

A WEAK TRAPEZOIDAL METHOD FOR A CLASS OF STOCHASTIC DIFFERENTIAL EQUATIONS*

DAVID F. ANDERSON[†] AND JONATHAN C. MATTINGLY[‡]

Abstract. We present a numerical method for the approximation of solutions for the class of stochastic differential equations driven by Brownian motions which induce stochastic variation in fixed directions. This class of equations arises naturally in the study of population processes and chemical reaction kinetics. We show that the method constructs paths that are second order accurate in the weak sense. The method is simpler than many second order methods in that it neither requires the construction of iterated Itô integrals nor the evaluation of any derivatives. The method consists of two steps. In the first an explicit Euler step is used to take a fractional step. The resulting fractional point is then combined with the initial point to obtain a higher order, trapezoidal like, approximation. The higher order of accuracy stems from the fact that both the drift and the quadratic variation of the underlying SDE are approximated to second order.

Key words. Weak trapezoidal, error analysis, reaction networks, higher-order methods, stochastic differential equations, numerical methods.

AMS subject classifications. 60H35, 65C30, 65L20.

1. Introduction

We consider the problem of constructing accurate approximations on bounded time intervals to solutions of the following family of stochastic differential equations (SDEs)

$$\begin{aligned} dX(t) &= b(X(t))dt + \sum_{k=1}^M \sigma_k(X(t)) \nu_k dW_k(t), \\ X(0) &= x \in \mathbb{R}^d, \end{aligned} \tag{1.1}$$

where $b: \mathbb{R}^d \rightarrow \mathbb{R}^d$, $\sigma_k: \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$, $\nu_k \in \mathbb{R}^d$, and $W_k(t)$ are one-dimensional Wiener processes. Thus, randomness enters the system in fixed directions ν_k , but at variable rates $\sigma_k(X(t))$. Precise regularity conditions on the coefficients will be presented with our main results in Section 2.

The algorithm developed in this paper is a trapezoidal-type method and consists of two steps; in the first an explicit Euler step is used to take a fractional step and in the second the resulting fractional point is used in combination with the initial point to obtain a higher order, trapezoidal like, approximation. We will prove that the method developed is second order accurate in the weak sense. Because the method developed here produces single paths, it is natural to allow variable step-sizes; this is in contrast to Richardson extrapolation techniques ([20]). Finally, it is important to note that while the method presented in this paper is applicable to only a sub-class of SDEs, that sub-class does include systems whose diffusion terms do not commute, which is a classical simplifying assumption to obtain higher order methods (see [9, 14]).

*Received: June 17, 2009; accepted (in revised version): June 10, 2010. Communicated by Weinan E.

[†]Department of Mathematics, University of Wisconsin-Madison, Madison, Wi. 53706, USA (anderson@math.wisc.edu).

[‡]Department of Mathematics, Center of Nonlinear and Complex systems, Center for Theoretical and Mathematical Science, and Department of Statistical Science, Duke University, Durham, N.C. 27708, USA (jonm@math.duke.edu).

The method we propose is in some sense similar to the classical predictor-corrector. There have already been a number of such methods proposed in the stochastic context to produce higher-order methods (see [18, 19, 5]). In a general way, all of these methods require the simulation of iterated Itô integrals and sometimes need derivatives of the diffusion terms. If one only cares about weak accuracy, it is possible to use random variables which make these calculations easier and computationally cheaper. That being said, the complexity and cost of such calculations is one of the main impediments to their wider use. By assuming a certain structure for (1.1), we are able to develop a numerical method which we hope is more easily applied and implemented.

Though a specific structure of (1.1) is assumed, it is a structure which arises naturally in a number of settings. For example, our method will be applicable whenever $d = 1$. Also, we note that diffusion approximations to continuous time Markov chain models of population processes, including (bio)chemical processes, satisfy (1.1). As stochastic models of biochemical reaction systems, and, in particular, gene regulatory systems, are becoming more prevalent in the science literature, developing algorithms that utilize the specific structure of such models has increased importance [1, 2]. Furthermore, in Section 8 we quote a result from the literature which states that any system with uniformly elliptic diffusion can be put in the form of (1.1) without changing its distribution.

The topic of this paper is a method that produces a weak approximation rather than a strong approximation in that the approximate trajectory is produced without reference to an underlying Wiener process trajectory. We see this as an advantage. Except for applications such as filtering or certain problems of collective motion for stochastic flows, one is usually simply interested in generating an accurate draw from the distribution on $C([0, T], \mathbb{R}^d)$ induced by (1.1). This is different than accurately reproducing the Itô map $W \mapsto X(t, W)$ implied by (1.1). The latter is referred to as strong approximation. In our opinion such approximations are usually unnecessary and lead to a concept of accuracy which is unnecessarily restrictive. In [7], it is discussed that without accurately estimating second order Itô integrals one cannot produce a strong method of order greater than $1/2$. If the vector fields commute, then this restriction does not apply and higher order strong methods are possible. While the term “strong approximation” is quite specific, the term “weak approximation” is used for a number of concepts. Here we mean that the joint distribution of the numerical method at a fixed number of time points converges to the true marginal distribution as the numerical grid converges to zero. If this error goes to zero at a rate equal to the numerical mesh size raised to the power p in some norm on measures, then we say the method is of order p . This should be contrasted with talking about the rate at which a given function of the path converges.

The outline of the paper is as follows. In Section 2 we present our algorithm together with our main results concerning its weak error properties. In Section 3 we give the intuition as to why the method should work. In Section 4 we give the delayed proof of the local error estimates for the method which were stated in Section 2. In Section 5 we provide examples illustrating the performance of the proposed algorithm. In Section 6 we discuss the effect of varying the size of the first fractional step of the algorithm. In Section 7 we compare one step of the algorithm to one step in a Richardson extrapolation type algorithm. In Section 8 we show how, at least theoretically, the method can be applied to any uniformly elliptic SDE. Finally, an appendix contains a tedious calculation needed in Section 4.

2. The numerical method and main results

Throughout the paper, we let $X(t)$ denote the solution to (1.1) and Y_i denote the computed approximation at the time t_i for the time discretization $0 = t_0 < t_1 < \dots$. We begin both from the same initial condition, namely $X(0) = Y_0 = x_0$. Let $\{\eta_{1k}^{(i)}, \eta_{2k}^{(i)} : k \in \{1, \dots, M\}, i \in \mathbb{N}\}$ be a collection of mutually independent Gaussian random variables with mean zero and variance one. It is notationally convenient to define $[x]^+ = x \vee 0 = \max\{x, 0\}$.

We propose the following algorithm to approximate the solutions of (1.1).

ALGORITHM. (Weak θ -Midpoint Trapezoidal) *Fixing $\theta \in (0, 1)$, we define*

$$\alpha_1 \stackrel{\text{def}}{=} \frac{1}{2} \frac{1}{\theta(1-\theta)} \quad \text{and} \quad \alpha_2 \stackrel{\text{def}}{=} \frac{1}{2} \frac{(1-\theta)^2 + \theta^2}{\theta(1-\theta)}. \tag{2.1}$$

Next, fixing a discretization step h , for each $i \in \{1, 2, 3, \dots\}$ we repeat the following steps in which we first compute a θ -midpoint y^ and then the new value Y_i :*

Step 1.
$$y^* = Y_{i-1} + b(Y_{i-1})\theta h + \sum_{k=1}^M \sigma_k(Y_{i-1}) \nu_k \eta_{1k}^{(i)} \sqrt{\theta h}$$

Step 2.

$$Y_i = y^* + (\alpha_1 b(y^*) - \alpha_2 b(Y_{i-1}))(1-\theta)h + \sum_{k=1}^M \sqrt{[\alpha_1 \sigma_k^2(y^*) - \alpha_2 \sigma_k^2(Y_{i-1})]^+} \nu_k \eta_{2k}^{(i)} \sqrt{(1-\theta)h}.$$

REMARK 2.1. Notice that on the i th-step y^* is the standard Euler approximation to $X(\theta h + (i-1)h)$ starting from Y_{i-1} at time $(i-1)h$ [13].

