809446348	0.0393	6.5e − 03
	SOR_CN	5.5710935496	3.1840	3.9e − 05	6.0871705422	3.1208	3.2e − 04
	INV_RE	5.5709809065	0.0559	7.4e − 05	6.0860502803	0.0601	1.4e − 03
	SOR_RE	5.5709809058	11.1796	7.4e − 05	6.0870218427	11.3238	4.7e − 04
	STS_RE	5.5710278176	0.5702	2.7e − 05	6.0872609232	0.5698	2.3e − 04
160	INV_CN	5.5710645318	0.1140	9.7e − 06	6.0841716547	0.0801	3.3e − 03
	SOR_CN	5.5710645311	3.7710	9.7e − 06	6.0873924444	3.7226	1.0e − 04
	INV_RE	5.5710359539	0.1083	1.9e − 05	6.0869151135	0.1193	5.8e − 04
	SOR_RE	5.5710359531	12.8102	1.9e − 05	6.0873331963	12.9760	1.6e − 04
	STS_RE	5.5710481552	1.1305	6.7e − 06	6.0875839962	1.1340	9.1e − 05
320	INV_CN	5.5710572768	0.2433	2.4e − 06	6.0858018523	0.1582	1.7e − 03
	SOR_CN	5.5710572766	4.6564	2.4e − 06	6.0874354049	4.6028	5.8e − 05
	INV_RE	5.5710500779	0.2224	4.8e − 06	6.0872476891	0.2408	2.5e − 04
	SOR_RE	5.5710500788	16.0405	4.8e − 06	6.0874324039	16.0467	6.1e − 05
	STS_RE	5.5710531897	2.2636	1.7e − 06	6.0874658348	2.2593	2.7e − 05
640	INV_CN	5.5710554631	0.5446	6.0e − 07	6.0866566365	0.4338	8.4e − 04
	SOR_CN	5.5710554630	6.0882	6.0e − 07	6.0874801971	5.9971	1.3e − 05
	INV_RE	5.5710536564	0.6477	1.2e − 06	6.0874145168	0.6917	7.9e − 05
	SOR_RE	5.5710536564	20.7108	1.2e − 06	6.0874829213	20.7715	1.0e − 05
	STS_RE	5.5710544422	4.4999	4.2e − 07	6.0874969531	4.5181	3.6e − 06
1280	INV_CN	5.5710550096	1.5557	1.5e − 07	6.0870731887	1.2253	4.2e − 04
	SOR_CN	5.5710550091	9.8862	1.5e − 07	6.0874890604	10.7403	4.3e − 06
	INV_RE	5.5710545571	2.2893	3.0e − 07	6.0874669422	2.3884	2.6e − 05
	SOR_RE	5.5710545561	30.8410	3.0e − 07	6.0874905967	31.6979	2.7e − 06
	STS_RE	5.5710547545	9.0119	1.0e − 07	6.0875020780	9.0130	8.8e − 06
2560	INV_CN	5.5710548963	3.3009	3.8e − 08	6.0872818618	2.9442	2.1e − 04
	SOR_CN	5.5710548963	24.5080	3.8e − 08	6.0874902341	24.0889	3.1e − 06
	INV_RE	5.5710547830	6.2863	7.5e − 08	6.0874838183	6.4247	9.5e − 06
	SOR_RE	5.5710547835	66.6839	7.5e − 08	6.0874908470	67.6991	2.5e − 06
	STS_RE	5.5710548325	17.8931	2.6e − 08	6.0874965528	17.9378	3.2e − 06
5120	INV_CN	5.5710548679	7.3201	9.4e − 09	6.0873884923	6.4074	1.0e − 04
	SOR_CN	5.5710548680	52.3461	9.5e − 09	6.0874926123	52.5042	7.1e − 07
	INV_RE	5.5710548396	14.6579	1.9e − 08	6.0874909960	14.9043	2.3e − 06
	SOR_RE	5.5710548399	145.4471	1.9e − 08	6.0874927814	145.9531	5.4e − 07
	STS_RE	5.5710548520	35.6537	6.5e − 09	6.0874926711	35.9482	6.5e − 07e
10 000	STD_RE	5.5710548540	2.6428	4.5e − 09	6.0874931545	3.2914	1.6e − 07

profiles for the American option. Best-fit values for β are obtained using equation (41).

