The Parallel Solution of Early-exercise Asian Options with Stochastic Volatility

A.K. Parrott [1], N.A.L. Clarke [2]

Introduction

This paper describes an parallel semi-Lagrangian finite difference approach to the pricing of early exercise Asian Options on assets with a stochastic volatility. A multigrid procedure is described for the fast iterative solution of the discrete linear complementarity problems that results. The accuracy and performance of this approach is improved considerably by a strike-price related analytic transformation of asset prices.

Asian options are contingent claims with payoffs that depend on the average price of an asset over some time interval. The payoff may depend on this average and a fixed strike price (Fixed Strike Asians) or it may depend on the average and the asset price (Floating Strike Asians). The option may also permit early exercise (American contract) or confine the holder to a fixed exercise date (European contract). The Fixed Strike Asian with early exercise is considered here where continuous arithmetic averaging has been used. Pricing such an option where the asset price has a stochastic volatility leads to the requirement to solve a tri-variate partial differential inequation in the three state variables of asset price, average price and volatility (or equivalently, variance). The similarity transformations [WDH93] used with Floating Strike Asian options to reduce the dimensionality of the problem are not applicable to Fixed Strikes

1 School of Computing & Mathematical Sciences, University of Greenwich, 30 Park Road, London SE10 9LS, U.K., email: pa10@gre.ac.uk
2 Oxford University Computing Laboratory, now with Goldman Sachs, London
Eleventh International Conference on Domain Decomposition Methods
Editors Choi-Hong Lai, Petter E. Bjørstad, Mark Cross and Olof B. Widlund

and so the numerical solution of a tri-variate problem is necessary. The computational challenge is to provide accurate solutions sufficiently quickly to support real-time trading activities at a reasonable cost in terms of hardware requirements.

American options with stochastic volatility

American Asian options with a stochastic volatility contain the vanilla American stochastic volatility option as a sub-problem; consequently the solution of this sub-problem (see [CP98] for more details) is summarised here. A standard American option with a stochastic volatility has a share price process S_t and its variance process Y_t (the variance has been used rather than the volatility $\sqrt{Y_t}$) which satisfy

$$\begin{aligned} dS_t &= \mu S_t dt + \sqrt{Y_t} S_t dB_t^1 && (1) \\ dY_t &= \alpha(\beta - Y_t)dt + \gamma\sqrt{Y_t} dB_t^2 && (2) \end{aligned}$$

where $B_t = (B_t^1, B_t^2)$ is a two-dimensional standard Brownian motion with correlation coefficient $\rho \in [-1, 1]$. The square-root mean-reverting stochastic volatility model described in [BR94] has been used where the reversion is to some mean value β at a rate determined by α, and $\gamma > 0$ governs the volatility of the variance process.

For an American option the following pricing in-equation must hold [WDH93] for the price u

$$\mathcal{D}u(s, y, t) \leq 0 \;\; \forall (s, y) \in \Omega, \quad t \in [0, T), \qquad (3)$$

$$u(s, y, T) = g(s, T) \qquad (4)$$

in the infinite quarter-plane $\Omega = \{(s, y) | s \geq 0, y \geq 0\}$, where s and y are the asset price and volatility variables respectively, g is the payoff and the pricing operator $\mathcal{D}$ (see [BR94]) is given by

$$\begin{aligned} \mathcal{D} = \frac{\partial}{\partial t} + \frac{1}{2}[s^2 y \frac{\partial^2}{\partial s^2} + 2\rho\gamma y s \frac{\partial^2}{\partial s \partial y} + \gamma^2 y \frac{\partial^2}{\partial y^2}] \\ + rs\frac{\partial}{\partial s} + [\alpha(\beta - y) - \lambda\gamma\sqrt{y}]\frac{\partial}{\partial y} - r, \end{aligned} \qquad (5)$$

where λ is the market price of risk and r is the risk-free interest rate. This leads to the following linear complementarity problem,

$$\left. \begin{aligned} u(s, y, t) &\geq g(s, t), \\ \mathcal{D}u(s, y, t) &\leq 0, \\ (u(s, y, t) - g(s, t))\mathcal{D}u(s, y, t) &= 0, \end{aligned} \right\} \; \forall (s, y, t) \in \Omega \times [0, T] \qquad (6)$$