REMARK 2.2. Notice that for all $\theta \in (0, 1)$ one has $\alpha_1 > \alpha_2$ and $\alpha_1 - \alpha_2 = 1$. It is reasonable to ask which θ is best. Notice that when $\theta = 1/2$ both α_1 and α_2 are minimized with values $\alpha_1 = 2$ and $\alpha_2 = 1$. This likely has positive stability implications. From the point of view of accuracy $\theta = 1/2$ also seems like a reasonable choice as it provides a central point for building a balanced trapezoidal approximation, as will be explained in Section 3. Further, picking a θ close to 1 or 0 increases the likelihood that the term $[\alpha_1 \sigma_k^2(y^*) - \alpha_2 \sigma_k^2(Y_{i-1})]^+$ will be zero, which will lower the accuracy of the method. If instability due to stiffness is a concern, one might consider a θ closer to one as that would likely give better stability properties, making the method closer to an implicit method. In general, $\theta = 1/2$ seems like a reasonable compromise, though this question requires further investigation and will be briefly revisited in Section 6.

For simplicity, we will restrict ourselves to the case when b and the σ_k are in $C^6(\mathbb{R}^d)$ — the space of bounded functions whose first through sixth derivatives are continuous and bounded. In general, we will denote by $C^k(\mathbb{R}^d)$ the space of bounded, continuous functions whose first k derivatives are bounded and continuous. For $f \in C^k(\mathbb{R}^d)$, we define the standard norm

$$\|f\|_k = \sup \{ |f(x)|, |\partial_\alpha f(x)| : x \in \mathbb{R}^d, \alpha = (\alpha_1, \dots, \alpha_j), \alpha_i \in \{1, \dots, d\}, j \leq k \}.$$

It is notationally convenient to define the Markov semigroup $\mathcal{P}_t : C^k \rightarrow C^k$ associated with (1.1) by

$$(\mathcal{P}_t f)(x) \stackrel{\text{def}}{=} \mathbb{E}_x f(X(t)), \tag{2.2}$$

where $X(0) = x$ and Markov semigroup $P_h: C^k \rightarrow C^k$ associated with a single full step of size h of the numerical method by

$$(P_h f)(y) \stackrel{\text{def}}{=} \mathbb{E}_y f(Y_1),$$

where $Y_0 = y$. Clearly $\|\mathcal{P}_h f\|_0 \leq \|f\|_0$ and $\|P_h f\|_0 \leq \|f\|_0$. It is also a standard fact, which we summarize in Appendix B, that in our setting, for any $t > 0$ and $k \in \mathbb{N}$, if $b, \sigma_1, \dots, \sigma_M \in C^k$ then there exists a $C = C(T, k, b, \sigma)$ so that $\|\mathcal{P}_t f\|_k \leq C\|f\|_k$ is true for all $t \leq T$. All of these can be rewritten succinctly in the induced operator norm from $C^k \rightarrow C^k$ as $\|\mathcal{P}_t\|_{k \rightarrow k} \leq C$, $\|\mathcal{P}_t\|_{0 \rightarrow 0} \leq 1$, and $\|P_h\|_{0 \rightarrow 0} \leq 1$. Analogously, for any linear operator $L: C^k \rightarrow C^\ell$ we will denote the induced operator norm from $C^k \rightarrow C^\ell$ by $\|L\|_{k \rightarrow \ell}$ which is defined by

$$\|L\|_{k \rightarrow \ell} = \sup_{f \in C^k, f \neq 0} \frac{\|Lf\|_\ell}{\|f\|_k}.$$

The following two theorems are the principle results of this article. They give, respectively, the weak local and global error of the Weak Trapezoidal method.

THEOREM 2.3 (One-step approximation). *Assume that $b \in C^6$ and for all k , $\sigma_k \in C^6$ with $\inf_x \sigma_k(x) > 0$. Then there exists a constant K so that*

$$\|\mathcal{P}_h - P_h\|_{6 \rightarrow 0} \leq Kh^3 \tag{2.3}$$

for all h sufficiently small.

From this one-step error bound, it is relatively straight-forward to obtain a global error bound. The following result shows that our approximation scheme gives a weak approximation of second order.

THEOREM 2.4 (Global approximation). *Assume that $b \in C^6$ and for all k , $\sigma_k \in C^6$ with $\inf_x \sigma_k(x) > 0$. Then for any $T > 0$ there exists a constant $C(T)$ such that*

$$\sup_{0 \leq n \leq T/h} \|\mathcal{P}_{nh} - P_h^n\|_{6 \rightarrow 0} \leq C(T)h^2. \tag{2.4}$$

Proof. We begin by observing that

$$P_h^n - \mathcal{P}_{nh} = \sum_{k=1}^n P_h^{k-1} (P_h - \mathcal{P}_h) \mathcal{P}_{h(n-k)}$$

and hence since $\sup_{0 \leq s \leq T} \|\mathcal{P}_s\|_{6 \rightarrow 6} \leq \tilde{C}(T)$ and $\|P_h^k\|_{0 \rightarrow 0} \leq 1$, using (2.3) we have that for any n with $0 \leq n \leq T/h$

$$\begin{aligned} \|\mathcal{P}_{nh} - P_h^n\|_{6 \rightarrow 0} &\leq \sum_{k=1}^n \|P_h^{k-1}\|_{0 \rightarrow 0} \|P_h - \mathcal{P}_h\|_{6 \rightarrow 0} \|\mathcal{P}_{h(n-k)}\|_{6 \rightarrow 6} \\ &\leq \sum_{k=1}^n \tilde{C}(T) Kh^3 = KT\tilde{C}(T)h^2 = C(T)h^2. \end{aligned}$$

□

REMARK 2.5. The restriction that $\inf_x \sigma_k(x) > 0$ can likely be relaxed if one has some control of the behavior of the solution around the degeneracies of $\sigma_k(x)$. This assumption is made to keep the proof simple with easily stated assumptions.

3. Why the method works

We now give two different, but related, explanations as to why the Weak θ -Midpoint Trapezoidal Algorithm is second order accurate in the weak sense.

3.1. A first point of view. Inserting the expression for y^* from Step 1 of the Weak θ -Midpoint Trapezoidal Algorithm into Step 2 and disregarding the diffusion terms yields

$$Y_i = Y_{i-1} + h \left[\frac{1}{2\theta} b(y^*) + \left(1 - \frac{1}{2\theta} \right) b(Y_{i-1}) \right] + \dots \tag{3.1}$$

$$= Y_{i-1} + b(Y_{i-1})h + \frac{b(y^*) - b(Y_{i-1})}{\theta h} \frac{h^2}{2} + \dots \tag{3.2}$$

Considering (3.1), we see that when $\theta \approx 1$ we recover the standard theta method (not to be confused with our use of θ) with theta = 1/2, which is known to be a second order method for deterministic systems. When $\theta = 1/2$, we recover the standard trapezoidal or midpoint method. For $\theta \neq 1/2$, we simply have a trapezoidal rule where a fractional point of the interval is used in the construction of the trapezoid. We will argue heuristically that the Weak Trapezoidal Algorithm handles the diffusion terms similarly. We also note that (3.2) shows that our algorithm can be understood as an approximation to the two-step Taylor series where θ is a parameter used to approximate the second derivative. This idea will be revisited in the proof of Theorem 2.3.

Equation (1.1) is distributionally equivalent to

$$X(t) = X(0) + \int_0^t b(X(s))ds + \sum_{k=1}^M \nu_k \int_0^\infty \int_0^t 1_{[0, \sigma_k^2(X(s))]}(u) Y_k(du \times ds), \tag{3.3}$$

where the Y_k are independent space-time white noise processes¹ and all other notation is as before, in that solutions to (3.3) are Markov processes that solve the same martingale problem as solutions to (1.1); that is, they have the same generator [6]. In order to approximate the diffusion term in (3.3) over the interval $[0, h]$, we must approximate $Y_k(A_{[0,h]}(\sigma_k^2))$, where $A_{[0,h]}(\sigma_k^2)$ is the region under the curve $\sigma_k^2(X(t))$, for $0 \leq t \leq h$.