High temporal accuracy solutions obtained using the binomial tree method yield $E_S = 0.0024711638$ for the European option and $E_S = 0.0028769064$ for the American option. Then $\alpha = \beta E_S^{\psi}$ follows from equation (41). The efficiency parameters (α, ψ, p) are now fully determined and presented in table 2.

The assumption of inverse scaling of W with Δt given by equation (36) is supported by the lower panels of

figures 6 and 7 except for the SOR schemes at low temporal resolution where the convergence properties of these schemes are poor.

5.4. Stability constraint for explicit schemes: determining ϵ

As discussed in sections 2.2 and 4.1, explicit schemes are limited by the CFL condition given by equation (26). The critical error, E_{crit}, below which this constraint applies, is derived from equation (45). ζ and η, as discussed in section 5.1, have been determined and are presented in table 1. All that remains to characterize the CFL constraint is to establish values for ϵ.

For the cases under consideration here, with $\sigma = 0.2$, $S_\infty = 5K$, and $K = 100$, equation (46) agrees with experiment to within 1% and we take $\epsilon_{\mathrm{STD\_RE}} = \epsilon_{\mathrm{STD\_RE}}^{\mathrm{theory}} = 10^{-4}$. On the other hand, STS_RE turns out to be particularly well suited to this problem and is stable well beyond the value of $\epsilon_{\mathrm{STS\_RE}}^{\mathrm{theory}} = 380\epsilon_{\mathrm{STD\_RE}}$ obtained from equation (47) with $N_{\mathrm{STS}} = 30$ and $\nu = 5 \times 10^{-4}$. In fact, we find experimentally that $\epsilon_{\mathrm{STS\_RE}} = 591\epsilon_{\mathrm{STD\_RE}}$. Presented in table 1, we now have (ϵ, ζ, η) as required for the CFL constraint on the explicit schemes STS_RE and STD_RE.

5.5. Semi-empirical efficiency functions

We now present the semi-empirical evaluation of the wall-time W in seconds, for each of the six schemes, as a function of the error E for the European and American variants of the option pricing problem under consideration. Equation (43) is valid for all cases except for the explicit schemes below the critical error E_{crit} (equation (45)) when the appropriate expression becomes equation (49) due to the CFL condition for stability. The parameters (α, ψ, p), required in the former instance, and (ϵ, ζ, η), as required in the latter instance, are presented in tables 1 and 2 respectively.

The results of the estimated optimal scheme wall-times are plotted in figure 8 in the range $10^{-6} \le E \le 10^{-2}$. The break in the power law can clearly be seen for STS_RE at $E_{\mathrm{crit}} = 0.0269$ in the European option plot and at 0.0168 for the American case. No such break is visible for the STD_RE lines because $E_{\mathrm{crit}} > 10^3$ in both cases: well above the maximum plotted error value.

Before interpreting these results it is important to emphasize that the results are biased heavily in favour of the implicit schemes. In particular, for the direct matrix inversion methods, the necessary matrix inversions are carried out only once before the timer is started. The inversions are usually slower than the time integration itself for $\Delta S \ll 1$, and, in general, the coefficients may be time varying and therefore this procedure will be required at each time-step. Secondly, for the SOR methods at $\Delta t \gtrsim 10^{-3}$, the wall-time to convergence plateaus at a value in excess of the assumed power law form.