with initial data $u(s, y, T) = g(s, T)$. The boundary conditions for this problem are:

$$
\begin{aligned}
u(0,y,t) &= g(0,t)\ \forall y \geq 0 \quad \text{and} \quad t \in [0,T]. && (7)\\
u(s,y,t) &\rightarrow 0 \text{ as } s \rightarrow \infty\ \forall y \geq 0 \text{ and } t \in [0,T], && (8)\\
\frac{\partial u(s,y,t)}{\partial y} &\rightarrow 0 \text{ as } y \rightarrow \infty\ \forall s \geq 0 \text{ and } t \in [0,T], && (9)
\end{aligned}
$$

A boundary condition is not required on $y = 0$ since the differential operator (5) becomes hyperbolic here. The domain is truncated at finite values $[s_{max}, y_{max}]$ for numerical solution, chosen sufficiently removed from the region of pricing interest for the solution to be unaffected.

The numerical procedure described in [CP98] uses an adaptive–upwind finite difference approximation together with a projected full approximation scheme (PFAS) Multigrid method [BC83]. Co-ordinate stretching transformations were also used for optimal efficiency,

$$
\begin{aligned}
\tilde{s} &= sinh^{-1}(s-k) - sinh^{-1}(-k);\ \forall s \in [0, s_{max}], && (10)\\
\tilde{y} &= ln(y') = ln(y/\beta)\ \forall y \in [y_0, y_{max}], && (11)
\end{aligned}
$$

where $y_0/\beta \ll 1$ and k is the strike price for the option. Both transformations induce rapid coarsening at large values of $\tilde{s}$ and $\tilde{y}$ enabling the truncation boundary to be placed far enough away for the asymptotic conditions (8), (9) to hold approximately. The co-ordinate stretching produced by both these transformations is equivalent to the use of a smooth and highly refined mesh about the strike.

A theta-method approximation of (5) on a uniform set of mesh points $(\tilde{s}_i, \tilde{y}_j) \in \Omega'_h \subseteq \Omega' = [0, s_{max}] \times [0, y_{max}]$, is then

$$
\mathcal{D}u(\tilde{s}_i, \tilde{y}_j) \simeq \mathcal{D}_h U^n_{ij} = (U^{n+1}_{ij} - U^n_{ij})/\Delta t^n + \theta \mathcal{L}_h U^{n+1}_{ij} + (1-\theta)\mathcal{L}_h U^n_{ij} \qquad (12)
$$

where $(0 \leq \theta \leq 1)$, U^n_{ij} is the discrete option price and $\mathcal{L}_h$ is an adaptive-upwind approximation to the drift-diffusion terms in (5). This leads to the following discrete linear complementarity problem, with initial data $U^N_{ij} = g(s_i, T)$, to be solved at a backwards sequence of times $t^n, n = N-1, N-2, ...0$,

$$
\left.
\begin{array}{rcc}
U^n_{ij} & \geq & g(s_i, t^n), \\
U^n_{ij} - (1-\theta)\Delta t^n \mathcal{L}_h U^n_{ij} & \geq & U^{n+1}_{ij} + \theta \Delta t^n \mathcal{L}_h U^{n+1}_{ij} \\
(U^n_{ij} - g(s_i, t^n))\mathcal{D}_h)U^n_{ij} & = & 0
\end{array}
\right\} \ \forall (i,j) \in \Omega'_h, \qquad (13)
$$

The value $\theta = 0.5 + \epsilon, \epsilon > 0$ was used since it gives unconditional stability (essential because of the coordinate stretching) and close to second order accuracy in the timestep $\Delta t^n = t^{n+1} - t^n$, provided $\epsilon \simeq O(h)$.

The conventional solution method for (13) is Projected Successive Over-Relaxation (PSOR) however the PFAS Multigrid method described in [CP98] is much superior as can be seen from in Table 1. Grid-independent convergence required the use of an adapted $\tilde{s}$-line smoother.