We consider a natural way to approximate $X(h)$ and focus on the double integral in (3.3) for a single k . We also take $\theta = 1/2$ for simplicity and simply note that the case $\theta \neq 1/2$ follows similarly. We begin by approximating the value $X(h/2)$ by y^* obtained via an Euler approximation of the system on the interval $[0, h/2]$. To do so, we hold $X(t)$ fixed at $X(0)$ and see that we need to calculate $Y_k(\text{Region 1})$, where Region 1 is the grey shaded region in Figure 3.1(a). Because

$$Y_k(\text{Region 1}) \stackrel{\mathcal{D}}{=} N(0, \sigma_k^2(X(0))h/2) \stackrel{\mathcal{D}}{=} \sigma_k(X(0))\sqrt{\frac{h}{2}} N(0, 1),$$

we see that this step is equivalent in distribution to Step 1 of Algorithm 1.1. (Here and in the sequel, “ $\stackrel{\mathcal{D}}{=}$ ” denotes “equal in distribution.”)

¹More precisely, the Y_k are random measures on $[0, \infty)^2$ such that if $A, B \subset [0, \infty)^2$ with $A \cap B = \emptyset$ then $Y_k(A)$ and $Y_k(B)$ are each independent, mean zero Gaussian random variables with variances $\text{Area}(A)$ and $\text{Area}(B)$, respectively. Integration with respect to this field can be defined in the standard way beginning with adapted simple functions which are fixed random variables on fixed rectangular sets and then extending by linearity after the appropriate Itô isometry is established.




Fig. 3.1: A graphical depiction of the Weak Trapezoidal Algorithm with $\theta = 1/2$. In (a) the region of space-time used in the first step of the Weak Trapezoidal Algorithm is depicted by the grey shaded Region 1. In (b) the desired region to use, in order to perform a trapezoidal approximation, would be Region 2. However we have used Region 3 in our previous calculation and this is analytically problematic to undo. In (c), where $V = \sigma_k^2(y^*) - \sigma_k^2(X(0))$, we see that Region 5 gives the correct amount of new area wanted as subtracting off the area of Region 4 “offsets” the used area of Region 3. The case $\theta \neq 1/2$ is similar.

If we were trying to determine the area under the curve $\sigma_k^2(X(t))$ using an estimated midpoint y^* for a deterministic $X(t)$, one natural (and common) way would be to use the area of Region 2, where Region 2 is the grey shaded region in Figure 3.1(b). Such a method would be equivalent to the trapezoidal rule given in (3.1). However, in our setting we would have to ignore, or subtract off, the area already accounted for in Region 3, which is depicted as the shaded green section of Figure 3.1(b). In doing so, the random variable needed in order to perform this step would necessarily be dependent upon the past (via Region 3), and our current analysis would break down. However, noting that Region 3 has the same area as Region 4, as depicted by the blue shaded region in Figure 3.1(c), we see that it would be reasonable to expect that if one only uses Region 5, as depicted as the grey shaded region in Figure 3.1(c), then the accuracy of the method should be improved as we have performed a trapezoidal type approximation. Because

$$\begin{aligned} Y_k(\text{Region 5}) &\stackrel{\mathcal{D}}{=} N\left(0, (\sigma_k^2(X(0)) + 2V) \frac{h}{2}\right) \stackrel{\mathcal{D}}{=} \sqrt{\sigma_k^2(X(0)) + 2V} \sqrt{\frac{h}{2}} N(0, 1) \\ &= \sqrt{2\sigma_k^2(y^*) - \sigma_k^2(X(0))} \sqrt{\frac{h}{2}} N(0, 1), \end{aligned}$$

where $V = \sigma_k^2(y^*) - \sigma_k^2(X(0))$, we see that this is precisely what is carried out by Step 2 of the Weak θ -Midpoint Trapezoidal Algorithm.

3.2. A second point of view. To obtain a higher order method one must both approximate well the expected drift term as well as the quadratic variation of the process. The basic idea of the Weak Trapezoidal Algorithm is to make a preliminary step using an Euler approximation and then use this step to make a higher order approximation to the drift integral and to the quadratic variation integral. Similar to (3.1), the desired one step approximation to the quadratic variation integrals are

$$\int_0^h \sigma_k^2(X(s)) ds \approx h \left[\frac{1}{2\theta} \sigma_k^2(y^*) + \left(1 - \frac{1}{2\theta}\right) \sigma_k^2(Y_{i-1}) \right],$$

where all notation is as before.

Considering just the variance terms of the quadratic variation, we let $\{e_i\}$ be an orthonormal basis and see that our method yields the approximation

$$\begin{aligned} \text{Var}(X(h) \cdot e_i) &\approx \sum_{k=1}^M \text{Var}\left(\sigma_k(Y_0)(\nu_k \cdot e_i)\eta_{1k}\sqrt{\theta h}\right. \\ &\quad \left. + \sqrt{[\alpha_1\sigma_k^2(y^*) - \alpha_2\sigma_k^2(Y_0)]^+}(\nu_k \cdot e_i)\eta_{2k}\sqrt{(1-\theta)h}\right) \\ &= \sum_{k=1}^M \mathbb{E}\left(\sigma_k^2(Y_0)\theta + [\alpha_1\sigma_k^2(y^*) - \alpha_2\sigma_k^2(Y_0)]^+(1-\theta)\right)(\nu_k \cdot e_i)^2 h. \end{aligned}$$

If the step-size is sufficiently small, then $[\alpha_1\sigma_k^2(y^*) - \alpha_2\sigma_k^2(Y_0)]^+$ is positive with high probability because of our uniform ellipticity assumption; hence,

$$\text{Var}(X(h) \cdot e_i) \approx \mathbb{E} \sum_{k=1}^M (\nu_k \cdot e_i)^2 \left(\frac{1}{2\theta} \sigma_k^2(y^*) + \left(1 - \frac{1}{2\theta}\right) \sigma_k^2(Y_{i-1}) \right) h,$$

which is a locally third order approximation to the true quadratic variation integral of

$$\text{Var}(X(h) \cdot e_i) = \mathbb{E} \sum_{k=1}^M (\nu_k \cdot e_i)^2 \int_0^h \sigma_k^2(X(s)) ds.$$

Notice that it was important in this simple analysis that the direction of variation ν_k stayed constant over the interval so that the two terms could combine exactly. Of course, one should really be computing the full quadratic variation, including terms such as $\text{Cov}(X(h) \cdot e_i, X(h) \cdot e_j)$, but they follow the same pattern as above because each is a linear combination of the integral terms $\int_0^h \sigma_k^2(X(s)) ds$.

4. Proof of local error estimate

We now give the proof of the local error estimate given in Theorem 2.3 which is the central result of this paper.

Proof. (of Theorem 2.3) We need to show that there exists a constant K so that for any $f \in C^6$ one has

$$|\mathbb{E}f(Y_1) - \mathbb{E}f(X(h))| \leq K \|f\|_6 h^3.$$

Hence for the remainder of the proof we fix an arbitrary $f \in C^6$. Observe that Step 1 of the Weak Trapezoidal Algorithm produces a value, y^* , that is distributionally equivalent to $y(\theta h)$, where $y(t)$ solves

$$\begin{aligned} dy(t) &= b(y(0))dt + \sum_{k=1}^M \sigma_k(y(0)) \nu_k dW_k(t), \\ y(0) &= x_0. \end{aligned} \tag{4.1}$$

Likewise, Step 2 of the Weak Trapezoidal Algorithm produces a value, Y_1 , that is distributionally equivalent to $y(h)$, where $y(t)$ solves

$$\begin{aligned} dy(t) &= (\alpha_1 b(y^*) - \alpha_2 b(x_0))dt + \sum_{k=1}^M \sqrt{[\alpha_1\sigma_k^2(y^*) - \alpha_2\sigma_k^2(x_0)]^+} \nu_k dW_k(t), \\ y(\theta h) &= y^*. \end{aligned} \tag{4.2}$$

Let \mathcal{F}_t denote the filtration generated by the Wiener processes $W_k(t)$ in (4.1) and (4.2). Then

$$\mathbb{E}f(y(h)) = \mathbb{E}[\mathbb{E}[f(y(h)) | \mathcal{F}_{\theta h}]] \stackrel{\text{def}}{=} \mathbb{E}[\mathbb{E}_{\theta h}f(y(h))], \quad (4.3)$$

where we have made the definition $\mathbb{E}_{\theta h}[\cdot] \stackrel{\text{def}}{=} \mathbb{E}[\cdot | \mathcal{F}_{\theta h}]$.