Even with these advantages, STS_RE is clearly the optimal performer when taken as a broad spectrum method for this problem. At all but low accuracies for
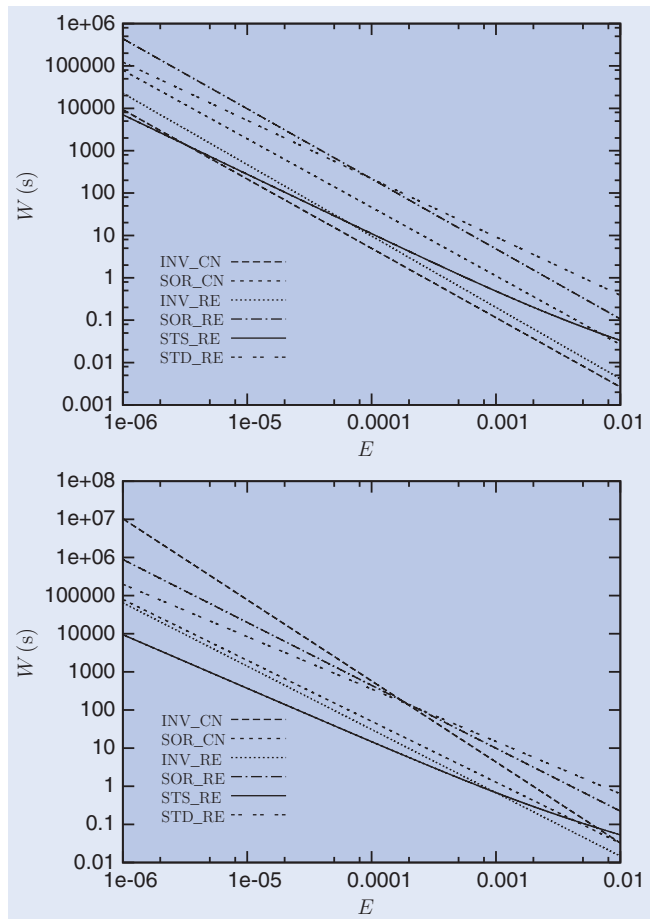

Figure 8. Evaluation of semi-empirical efficiency functions for European and American put options with $\sigma = 0.2$, $r = 0.05$, $T = 1.0$, $S_\infty = 5K$, and $K = 100$. Wall-time W in seconds is plotted against required total error E in the solution for value at strike price $V(K)$. Upper panel: European put option. For $E \gtrsim 10^{-4}$, STS_RE is inferior only to the direct matrix inversion methods, otherwise it is equivalent or superior. Note the break at $E_{\mathrm{crit}} = 0.0269$ in the case of STS_RE below which the optimal relationship between the spatial and temporal errors, E_S and E_t, is constrained by the CFL condition. Lower panel: American put option. In this test, STS_RE is superior for $E \lesssim 10^{-2}$. Note the break at $E_{\mathrm{crit}} = 0.0168$ for STS_RE.

some of the alternative implicit methods where wall-time is low in any case, it is the most efficient option. When the simplicity of the method and the above provisos are taken into consideration, the case for using STS_RE as opposed to any of the presented implicit methods is overwhelming.

6. Conclusion

An acceleration technique, known as Super-Time-Stepping (STS), for explicit finite difference algorithms is introduced for the first time in computational finance. We demonstrate the efficacy of the method by pricing European and American put options in a series of bench-tests with several well-known finite difference techniques. Simple vanilla options are chosen as case studies for their inherent simplicity. However, in practice, the target application for this method is the numerical modeling

of more complex systems for which implicit methods quickly become prohibitively difficult to implement. For example, we suggest multi-dimensional systems requiring decomposed and/or adaptive meshes. Stability is formally demonstrated for a novel operator split implementation of the technique. This composite method is of greater generality and it is suggested that it may be the appropriate approach when considering pricing of more exotic financial instruments.

A novel methodology to assess and compare the schemes' efficiencies is also introduced. Applying this technique to test cases we demonstrate degrees of acceleration provided by the STS method which yield comparable, and even superior, efficiencies to implicit differencing methods. The implicit methods considered are PSOR and direct matrix inversion in both European and American cases. Of central importance, this is achieved with no significant increase in implementation complexity over and above that of the underlying standard explicit algorithm.

Given that STS accelerated methods inherit the simplicity of explicit methods whilst achieving high accuracy at low computational cost, we conclude that when faced with complex pricing models this approach offers a highly attractive alternative to the substantial challenges presented by conventional implicit techniques. Promising targets include models involving multi-dimensional parameter spaces, variable meshes, or moving boundaries.