Table 1 Comparison of Multigrid with PSOR

	Avg. Iterations		CPU (s)	
Grid(Ω_h)	MG	PSOR	MG	PSOR
33 × 13	1.52	22.0	0.59	1.44
65 × 25	2.04	42.7	2.62	5.25
129 × 49	2.08	72.6	9.76	42.1
257 × 97	2.52	176.4	49.3	631.6

American Asian Options

The value of Asian options depend additionally on the average price; the continuously sampled average A_t is defined as

$$A_t = \frac{1}{t}\int_0^t S_\tau d\tau, \quad \text{i.e.} \quad dA_t = \frac{1}{t}(S_t - A_t)dt$$

where $A_0 = S_0$, leading to the following pricing operator

$$\begin{aligned} \mathcal{D}_A u(s,a,y,t) &= \frac{\partial u}{\partial t} + \frac{1}{2}[s^2 y \frac{\partial^2 u}{\partial s^2} + 2\rho\gamma y s \frac{\partial^2 u}{\partial s \partial y} + \gamma^2 y \frac{\partial^2 u}{\partial y^2}] \\ &\quad + rs\frac{\partial u}{\partial s} + [\alpha(\beta - y) - \lambda\gamma\sqrt{y}]\frac{\partial u}{\partial y} + \frac{1}{t}(s-a)\frac{\partial u}{\partial a} - ru \end{aligned} \qquad (14)$$

for the option price u on $\Omega = \{s, a, y : s \geq 0, a \geq 0, y \geq 0\} \times [0, T]$, with $u(s, a, y, T) = g(a, T)$ at $t = T$, for some payoff g and where a is the average-price variable. A fixed-strike Asian has a payoff of the form $g(s, a, t) = max(a - k, 0)$. Early exercise means that the option price satisfies the following LCP

$$u \geq g(a,t), \qquad (15)$$

$$\mathcal{D}_A u \leq 0 \qquad (16)$$

$$(u - g)\mathcal{D}_A u = 0 \qquad (17)$$

with appropriate boundary conditions [Cla98]. It can be seen that the pricing operator (14) is hyperbolic in the average price direction. Accurate approximation in this direction requires special care so a semi-Lagrangian approach was used. The Lagrangian derivative of u along any trajectory $\mathcal{T}(a, t)$ is

$$\frac{du}{dt} = \frac{\partial u}{\partial t} + \frac{\partial u}{\partial a}\frac{da}{dt}$$

and $\mathcal{D}_A$ can be simplified by choosing the trajectory to satisfy

$$\frac{da}{dt} = \frac{1}{t}(s - a) \qquad (18)$$

Consider a uniform 3-D set of meshpoints $(s_i, a_j, y_k) \in \Omega_h$. Let $\mathcal{T}(\bar{a}_j, t^n + \Delta t; a_j, t^n)$ be a trajectory satisfying (18) where (a_j, t^n) is the departure point and $(\bar{a}_j, t^n + \Delta t)$

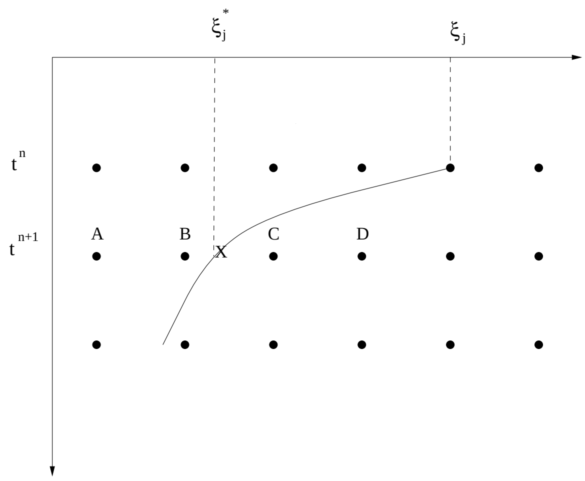

Figure 1 Lagrangian trajectory from (a_j, t^n) to $(\bar{a}_j, t^n + \Delta t)$.