Let A denote the generator for the process (1.1), B_1 denote the generator for the process (4.1), and B_2 denote the generator for the process (4.2) conditioned upon $\mathcal{F}_{\theta h}$. Then

$$(Af)(x) = f'[b](x) + \frac{1}{2} \sum_k \sigma_k^2 f''[\nu_k, \nu_k](x),$$

$$(B_1f)(x) = f'[b(x_0)](x) + \frac{1}{2} \sum_k \sigma_k(x_0)^2 f''[\nu_k, \nu_k](x), \quad \text{and}$$

$$(B_2f)(x) = f'[\alpha_1 b(y^*) - \alpha_2 b(x_0)](x) + \frac{1}{2} \sum_k [\alpha_1 \sigma_k(y^*)^2 - \alpha_2 \sigma_k(x_0)^2]^+ f''[\nu_k, \nu_k](x),$$

where $f'[\xi](z)$ is the derivative of f in the direction ξ evaluated at the point z . Note that $(Af)(x_0) = (B_1f)(x_0)$. For any integer $k \geq 2$ we define recursively $(A^k f)(x) \stackrel{\text{def}}{=} (A(A^{k-1}f))(x)$, and similarly for B_1 and B_2 . By repeated application of the Itô-Dynkin formula (see [17]), we have

$$\begin{aligned} \mathbb{E}_{\theta h}f(y(h)) &= f(y^*) + \int_{\theta h}^h \mathbb{E}_{\theta h}(B_2f)(y(s)) ds \\ &= f(y^*) + (B_2f)(y^*)(1-\theta)h + \int_{\theta h}^h \int_{\theta h}^s \mathbb{E}_{\theta h}(B_2^2f)(y(r)) dr ds \\ &= f(y^*) + (B_2f)(y^*)(1-\theta)h + (B_2^2f)(y^*) \frac{(1-\theta)^2 h^2}{2} \\ &\quad + \int_{\theta h}^h \int_{\theta h}^s \int_{\theta h}^r \mathbb{E}_{\theta h}(B_2^3f)(y(u)) du dr ds. \end{aligned} \quad (4.4)$$

The term $(B_2^3f)(y(u))$ depends on the first six derivatives of f . Therefore, since $f \in C^6$,

$$\left| \int_{\theta h}^h \int_{\theta h}^s \int_{\theta h}^r \mathbb{E}_{\theta h}(B_2^3f)(y(u)) du dr ds \right| \leq C \|f\|_6 h^3 \quad (4.5)$$

for some constant C . Combining (4.3), (4.4), (4.5), and recalling that $\mathbb{E}f(Y_1) = \mathbb{E}f(y(h))$, gives

$$\begin{aligned} \mathbb{E}f(Y_1) &= \mathbb{E}[\mathbb{E}_{\theta h}f(y(h))] \\ &= \mathbb{E}f(y^*) + \mathbb{E}(B_2f)(y^*)(1-\theta)h \\ &\quad + \mathbb{E}(B_2^2f)(y^*) \frac{(1-\theta)^2 h^2}{2} + O(h^3). \end{aligned} \quad (4.6)$$

Here and in the sequel, we will write $F = G + O(h^p)$ to mean that there exist a constant K depending on only σ and b so that for all initial conditions x_0 ,

$$|F - G| \leq K \|f\|_6 h^p \quad (4.7)$$

for h sufficiently small. In the spirit of the preceding calculation, repeated application of the Itô-Dynkin formula to (1.1) produces

$$\mathbb{E}f(X(h)) = f(x_0) + (Af)(x_0)h + (A^2f)(x_0)\frac{h^2}{2} + O(h^3).$$

The proof of the theorem is then completed by Lemma 4.1 given below. Its proof, which is straightforward but tedious, is given in the appendix. \square

LEMMA 4.1. *Under the assumptions of Theorem 2.3, for all $h > 0$ sufficiently small and $f \in C^6$ one has*

$$\begin{aligned} & \mathbb{E} \left[f(y^*) + (B_2f)(y^*)(1 - \theta)h + (B_2^2f)(y^*)\frac{(1 - \theta)^2h^2}{2} \right] \\ &= f(x_0) + (Af)(x_0) + (A^2f)(x_0)\frac{h^2}{2} + O(h^3). \end{aligned}$$

REMARK 4.2. Comparing Equation (3.2) and Lemma 4.1 shows that our algorithm can be viewed as providing an approximation to the two step Taylor series approximation.

5. Examples

We present two examples that demonstrate the rate of convergence of the Weak Trapezoidal Algorithm with $\theta = 1/2$. In each example we shall compare the accuracy of the proposed algorithm to that of Euler’s method and a “midpoint drift” algorithm defined via repetition of the following steps

$$\begin{aligned} y^* &= Y_{i-1} + b(Y_{i-1})\frac{h}{2}, \\ Y_i &= Y_{i-1} + b(y^*)h + \sum_{k=1}^M \sigma_k(Y_{i-1})\nu_k \eta_k \sqrt{h}, \end{aligned} \tag{5.1}$$

where the notation is as before. We compare the proposed algorithm to that given via (5.1) to point out that the gain in efficiency being demonstrated is not solely due to the fact that we are getting better approximations to the drift terms, but also because of the superior approximation of the diffusion terms.

5.1. First example. Consider the system

$$\begin{bmatrix} dX_1(t) \\ dX_2(t) \end{bmatrix} = \begin{bmatrix} X_1(t) \\ 0 \end{bmatrix} + X_1(t) \begin{bmatrix} 0 \\ 1 \end{bmatrix} dW_1(t) + \frac{1}{10} \begin{bmatrix} 1 \\ 1 \end{bmatrix} dW_2(t), \tag{5.2}$$

where $W_1(t)$ and $W_2(t)$ are standard Weiner processes. In our notation $b_1(x) = x_1$, $b_2(x) = 0$, $\sigma_1(x) = x_1$, $\sigma_2(x) = 1/10$, and $\nu_1 = [0, 1]^T$, $\nu_2 = [1, 1]^T$. Note that the noise does not commute. It is an exercise to show that

$$\mathbb{E}X_2(t)^2 = \mathbb{E} X_2(0)^2 - \frac{1}{2}\mathbb{E} X_1(0)^2 + \frac{1}{400}e^{2t}(200\mathbb{E}X_1(0)^2 + 1) + \frac{t}{200} - \frac{1}{400}. \tag{5.3}$$

For both Euler’s method and the midpoint drift method (5.1) we used step sizes $h_k = 1/3^k$, $k \in \{1, 2, 3, 4, 5\}$ and initial condition $X_1(0) = X_2(0) = 1$ to generate 500,000 sample paths of the system (5.2). We then computed

$$\text{error}_k(t) = \mathbb{E}X_2(t)^2 - \frac{1}{5 \times 10^5} \sum_{i=1}^{5 \times 10^5} \bar{X}_2^{h_k}(t)^2, \tag{5.4}$$

where $\bar{X}^{h_k}(t)$ is the sample path generated numerically and $\mathbb{E}X_2(t)^2$ is given via (5.3). We also generated 10,000,000 sample paths using the Weak Trapezoidal Algorithm with the same initial condition and step sizes $h_k = 1/(4k)$, $k \in \{1, 2, 3, 4\}$. We then computed $\text{error}_k(t)$ similarly to before. The outcome of the numerical experiment is summarized in Figure 5.1a where we have plotted $\log(h_k)$ versus $\log(|\text{error}_k(1)|)$ for the different algorithms. As expected, we see that the Weak Trapezoidal Algorithm gives an error that decreases proportional to h^2 , whereas the other two algorithms give errors that decrease proportional to h .