Acknowledgements

References

Alexiades, V., Amiez, G. and Gremaud, P., Super-Time-Stepping acceleration of explicit schemes for parabolic problems. *Com. Numer. Meth. Engng.*, 1996, **12**, 31–42.

Black, F. and Scholes, M.S., The pricing of options and corporate liabilities. *J. Polit. Econ.*, 1973, **81**, 637–654.

Brennan, M. and Schwartz, E., The valuation of American put options. *J. Finance*, 1977, **32**, 449–462.

Broadie, M. and Detemple, J., American option valuation: new bounds, approximations, and a comparison of existing methods. *Rev. Financial Stud.*, 1996, **9**, 1211–1259.

Crank, J., *Free and Moving Boundary Problems*, 1984 (Clarendon Press: Oxford).

Duffie, D., *Dynamic Asset Pricing Theory*, 1996 (Princeton University Press: Princeton).

Evje, S., Karlsen, K.H. and Risebro, N.H., A continuous dependence result for nonlinear degenerate parabolic equations with spatially dependent flux function. In *Hyperbolic Problems: Theory, Numerics, Applications*, edited by H. Freistüler and G. Warnecke, pp. 337–346, 2001 (Birkhäuser).

Gentzsch, W., Numerical solution of linear and non-linear parabolic differential equations by a time-discretisation of third order accuracy. In *Proceeding of the Third GAMM Conference on Numerical Methods in Fluid Mechanics*, edited by E.H. Hirschel, pp. 109–117, 1979 (Friedr. Vieweg & Sohn).

Gerschgorin, S., Über die Abgrenzung der Eigenwerte einer Matrix. *Izv. Akad. Nauk SSSR, Ser. Mat.*, 1931, **7**, 749–754.

Geske, R. and Johnson, H.E., The American put option valued analytically. *J. Finance*, 1984, **39**, 1511–1524.

Hager, W.W., *Applied Numerical Linear Algebra*, 1988 (Prentice-Hall: Englewood Cliffs, NJ).

Ikonen, S. and Toivanen, J., Pricing American options using LU decomposition. *Appl. Math. Sci.*, 2007a, **1**, 2529–2551.

Ikonen, S. and Toivanen, J., Efficient numerical methods for pricing American options under stochastic volatility. *Numer. Meth. Partial Differ. Eqns*, 2007b, **24**, 104–126.

Markoff, W., Über Polynome, die in einem gegebenen Intervalle möglichst wenig von null abweichen. *Math. Ann.*, 1916, **77**, 213–258, translated by J. Grossman from original Russian article published in 1892.

Merton, R.C., The theory of rational option pricing. *Bell J. Econ. Mgmt Sci.*, 1973, **4**, 141–183.

Mignone, A., Bodo, G., Massaglia, S., Matsakos, T., Tesileanu, O., Zanni, C. and Ferrari, A., PLUTO: a numerical code for computational astrophysics. *ApJS*, 2007, **170**, 228.

Morton, K.W., Stability of finite difference approximations to a diffusion–convection equation. *Int. J. Numer. Meth. Engng*, 1980, **15**, 677–683.

Nielsen, B., Skavhaug, O. and Tveito, A., Penalty and front-fixing methods for the numerical solution of American option problems. *J. Comput. Finance*, 2002, **5**, 69–97.

Ødegaarde, B.A., Financial numerical recipes in C++, 2007. Available online at: http://finance.bi.no/~bernt/gcc_prog/recipes/recipes.pdf

O'Sullivan, S. and Downes, T.P., An explicit scheme for multifluid magnetohydrodynamics. *MNRAS*, 2006, **366**, 1329.

O'Sullivan, S. and Downes, T.P., A three-dimensional numerical method for modelling weakly ionized plasmas. *MNRAS*, 2007, **376**, 1648.

Sbalzarini, I.F., Hayer, A., Helenius, A. and Koumoutsakos, P., Simulations of (an)isotropic diffusion on curved biological surfaces. *Biophys. J.*, 2006, **90**, 8478–8485.