is the arrival point of the trajectory (see Figure 1). Note that equation (18) can be solved exactly. Integrating (16) along this trajectory gives

$$u(s_i, \bar{a}_j, y_k, t^n + \Delta t) - u(s_i, a_j, y_k, t^n) \leq -\int_{t^n}^{t^n+\Delta t} \mathcal{L} u d\tau,$$

where $\mathcal{L}$ contains only asset-price/volatility drift diffusion terms. Approximating this integral with a weighted average of the end-point values leads to the discrete approximation

$$\tilde{U}_{ijk}^{n+1} - U_{ijk}^n \leq -\Delta t((1-\theta)\mathcal{L}_h U_{ijk}^n + \theta \mathcal{L}_h \tilde{U}_{ijk}^{n+1})$$

where $\mathcal{L}_h$ is the adaptive-upwind approximation described in the previous section of the asset price - volatility component of the pricing operator, and $\tilde{U}_{ijk}^{n+1}$ is obtained via cubic interpolation.

This semi-Lagrangian approach leads to the following discrete linear complementarity problems, with initial data $U_{ijk}^N = g(a_j, T)$, to be solved at a backwards sequence of times $t^n, n = N-1, N-2, ..., 0$,

$$\left.\begin{array}{ccc} U_{ijk}^n & \geq & g(a_j, t^n), \\ U_{ijk}^n - (1-\theta)\Delta t^n \mathcal{L}_h U_{ijk}^n & \geq & \tilde{U}_{ijk}^{n+1} + \theta \Delta t^n \mathcal{L}_h \tilde{U}_{ijk}^{n+1} \\ (U_{ijk}^n - g(a_j, t^n) \mathcal{D}_h U_{ijk}^n & = & 0 \end{array}\right\} \ \forall (i,j,k) \in \Omega_h,$$

The system of equations has a block-diagonal structure so that each $s - y$ plane can be solved independently using a modified form of the PFAS procedure described in Section 48. Co-ordinate stretching transformations are again essential for optimal efficiency. The results in Table 2 demonstrate the convergence of the method for the stochastic volatility American Asian and indicate that acceptable accuracy can be obtained on the $65 \times 65 \times 25$ mesh.

Table 2 Fixed Strike Asian Put Convergence results, $K = 10$

Asset Price (s)	Grid Size	Volatility ($\sqrt{y}$)				
		0.1	0.2	0.3	0.4	0.5
9.5	33 × 33 × 13	0.7733	0.7870	0.8104	0.8353	0.8840
	65 × 65 × 25	0.6372	0.6702	0.7300	0.8036	0.8869
	97 × 97 × 33	0.6216	0.6646	0.7269	0.8038	0.8881
10	33 × 33 × 13	0.2281	0.2762	0.3569	0.4477	0.5444
	65 × 65 × 25	0.2136	0.2762	0.3599	0.4535	0.5535
	97 × 97 × 33	0.2129	0.2756	0.3594	0.4538	0.5531
10.5	33 × 33 × 13	0.0638	0.1039	0.1686	0.2453	0.3284
	65 × 65 × 25	0.0592	0.0978	0.1604	0.2388	0.3287
	97 × 97 × 33	0.0583	0.0972	0.1602	0.2386	0.3290

Parallelization

The algebraic structure of the semi-Lagrangian approach makes it very suitable for parallelization. An s-line smoother is used in the PFAS multigrid method so it is advantageous on small-scale parallel systems to partition in the y-direction only, since this will not affect the smoother performance (a zebra ordering was used), non-local memory requests can be arranged to be for contiguous data, and this choice will keep trajectories within the partition. The surface to volume ratio for a 1-D partition degrades fairly quickly with the number of processors p and grid-coarsening so it was important to model the performance of this approach to establish its efficiency on different type of parallel systems.

The American Asian algorithm described above can be described in terms of the performance cost of a single iteration decomposed into its floating point cost, its remote memory access cost G, and some multiple of the processor network synchronization cost L . G and L are expressed in units of flop using a representative megaflop rate s. G can be written as gh where h is the maximum number of bytes read and written by any one processor during the iteration (termed an *h–relation*) so that g is the average cost in flop of accessing one remote word. This communication model is conservative, emphasizing global access; however low-level parallel library optimizations (which aggregate remote access between synchonisations to minimise latency and maximise bandwidth) make this a practical approach. The parallelisation relied on the BSP Library [Hil95] which uses such optimisations and for which estimates of L and g are available for many systems (see Table 3).