Fig. 5.1: Log-log plots of the step-size versus the error for the three different algorithms. In (a) the example (5.2) is considered. The best fit lines for the data (shown) have slopes 2.029, .998, and 1.030, for the Weak Trapezoidal Algorithm, Euler's method, and the midpoint drift method, respectively. In (b) the example in (5.5) is considered. The best fit lines for the data (shown) have slopes 2.223, .952, and 1.098, for the Weak Trapezoidal Algorithm, Euler's method, and the midpoint drift method, respectively. In both examples all results agree with what was expected.

5.2. Second example. Now consider the following system that is similar to one considered in [20]:

$$\begin{aligned} \begin{bmatrix} dX_1(t) \\ dX_2(t) \end{bmatrix} &= \begin{bmatrix} -X_2(t) \\ X_1(t) \end{bmatrix} + \sqrt{\frac{\sin^2(X_1(t) + X_2(t)) + 6}{t+1}} \begin{bmatrix} 1 \\ 0 \end{bmatrix} dW_1(t) \\ &+ \sqrt{\frac{\cos^2(X_1(t) + X_2(t)) + 6}{t+1}} \begin{bmatrix} 0 \\ 1 \end{bmatrix} dW_2(t), \end{aligned} \quad (5.5)$$

where $W_1(t)$ and $W_2(t)$ are independent Wiener processes. It is simple to show that

$$\mathbb{E}|X(t)|^2 = \mathbb{E}X(0)^2 + 13 \log(1+t). \quad (5.6)$$

We used step sizes $h_k = 1/(2k)$, $k \in \{1, 2, \dots, 8\}$, to generate five million approximate sample paths of the system (5.5) using each of: (a) Weak Trapezoidal Algorithm, (b) Euler's method, and (c) the midpoint drift method (5.1). We then computed

$$\text{error}_k(t) = \mathbb{E}|X(t)|^2 - \frac{1}{5 \times 10^6} \sum_{i=1}^{5 \times 10^6} |\bar{X}^{h_k}(t)|^2,$$

where $\bar{X}^{h_k}(t)$ is the sample path generated numerically and $\mathbb{E}|X(t)|^2$ is given via (5.6). The outcome is summarized in Figure 5.1b where we have plotted $\log(h_k)$ versus $\log(|\text{error}_k(1)|)$ for the different algorithms. As before, we see that the Weak Trapezoidal Algorithm gives an error that decreases proportional to h^2 , whereas the other two algorithms give errors that decrease proportional to h .

REMARK 5.1. We note that in both examples we needed to average over an extremely large number of computed sample paths in order to estimate $\text{error}_k(t)$ for the Weak Trapezoidal Algorithm. This is due to the fact that the increased accuracy of the method quickly makes sampling error the dominant error.

6. The effect of varying θ

The term $[\alpha_1\sigma_k^2(y^*) - \alpha_2\sigma_k^2(Y_{i-1})]^+$ in Step 2 of the Weak Trapezoidal Algorithm will yield zero, and the given step will have a local error of only $O(h^2)$, if

$$\alpha_1\sigma_k^2(y^*) < \alpha_2\sigma_k^2(Y_{i-1}) \iff \sigma_k^2(y^*) < \frac{\alpha_2}{\alpha_1}\sigma_k^2(Y_{i-1}) = (1 - 2\theta + 2\theta^2)\sigma_k^2(Y_{i-1}).$$

We will call such a step a “degenerate” step. The function $g(\theta) = 1 - 2\theta + 2\theta^2$ is minimized at $g(1/2) = 1/2$, and $g(\theta) \rightarrow 1$ as $\theta \rightarrow 0$ or $\theta \rightarrow 1$. Therefore, as mentioned Remark 2.2, one would expect that as $\theta \rightarrow 0$ or $\theta \rightarrow 1$ more steps will be degenerate, and a decrease in accuracy, together with a bias against σ_k decreasing, would follow. Using a step-size of $h = 1/10$, we tracked the percentage of degenerate steps for the simple system

$$dX(t) = \sqrt{X(t)^2 + 1} dW(t), \quad X(0) = 1, \tag{6.1}$$

where $W(t)$ is a standard Weiner process. To do so, we computed 10,000 sample paths over the time interval $[0, 1]$ for each of $\theta = .02k$, $k \in \{1, \dots, 49\}$. The results are shown in Figure 6.1a where the behavior predicted above is seen. Curiously, the minimum number of rejections takes place at $\theta = .42$. It is also worth noting that




Fig. 6.1: (a) The number of degenerate steps for the Weak Trapezoidal Algorithm applied to (6.1) with $h = 1/10$ and different values of θ . (b) The $\log h$ vs $\log(|\text{error}|)$ plot is given for different choices of θ for the Weak Trapezoidal Algorithm applied to (5.2) where the error is defined similarly to (5.4). The best fit lines for the data (shown) have slopes 1.865, 1.996, 2.029, and 2.033 for $\theta = .05, .25, .50, .75$, respectively.

one can check on computer software that in the general case the coefficient of h^3 for the one-step error grows like $1/\theta$ as $\theta \rightarrow 0$. This does not happen in the deterministic case (3.1).

While the above considerations give some interesting insight into the effect of various θ , the situation is more complex. A θ closer to one should give the method more stability, albeit at an expense as the rejection fraction increases as θ approaches one. It would be interesting to perform a stability analysis in the spirit of [8] to better understand the effect of θ . In lieu of this, Figure 6.1b gives the result of a convergence analysis of the Weak Trapezoidal Algorithm applied to (5.2) with different choices of θ . Interestingly, larger θ seem to result in smaller (and hence better) convergence rate prefactors. This seems to indicate that in at least this example stability is an issue.

The performance of the Weak Trapezoidal Algorithm as a function of θ is a topic deserving further consideration, but combining the above shows that $\theta = 1/2$ is a reasonable first choice, though stability considerations might lead one to consider a θ closer to 1.

7. Comparison to richardson extrapolation

It is illustrative to compare the Weak Trapezoidal Algorithm to Richardson extrapolation, which from a certain point of view is the method in the literature that is most similar to ours. See [20] for complete details of Richardson extrapolation in the SDE setting.

Let $Z_{h/2}(t)$ and $Z_h(t)$ denote approximate sample paths of (1.1) generated using Euler's method with step sizes of $h/2$ and h , respectively. For all f satisfying mild assumptions, both $\mathbb{E}f(Z_{h/2}(t))$ and $\mathbb{E}f(Z_h(t))$ will approximate $\mathbb{E}f(X(t))$ with an order of $O(h)$. However, Richardson extrapolation may be used and the linear combination $2\mathbb{E}f(Z_{h/2}(t)) - \mathbb{E}f(Z_h(t))$ will approximate $\mathbb{E}f(X(t))$ with an order of $O(h^2)$ (see [20]). Of course, taking f to be the identity shows that the linear combination $2Z_{h/2}(t) - Z_h(t)$ gives an $O(h^2)$ approximate of the mean of the process. As Richardson extrapolation does not compute a single path, but instead uses the statistics from two to achieve a higher order of approximation for a given statistic, we will compare one step of the Weak Trapezoidal Algorithm with a step-size of h , to one step of size h of the process $2Z_{h/2}(t) - Z_h(t)$ with the clear understanding that $2Z_{h/2}(t) - Z_h(t)$ is only $O(h)$ accurate for higher moments.