Shi, Y., Li, L. and Liang, C.H., Multidomain pseudospectral time-domain algorithm based on super-time-stepping method. *IEE Proc. Microwaves, Antennas Propagation*, 2006, **153**, 55–60.

Sommeijer, B.P., Shampine, L.F. and Verwer, J.G., RKC: An explicit solver for parabolic PDEs. Technical Report MAS-R9715, CWI Amsterdam, 1997.

Tavella, D. and Randall, C., *Pricing Financial Instruments: The Finite Difference Method*, 2000 (Wiley: New York).

van der Houwen, P.J., *Construction of Integration Formulas for Initial Value Problems*, 1977 (North-Holland: Amsterdam).

van der Houwen, P.J. and Sommeijer, B.P., On the internal stability of explicit m-stage Runge–Kutta methods for large values of m. *Z. Angew. Math. Mech.*, 1980, **60**, 479–485.

Verwer, J.G., Explicit Runge–Kutta methods for parabolic partial differential equations. *Appl. Numer. Math.*, 1996, **22**, 359–379.

Verwer, J.G., Hundsdorfer, W.H. and Sommeijer, B.P., Convergence properties of the Runge–Kutta–Chebyshev method. *Numer. Math.*, 1990, **57**, 157–178.

Wilmott, P., Howison, S. and Dewynne, J., *Option Pricing: Mathematical Models and Computation*, 1995 (Oxford Financial Press: Oxford).

Zhu, S.P., An exact and explicit solution for the valuation of American put options. *Quant. Finance*, 2006, **6**, 229–242.

Appendix A: Composite acceleration method

The fully explicit ($\theta = 1$) expression of equation (22) is

$$\mathbf{V}^n = (\mathbf{I} - \Delta t \mathbf{P})\mathbf{V}^{n+1} - \Delta t \mathbf{b}^{n+1}, \qquad (A1)$$

with all quantities defined as in section 2.2.

Although in sample cases STS applied to a non-symmetric operator has been shown to result in a slight flattening of the stability region (Alexiades *et al.* 1996), formally stability may only be established for the symmetric positive definite case. For the problems considered in the present work, we have found that STS may be applied to the weakly non-symmetric matrix $\mathbf{P}$ with negligible impact on the scheme's stability. For completeness, however, we now proceed to determine conditions for strict stability of an alternate finite differencing of the Black–Scholes equation which does incorporate a symmetric positive definite operator. This scheme formally admits application of the STS acceleration method.

$\mathbf{P}$ may be trivially decomposed according to

$$\mathbf{P} = \mathbf{A} + \mathbf{K}, \qquad (A2)$$

where $\mathbf{A} \equiv (\mathbf{P} + \mathbf{P}^{\mathrm{T}})/2$ and $\mathbf{K} \equiv (\mathbf{P} - \mathbf{P}^{\mathrm{T}})/2$ are the symmetric and skew-symmetric parts of $\mathbf{P}$, respectively, with $\cdot^{\mathrm{T}}$ denoting a transpose operation. Explicitly, these matrices are given by

$$\mathbf{A} = \begin{bmatrix} b_1 & \frac{1}{2}(a_2 + c_1) & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & \frac{1}{2}(a_{J-1} + c_{J-2}) & b_{J-1} \end{bmatrix}, \qquad (A3)$$

and

$$\mathbf{K} = \begin{bmatrix} 0 & -\frac{1}{2}(a_2 - c_1) & \cdots & 0 \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & \frac{1}{2}(a_{J-1} - c_{J-2}) & 0 \end{bmatrix}. \qquad (A4)$$

Prescribing the update of $\mathbf{V}^n$ to $\mathbf{V}^{n+1}$ via Multiplicative Operator Splitting (MOS) we write

$$\mathbf{V}^n = \mathbf{RTV}^{n+1} - \Delta t \mathbf{b}^{n+1}. \qquad (A5)$$

Strict stability requires that $\|\mathbf{RT}\| \leq 1$ for any operator norm (Morton 1980). By multiplicative dominance $\|\mathbf{RT}\| \leq \|\mathbf{R}\| \|\mathbf{T}\|$. For stability we therefore demand $\|\mathbf{R}\| \leq 1$ and $\|\mathbf{T}\| \leq 1$ for operators $\mathbf{R}$ and $\mathbf{T}$ consistent with equation (A1). We shall adopt the Euclidean operator norm, denoted $\|\cdot\|_2$, hereafter.