The iteration cost model presented assumes that the partitioning is load balanced and ignores the costs of restriction and prolongation operations within the multigrid iteration (they are small in comparison to the smoothing costs). These assumptions lead to the following

$$\text{iteration cost} = \sum_{i=1}^{N_{grids}} \left[4B_1 \left(\frac{Mr^{i-1}}{p} \right) + \left(32(p-1)\tilde{M}\tilde{r}^{i-1} \right) g + 16L \right].$$

Table 3 BSP Parameters

Machine	s (Mflops)	p	L	g
SGI Power-Challenge	55	2	627	7.6
		4	1248	7.4
Pentium (400MhZ)	40	2	20500	51
		4	29500	65
		8	39500	66

Table 4 Parallel efficiency (%) for the SGI Power Challenge

Grid Size		No. of Processors (p)		
		p=2	p=3	p=4
13 × 33 × 33	predicted	93	82	–
13 × 33 × 33	measured	92	83	–
25 × 65 × 65	predicted	97	90	83
25 × 65 × 65	measured	90	86	82

where r and $\tilde{r}$ are the grid/interface coarsening factors, M and $\tilde{M}$ are the total number of grid and interface mesh points and $B_1 \approx 300$ is the flop count. The cost of the semi-Lagrangian integration is an additional $B_2 M/p$ flop where $B_2 \approx 100$.

The parallel performance of the algorithm is shown in Table 4; the trend is comparable to those predicted by the performance model. The 1-D partitioning and grid coarsening limitations are visible in the trend but do not affect performance significantly on the 4-processor system. The predicted efficiency (not shown) for the $13 \times 33 \times 33$ is 77%, signalling the expected turndown due to 1-D partitioning. Work is in progress to repeat these results to other parallel systems. The fine mesh results can be obtained in 1.5 minutes (see Table 5) for a short-dated options, demonstrating that the semi-Lagrange algorithm combined with the PFAS multigrid approach has real-time performance on modestly parallel systems. Furthermore the SGI system performance can now be reproduced much more cheaply using Pentium based multiprocessor systems making this a widely affordable approach. The longer dated options are still rather slow to compute however the algorithm could be made more efficient with the use of adaptive time-stepping in these cases [CP98].

Conclusions

The semi-Lagrangian time integration has been shown to be well-suited to American Asian-options providing accurate integration of the hyperbolic terms in the pricing operator and making two-factor early-exercise solution methods applicable. Stochastic

Table 5 CPU times (in seconds) required for solution.

Time to Expiry	$33 \times 33 \times 13$			$65 \times 65 \times 25$			
	p=1	p=2	p=3	p=1	p=2	p=3	p=4
0.25	20	11	8	325	183	126	89
0.50	40	22	16	630	351	242	195
1.00	80	43	32	1264	695	498	382

volatility Asian options with early exercise can be solved in parallel by partitioning in the volatility direction and using a zebra-ordering; good performance was obtained on small parallel systems giving real-time computation of option prices for a reasonable cost. Finally the BSP performance model gave a useful indication of the portability of this parallel application.

REFERENCES

[BC83] Brandt C. and Cryer C. W. (1983) Multigrid algorithms for the solution of linear complementarity problems arising from free boundary problems. *SIAM Journal of Scientific and Statistical Computing* 4: 655–684.

[BR94] Ball C. A. and Roma A. (1994) Stochastic volatility option pricing. *Journal of Financial and Quantitative Analysis* 29: 589–607.

[Cla98] Clarke N. A. L. (1998) *The Numerical Solution of Financial Derivatives.* PhD thesis, University of Oxford.

[CP98] Clarke N. A. L. and Parrott A. K. (1998) Multigrid for American option pricing with stochastic volatility. *To be published in Applied Mathematical Finance* .

[Hil95] Hill J. M. D. (November 1995) The Oxford BSP toolset users guide. *Oxford University Computing Laboratory* .

[WDH93] Wilmott P., Dewynne J., and Howison S. (1993) *Options Pricing: Models and Computation.* Oxford Financial Press.