Recall that systems of the form (1.1) are equivalent to those driven by space-time white noise processes (3.3). As in Section 3.1, we consider how each method uses the areas of $[0, \infty)^2$ associated to $Y_k(du \times ds)$ from (3.3) during one step. We will proceed considering a single k since it is sufficient to illustrate the point. For $A_i \subset [0, \infty)^2$, we denote by η_{A_i} a normal random variable with mean 0 and variance $\text{area}(A_i)$. Recall that η_{A_i} and η_{A_j} are independent as long as $A_i \cap A_j$ has Lebesgue measure zero. Consider (7.1)(a) in which we are supposing that $\sigma_k^2(X(t))$ increases over a single time-step. The change in the process $Z_{h/2}$ due to this k would be ν_k times

$$\eta_{A_1} + \eta_{A_2} + \eta_{A_3}.$$

Similarly, the change in Z_h would be ν_k times $\eta_{A_1} + \eta_{A_2}$. Therefore, the change in the process $2Z_{h/2}(t) - Z_h(t)$ would be ν_k times

$$\eta_{A_1} + \eta_{A_2} + 2\eta_{A_3}.$$

On the other hand, the change in the process generated by the Weak Trapezoidal Algorithm due to this k is ν_k times

$$\eta_{A_1} + \eta_{A_2} + \eta_{A_3} + \eta_{A_4}.$$




Fig. 7.1: The areas of space-time utilized by $2Z_{h/2} - Z_h$ and the Weak Trapezoidal Algorithm for a single k and a single step. In 7.1(a), $\sigma_k^2(X(t))$ increases and $2Z_{h/2} - Z_h$ uses $\eta_{A_1} + \eta_{A_2} + 2\eta_{A_3}$, whereas the Weak Trapezoidal Algorithm uses $\eta_{A_1} + \eta_{A_2} + \eta_{A_3} + \eta_{A_4}$. In the case when $\sigma_k^2(X(t))$ decreases, 7.1(b) above, the processes use $\eta_{A_1} + \eta_{A_2} + \eta_{A_3} - \eta_{A_4}$ and $\eta_{A_1} + \eta_{A_2}$, respectively. In both cases, it is the better use of the areas by the Weak Trapezoidal Algorithm that achieves a higher order of convergence.

Therefore, and as expected, the means should be the same, but the variances should not as

$$\text{Var}(2\eta_{A_3}) = 4\text{Var}(\eta_{A_3}) = 2\text{Var}(\eta_{A_3} + \eta_{A_4}).$$

Similarly, in the case in which $\sigma_k^2(X(t))$ decreases as depicted in (7.1)(b), the process $2Z_{h/2}(t) - Z_h(t)$ would use $\eta_{A_1} + \eta_{A_2} + \eta_{A_3} - \eta_{A_4}$, whereas the Weak Trapezoidal Algorithm would use $\eta_{A_1} + \eta_{A_2}$. Again, the means will be the same, but the variances will not. In both cases, the Weak Trapezoidal Algorithm makes better use of the areas to approximate the quadratic variation of the true process, and thus achieves a higher order of convergence.

8. Extension to general uniformly elliptic systems

For a moment let us consider the setting of general uniformly elliptic SDEs

$$\begin{aligned} dX(t) &= b(X(t))dt + \sum_{k=1}^M g_k(X(t)) dW_k(t), \\ X(0) &= x \in \mathbb{R}^d, \end{aligned} \tag{8.1}$$

where b and W are as before and $g_k: \mathbb{R}^d \rightarrow \mathbb{R}^d$ is such that if $G(x) = (g_1(x), \dots, g_M(x))(g_1(x), \dots, g_M(x))^T$ then there exist positive λ_- and λ_+ such that

$$\lambda_- |\xi|^2 \leq G(x)\xi \cdot \xi \leq \lambda_+ |\xi|^2$$

for all $x, \xi \in \mathbb{R}^d$. For such a family of uniformly elliptic matrices a lemma of Motzkin and Wasow [15], whose precise formulation we take from Kurtz [10], states that if the

entries of G are C^k then there exists an M and $\{\sigma_k: \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0} : k = 1, \dots, M\}$, $\{\nu_k \in \mathbb{R}^d : k = 1, \dots, M\}$ with $\sigma_k \in C^k$ and strictly positive so that

$$G(x) = \sum \sigma_k^2(x) \nu_k \nu_k^T.$$

Hence (8.1) has the same law on path space as (1.1) with these σ_k and ν_k . Of course M might be arbitrarily large (depending on the ratio of λ_+/λ_-) and hence it is more subtle to compare the total work for our method with a standard scheme based directly on (8.1). Furthermore, depending on the dependence on x , it is not transparent how to obtain the vectors ν and functions σ exactly. Approximations could be obtained using the SVN of the matrix $G(x)$ for fixed x but we do not explore this further here.

9. Conclusions and further extensions

We have presented a relatively simple method directly applicable to a wide class of systems which is weakly second order. We have also shown how, at least theoretically, it should be applicable to systems which do not satisfy our structural assumptions but are uniformly elliptic. We have picked a particularly simple setting to perform our analysis to make the central points clearer. The assumption that b and σ_k are uniformly bounded can be relaxed to a local Lipschitz condition. That is to say, if b and σ and their needed derivatives are not bounded uniformly, but rather are bounded by an appropriate Lyapunov function, then it should be possible to extend the method directly to the setting of unbounded coefficients provided the method is stable for the given SDE (see for instance [12]). If the SDE is not globally Lipschitz then using an implicit drift split-step method as in [12], an adaptive method as in [11], or a truncation method as in [14] should extend to our current setting. More interesting is relaxing the non-degeneracy assumption on the σ_k , which was used to minimize the probability of the diffusion correction being negative. This tact is in some ways reminiscent of [14] in that a modification of the method is made on a small set of paths, though the take here is quite different. It would be interesting to use the probability that the correction to the diffusion is negative to adapt the step-size much in the spirit of [11]. Lastly, there is some similarity of our method with predictor corrector methods. In the deterministic setting, predictor corrector methods not only have a higher order of accuracy but also have better stability properties. There have been a number of papers exploring this in the stochastic context (see [5, 4, 19, 8]). It would be interesting to do the same with the method presented here.

Appendix A. Proof of Lemma 4.1. The proof of Lemma 4.1 requires the replacement of the terms of the form $[\alpha_1 \sigma_k^2(y^*) - \alpha_2 \sigma_k^2(x_0)]^+$ with $[\alpha_1 \sigma_k^2(y^*) - \alpha_2 \sigma_k^2(x_0)]$. The following two lemmas show that this can be done at the cost of an error whose size is $O(h^3)$. Here $O(h^3)$ has the same meaning described earlier around (4.7). We begin with an abstract technical lemma where p and q satisfy $1/p + 1/q = 1$.

PROPOSITION A.1. *Let X and Y be a real valued random variables on a probability space (Ω, \mathbb{P}) with $|XY|_{L^p(\Omega)} < \infty$ for some $p \in (1, \infty]$. Then $|\mathbb{E}Y[X]^+ - \mathbb{E}YX| \leq |YX|_{L^p(\Omega)}(\mathbb{P}\{X < 0\})^{1/q}$. Similarly if X, Y and Z are real valued random variables with $|ZXY|_{L^p(\Omega)} < \infty$ and $A = \{X < 0\} \cup \{Z < 0\}$ then $|\mathbb{E}Y[X]^+[Z]^+ - \mathbb{E}YXZ| \leq 2|ZYX|_{L^p(\Omega)}(\mathbb{P}\{A\})^{1/q}$.*

Proof. Let $A = \{X < 0\}$ and $q = p/(p-1)$. Then $|\mathbb{E}Y([X]^+ - X)| \leq \mathbb{E}|Y| |[X]^+ - X| \mathbf{1}_A \leq |YX|_{L^p(\Omega)}(\mathbb{P}(A))^{1/q}$, showing the first claim. For the second notice that $\mathbb{E}Y[X]^+[Z]^+ - \mathbb{E}YXZ = (\mathbb{E}Y[X]^+Z - \mathbb{E}YXZ) + (\mathbb{E}Y[X]^+[Z]^+ - \mathbb{E}Y[X]^+Z)$ and that each of the terms in parentheses can be bounded by the first result. \square