Note that, for any complex square matrix $\mathbf{X}$, we may write $\|\mathbf{X}\|_2 = \sqrt{\rho(\mathbf{X}^{\dagger}\mathbf{X})}$, where $\mathbf{X}^{\dagger}$ denotes the complex conjugate of $\mathbf{X} \in \mathbb{C}^n \times \mathbb{C}^n$ and $\rho(\cdot)$ indicates the spectral radius of an operator (i.e. the maximum of the absolute values of the operator's eigenvalues).

We propose $\mathbf{R} = \mathbf{I} - \Delta t \mathbf{A}$, which is symmetric by construction and clearly consistent with equation (A1). We now proceed to establish the positive definiteness and stability properties for this definition.

By Sylvester's criterion, a matrix is positive definite if and only if the determinants of all upper-left sub-matrices are positive. The determinant, D_j, for the upper-left $j \times j$ sub-matrix of the tri-diagonal matrix (23) is easily expressed by means of a simple three-term recurrence relation (e.g., Hager 1988)

$$D_j = b_j D_{j-1} - \frac{1}{4}(a_j + c_{j-1})^2 D_{j-2}, \qquad (A6)$$

with $D_{-1} \equiv 0$ and $D_0 \equiv 1$. It is straightforward to show that

$$D_j = \sum_{k=0}^{j} d_{jk} \sigma^{2k} r^{j-k}, \qquad (A7)$$

where d_{jk} is a positive real number for all values of k.

Under the condition $r > 0$, $\mathbf{A}$ is therefore strictly positive-definite. Note that this result does not depend on uniform $\sigma(S)$ and $r(S)$. Incidentally, this result implies that $\mathbf{P}$ must be positive stable.

In order to determine the stability limit Δt_{STD}, Gerschgorin's (1931) theorem may be invoked. The eigenvalues λ_j^A of $\mathbf{A}$ are bounded by the relation

$$|\lambda_j^A - b_j| \leq |\tfrac{1}{2}(a_j + c_{j-1})| + |\tfrac{1}{2}(a_{j+1} + c_j)|, \qquad (A8)$$

for $j = \{1, \ldots, J-1\}$.

Under the usual assumption of $\sigma^2 S_\infty / r \gg 1$, we find $\lambda_{\max}^A \leq 2\sigma^2 S_\infty^2 / \Delta S^2$. Therefore, since $\mathbf{A}$ is symmetric positive definite we have $\|\mathbf{R}\|_2 = \sqrt{\rho(\mathbf{R}^2)} = 1 - \Delta t \lambda_{\max}^A$. The resultant stability constraint is again given by equation (26), i.e. $\Delta t \leq \Delta S^2 / \sigma^2 S_\infty^2$.

We now require a suitable prescription for $\mathbf{T}$ to complete the definition of the scheme. The simple choice of $\mathbf{T} = \mathbf{I} - \Delta t \mathbf{K}$ is unconditionally unstable since the eigenvalues of any skew symmetric real matrix are purely imaginary (or zero).

However, applying an update derived from a full-step predictor for a skew-symmetric operator is known to retrieve conditional stability for explicit schemes (O'Sullivan and Downes 2006, 2007), i.e. $\mathbf{TV}^{n+1} = \mathbf{V}^{n+1} - \Delta t \mathbf{K} \hat{\mathbf{V}}^{n+1}$ where $\hat{\mathbf{V}}^{n+1} = (\mathbf{I} - \Delta t \mathbf{K})\mathbf{V}^{n+1}$.

Hence, we adopt the following prescription for $\mathbf{T}$:

$$\mathbf{T} = \mathbf{I} - \Delta t \mathbf{K} + \Delta t^2 \mathbf{K}^2. \qquad (A9)$$

We proceed by establishing the stability properties of $\mathbf{T}$.