COROLLARY A.2. Let $\sigma_k \in C^2$ with $\inf_x \sigma_k(x) > 0$ for all k and let Y be a random variable with $|Y| \leq C$ a.s. for some C . Then for any $p \geq 1$ there exists an h_0 so that

$$\begin{aligned} \mathbb{E}Y[\alpha_1\sigma_k^2(y^*) - \alpha_2\sigma_k^2(x_0)]^+ &= \mathbb{E}Y[\alpha_1\sigma_k^2(y^*) - \alpha_2\sigma_k^2(x_0)] + O(h^p) \quad \text{and} \\ \mathbb{E}Y[\alpha_1\sigma_k^2(y^*) - \alpha_2\sigma_k^2(x_0)]^+[\alpha_1\sigma_\ell^2(y^*) - \alpha_2\sigma_\ell^2(x_0)]^+ \\ &= \mathbb{E}Y[\alpha_1\sigma_k^2(y^*) - \alpha_2\sigma_k^2(x_0)][\alpha_1\sigma_\ell^2(y^*) - \alpha_2\sigma_\ell^2(x_0)] + O(h^p) \end{aligned}$$

for all $h \in (0, h_0]$ and $k, \ell \in \{1, \dots, M\}$, where y^* is defined via Step 1 of the Weak Trapezoidal Algorithm.

Proof. Define the event $A_k = \left\{ \sigma_k(y^*) < \frac{\alpha_2}{\alpha_1} \sigma_k(x_0) \right\}$. In light of Proposition A.1, it is sufficient to show that for any $p > 1$ there exists a C_p such that $\mathbb{P}(A_k) \leq C_p h^p$. Because σ_k is Lipschitz there exists a positive C such that

$$\sigma_k^2(x_0 + \delta) - \frac{\alpha_2}{\alpha_1} \sigma_k^2(x_0) > \left(1 - \frac{\alpha_2}{\alpha_1}\right) \sigma_k^2(x_0) - C|\delta|,$$

for any $\delta > 0$. In particular, setting $\delta = y^* - x_0 = b(x_0)\theta h + \sum_j \sigma_j(x_0)\sqrt{\theta h} \nu_j \eta_{1j}^{(1)}$, and noting that $\alpha_2 < \alpha_1$ and that the σ 's are uniformly bounded from both above and below, the result follows from the Gaussian tails of the η 's. \square

Proof. (of Lemma 4.1) From Taylor's theorem and the definition of the operators involved one has

$$\begin{aligned} \mathbb{E}f(y^*) &= f(x_0) + (B_1f)(x_0)\theta h + (B_1^2f)(x_0)\frac{\theta^2 h^2}{2} + O(h^3) \\ &= f(x_0) + (Af)(x_0)\theta h + (B_1^2f)(x_0)\frac{\theta^2 h^2}{2} + O(h^3). \end{aligned}$$

In the last line, we have used the observation that $(B_1f)(x_0) = (Af)(x_0)$. Now we turn to $\mathbb{E}(B_2f)(y^*)$. We begin by using Lemma A.2 to remove the $[\cdot]^+$. Then we use the fact that $\alpha_1 - \alpha_2 = 1$ and Taylor's theorem to expand various terms to produce the following:

$$\begin{aligned} &\mathbb{E}(B_2f)(y^*) \\ &= \mathbb{E}f'(y^*)[\alpha_1 b(y^*) - \alpha_2 b(x_0)] + \frac{1}{2} \mathbb{E} \sum_k [\alpha_1 \sigma_k^2(y^*) - \alpha_2 \sigma_k^2(x_0)]^+ f''[\nu_k, \nu_k](y^*) \\ &= \mathbb{E}f'(y^*)[\alpha_1 b(y^*) - \alpha_2 b(x_0)] + \frac{1}{2} \mathbb{E} \sum_k [\alpha_1 \sigma_k^2(y^*) - \alpha_2 \sigma_k^2(x_0)] f''[\nu_k, \nu_k](y^*) + O(h^2) \\ &= f'(x_0)[b(x_0)] + \frac{1}{2} \sum_k \sigma_k(x_0)^2 f''(x_0)[\nu_k, \nu_k] \\ &\quad + \mathbb{E}B_1\left(f'[\alpha_1 b - \alpha_2 b(x_0)] + \frac{1}{2} \sum_k (\alpha_1 \sigma_k^2 - \alpha_2 \sigma_k^2(x_0)) f''[\nu_k, \nu_k]\right)(x_0)\theta h + O(h^2) \\ &= (Af)(x_0) + \alpha_1 (B_1(Af))(x_0)\theta h - \alpha_2 (B_1^2f)(x_0)\theta h + O(h^2) \\ &= (Af)(x_0) + \alpha_1 (A^2f)(x_0)\theta h - \alpha_2 (B_1^2f)(x_0)\theta h + O(h^2). \end{aligned}$$

Similar reasoning produces

$$\begin{aligned}
& \mathbb{E}(B_2^2 f)(y^*) \\
&= \mathbb{E}\left(B_2(f'[\alpha_1 b(y^*) - \alpha_2 b(x_0)] + \frac{1}{2} \sum_k [\alpha_1 \sigma_k^2(y^*) - \alpha_2 \sigma_k^2(x_0)]^+ f''[\nu_k, \nu_k])(y^*)\right) \\
&= f''[b(x_0), b(x_0)](x_0) + \mathbb{E} \sum_k [\alpha_1 \sigma_k^2(y^*) - \alpha_2 \sigma_k^2(x_0)]^+ f'''[\nu_k, \nu_k, b(x_0)](x_0) \\
&+ \frac{1}{4} \mathbb{E} \sum_{k,j} [\alpha_1 \sigma_k^2(y^*) - \alpha_2 \sigma_k^2(x_0)]^+ [\alpha_1 \sigma_j^2(y^*) - \alpha_2 \sigma_j^2(x_0)]^+ f''''[\nu_k, \nu_k, \nu_j, \nu_j](x_0) + O(h) \\
&= f''[b(x_0), b(x_0)](x_0) + \sum_k \sigma_k^2(x_0) f'''[\nu_k, \nu_k, b(x_0)](x_0) \\
&\quad + \frac{1}{4} \sum_{k,j} \sigma_k^2(x_0) \sigma_j^2(x_0) f''''[\nu_k, \nu_k, \nu_j, \nu_j](x_0) + O(h) \\
&= (B_1^2 f)(x_0) + O(h).
\end{aligned}$$

Combining these estimates and the fact that $2(1-\theta)\theta\alpha_2 = \theta^2 + (1-\theta)^2$ and $2(1-\theta)\theta\alpha_1 = 1$ produces the quoted result after some algebra. \square

Appendix B. Operator bound for $\mathcal{P}_t: C^k \rightarrow C^k$. In this section, we show that if $b, \sigma_\ell \in C^k$ then \mathcal{P}_t is a bounded operator from C^m to C^m for $m \in \{0, \dots, k\}$. The $k=0$ case follows immediately from $|f(x)| \leq \|f\|_0$ for all $x \in \mathbb{R}^d$. To address the higher k , we introduce the first k variations of Equation (1.1).