Since $\mathbf{K}$ is real and skew symmetric, we have $\mathbf{T}^{\dagger}\mathbf{T} = (\mathbf{I} + \Delta t \mathbf{K} + \Delta t^2 \mathbf{K}^2)(\mathbf{I} - \Delta t \mathbf{K} + \Delta t^2 \mathbf{K}^2) = \mathbf{I} + \Delta t^2 \mathbf{K}^2 + \Delta t^4 \mathbf{K}^4$. Furthermore, $\mathbf{K}$ is normal and may therefore be represented via the relation $\mathbf{K} = \mathbf{U}^{\dagger} \Lambda^K \mathbf{U}$ where $\mathbf{U} \in \mathbb{C}^{J-1} \times \mathbb{C}^{J-1}$ is unitary and $\Lambda^K \in \mathbb{C}^{J-1} \times \mathbb{C}^{J-1}$ is a diagonal matrix formed from the eigenvalues of $\mathbf{K}$, which occur in imaginary conjugate pairs $\lambda^K = \pm i \hat{\lambda}^K$ (plus 0 if the dimension is an odd integer). Now, $\mathbf{T}^{\dagger}\mathbf{T} = \mathbf{U}^{\dagger}(\mathbf{I} + \Delta t^2 (\Lambda^K)^2 + \Delta t^4 (\Lambda^K)^4)\mathbf{U}$ and therefore $\|\mathbf{T}\|_2 = [(1 - \Delta t^2 (\hat{\lambda}^K)^2 + \Delta t^4 (\hat{\lambda}^K)^4)_{\max}]^{1/2}$ since $\|\mathbf{U}\|_2 = 1$. We apply Gerschgorin's theorem to establish that $|\hat{\lambda}^K|_{\max} \leq |\sigma^2 - r| S_\infty / \Delta S$ and hence the condition for stability $\Delta t \leq \Delta S / |\sigma^2 - r| S_\infty$. This will in general be an insignificant constraint with respect to equation (26).

The conditionally stable split scheme is therefore

$$\mathbf{V}^n = (\mathbf{I} - \Delta t \mathbf{A}(\mathbf{I} - \Delta t \mathbf{K} + \Delta t^2 \mathbf{K}^2)\mathbf{V}^{n+1} - \Delta t \mathbf{b}^{n+1}. \quad (A10)$$

Since $\mathbf{A}$ is proven positive definite, acceleration via STS is formally admissible with the inner sub-steps, Δt_j, derived from the normal explicit time-step limit Δt_{STD} using equation (31). In particular, recall that $\Delta t_{\mathrm{STS}} \to N_{\mathrm{STS}}^2 \Delta t_{\mathrm{STD}}$ as $\nu \to 0$ where ν is a user-defined damping

parameter, and $\Delta t_{\text{STS}} = \sum_{j=1}^{N_{\text{STS}}} \Delta t_j$. The full accelerated scheme may therefore be written as

$$\mathbf{V}^n = \left(\prod_{j=1}^{N_{\text{STS}}} (\mathbf{I} - \Delta t_j \mathbf{A}) \right) (\mathbf{I} - \Delta t \mathbf{K} + \Delta t^2 \mathbf{K}^2) \mathbf{V}^{n+1} - \Delta t \mathbf{b}^{n+1},$$

(A11)

with the usual CFL condition $\Delta t \leq \Delta S^2 / \sigma^2 S_\infty^2$.

In summary, via MOS of the base scheme (A1), we have established strict stability of the consistent scheme (A11) under the usual time-step constraint (26). However, in this case, STS is formally admissible and yields acceleration in line with the results demonstrated in the main body of this work. In trials, we find that for the tests presented in this paper, the scheme (A11) yields negligible difference from the unsplit case (A1). This is in agreement with other authors who have noted in test cases that relaxing the symmetric positive definite requirement for the application of STS results only in a slight reduction of the stability region in the complex plane along the imaginary direction as long as the evolution operator is only *weakly* non-symmetric.