For any $\xi \in \mathbb{R}^d$ we denote the first variation of (1.1) in the direction ξ by $J^{(1)}(t, x)[\xi]$ which solves the linear equation

$$\begin{aligned}
dJ^{(1)}(t, x)[\xi] &= (\nabla b)(X(t))[J^{(1)}(t, x)[\xi]] dt + \sum_{k=1}^M \nu_k (\nabla \sigma_k)(X(t))[J^{(1)}(t, x)[\xi]] dW_k(t), \\
J^{(1)}(0, x)[\xi] &= \xi \quad \text{and} \quad X(0) = x.
\end{aligned}$$

Similarly for $\xi = (\xi_1, \xi_2) \in \mathbb{R}^2$, the second variation of $X(t)$ (in the directions ξ) will be denoted by $J^{(2)}(t, x)[\xi]$ and defined by

$$\begin{aligned}
dJ^{(2)}(t, x)[\xi] &= (\nabla b)(X(t))[J^{(2)}(t, x)[\xi]] dt + \sum_{k=1}^M \nu_k (\nabla \sigma_k)(X(t))[J^{(2)}(t, x)[\xi]] dW_k(t) \\
&\quad + (\nabla^2 b)(X(t))[J^{(1)}(t, x)[\xi_1], J^{(1)}(t, x)[\xi_2]] \\
&\quad + \sum_{k=1}^M (\nabla^2 \sigma_k)(X(t))[J^{(1)}(t, x)[\xi_1], J^{(1)}(t, x)[\xi_2]] dW_k(t) \\
J^{(2)}(0, x)[\xi] &= 0 \quad \text{and} \quad X(0) = x.
\end{aligned}$$

These equations were obtained from successive formal differentiation of (1.1). By further formal differentiation we obtain analogous equations for the k -variation $J^{(k)}(t, x)[\xi]$ where $\xi = (\xi_1, \dots, \xi_k) \in \mathbb{R}^k$ is the vector of directions. It is a standard fact that if the coefficients b, σ_j are in C^k then for any $t > 0$

$$\sup_x \mathbb{E}_x \sup \left\{ \sup_{s \in [0, t]} |J^{(n)}(s, x)[\xi_1, \dots, \xi_n]|^p : \xi_i \in \mathbb{R}^d \text{ with } |\xi_i| = 1 \right\} < \infty.$$

This can be found in Lemma 2 in [3] on p. 196 or in a slightly different context in Proposition 1.3 in [16]² With these definitions in hand, we have that for any $f \in C^1$ that

$$\begin{aligned}\nabla(\mathcal{P}_t f)(x)[\xi] &= \mathbb{E}_x f'(X(t))[J^{(1)}(t, x)[\xi]], \\ \nabla^2(\mathcal{P}_t f)(x)[\xi] &= \mathbb{E}_x f''(X(t))[J^{(2)}(t, x)[\xi]] + \mathbb{E}_x f^{(2)}(X(t))[J^{(1)}(t, x)[\xi_1], J^{(1)}(t, x)[\xi_2]].\end{aligned}$$

Using the moment bounds we have that for $q \geq 1$ and an ever changing constant C ,

$$\begin{aligned}\mathbb{E} \sup_{|\xi|=1} |\nabla(\mathcal{P}_t f)(x)[\xi]|^q &\leq C \|f\|_{C^1}^q \sup_{|\xi|=1} |J_t^{(1)}[\xi]|^q \leq C \|f\|_{C^1}^q < \infty \\ \mathbb{E} \sup_{|\xi_i|=1} |\nabla^2(\mathcal{P}_t f)(x)[\xi_1, \xi_2]|^q &\leq C \|f\|_{C^2}^q \left(\mathbb{E} \sup_{|\xi_1|=1} |J^{(1)}(t, x)[\xi_1]|^{2q} \right)^{\frac{1}{2}} + \mathbb{E} \sup_{|\xi_i|=1} |J^{(2)}(t, x)[\xi_1, \xi_2]|^q \\ &\leq C \|f\|_{C^2}^q < \infty.\end{aligned}$$

Continuing in this manner we see that for any positive integer m , if $f, b, \sigma_\ell \in C^m$ then for any $q \geq 1$ one has

$$\mathbb{E} \sup_{|\xi_i|=1} |\nabla^m(\mathcal{P}_t f)(x)[\xi_1, \dots, \xi_m]|^q \leq C \|f\|_{C^m}^q < \infty$$

for some C . Now observe that taking $q = 1$ proves the desired claim on the operator norm of \mathcal{P}_t from C^k to C^k since

$$\|\mathcal{P}_t f\|_k \leq C \sum_{j=0}^k \mathbb{E} \sup_{|\xi_i|=1} |(\nabla^j \mathcal{P}_t f)(x)[\xi_1, \dots, \xi_j]| \leq C \sum_{j=0}^k \|f\|_{C^j} \leq C \|f\|_{C^k}.$$

Acknowledgments. DFA was supported through grant NSF-DMS-0553687 and JCM through grants NSF-DMS-0449910 and NSF-DMS-06-16710 and a Sloan Foundation Fellowship. We would like to thank Andrew Stuart for useful comments on an early draft and Martin Clark for stimulating questions about Richardson Extrapolation. We also thank Thomas Kurtz for pointing out that all uniformly elliptic SDEs can be represented in the form considered in this paper.

REFERENCES

- [1] D.F. Anderson, *Incorporating postleap checks in tau-leaping*, J. Chem. Phys., 128, 054103, 2008.
- [2] D.F. Anderson, A. Ganguly and T.G. Kurtz, *Error analysis of the tau-leap simulation method for stochastically modeled chemical reaction systems*, submitted.
- [3] D.R. Bell, *The Malliavin Calculus*, Pitman Monographs and Surveys in Pure and Applied Mathematics, Longman Scientific & Technical, Harlow, 34, 1987.
- [4] N. Bruti-Liberati and E. Platen, *Strong predictor-corrector Euler methods for stochastic differential equations*, Stoch. Dyn., 8, 561–581, 2008.
- [5] K. Burrage and T. Tian, *Predictor-corrector methods of Runge-Kutta type for stochastic differential equations*, SIAM J. Numer. Anal., 40, 1516–1537, (electronic) 2002.
- [6] S.N. Ethier and T.G. Kurtz, *Markov Processes: Characterization and Convergence*, John Wiley & Sons, New York, 1986.

²The statement of the Proposition demands coefficients in C^∞ . However the bounds on $J^{(k)}$ only require C^k coefficients.

- [7] J.G. Gaines and T.J. Lyons, *Variable step size control in the numerical solution of stochastic differential equations*, SIAM J. Appl. Math., 57, 1455–1484, 1997.
- [8] D.J. Higham, *Mean-square and asymptotic stability of the stochastic theta method*, SIAM J. on Numerical Analysis, 38, 753–769, 2000.
- [9] P.E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*, Applications of Mathematics (New York), Springer-Verlag, Berlin, 23, 1992.
- [10] T.G. Kurtz, *Representations of markov processes as multparameter time changes*, Ann. Prob., 8, 682–715, 1980.
- [11] H. Lamba, J.C. Mattingly and A.M. Stuart, *An adaptive Euler-Maruyama scheme for SDEs: convergence and stability*, IMA J. Numer. Anal., 27, 479–506, 2007.
- [12] J.C. Mattingly, A.M. Stuart and D.J. Higham, *Ergodicity for SDEs and approximations: locally Lipschitz vector fields and degenerate noise*, Stochastic Process. Appl., 101, 185–232, 2002.
- [13] G.N. Milstein, *Numerical Integration of Stochastic Differential Equations*, Kluwer Academic Press, Dordrecht, The Netherlands, 1995.
- [14] G.N. Milstein and M.V. Tretyakov, *Numerical integration of stochastic differential equations with nonglobally Lipschitz coefficients*, SIAM J. Numer. Anal., 43, 1139–1154, (electronic) 2005.
- [15] T.S. Motzkin and W. Wasow, *On the approximation of linear elliptic differential equations by difference equations with positive coefficients*, J. Math. Physics, 31, 253–259, 1953.
- [16] J. Norris, *Simplified Malliavin calculus*, in Séminaire de Probabilités, XX, 1984/85, Lecture Notes in Math., Springer, Berlin, 1204, 101–130, 1986.
- [17] B. Øksendal, *Stochastic Differential Equations: an Introduction with Applications*, Springer, Berlin, sixth ed., 2003.
- [18] É. Pardoux and D. Talay, *Discretization and simulation of stochastic differential equations*, Acta Appl. Math., 3, 23–47, 1985.
- [19] E. Platen, *On weak implicit and predictor-corrector methods*, Math. Comput. Simulation, 38, 69–76, 1995. Probabilités numériques Paris, 1992.
- [20] D. Talay and L. Tubaro, *Expansion of the global error for numerical schemes solving stochastic differential equations*, Stochastic Analysis and Applications, 8, 483–509, 1990.