

Multilevel and Multi-index Monte Carlo methods for the McKean–Vlasov equation

Abdul-Lateef Haji-Ali¹  · Raúl Tempone²

Received: 31 October 2016 / Accepted: 31 July 2017 / Published online: 12 September 2017
© The Author(s) 2017. This article is an open access publication

Abstract We address the approximation of functionals depending on a system of particles, described by stochastic differential equations (SDEs), in the mean-field limit when the number of particles approaches infinity. This problem is equivalent to estimating the weak solution of the limiting McKean–Vlasov SDE. To that end, our approach uses systems with finite numbers of particles and a time-stepping scheme. In this case, there are two discretization parameters: the number of time steps and the number of particles. Based on these two parameters, we consider different variants of the Monte Carlo and Multilevel Monte Carlo (MLMC) methods and show that, in the best case, the optimal work complexity of MLMC, to estimate the functional in one typical setting with an error tolerance of TOL , is $\mathcal{O}(TOL^{-3})$ when using the partitioning estimator and the Milstein time-stepping scheme. We also consider a method that uses the recent Multi-index Monte Carlo method and show an improved work complexity in the same typical setting of $\mathcal{O}(TOL^{-2} \log(TOL^{-1})^2)$. Our numerical experiments are carried out on the so-called Kuramoto model, a system of coupled oscillators.

Keywords Multi-index Monte Carlo · Multilevel Monte Carlo · Monte Carlo · Particle systems · McKean–Vlasov · Mean-field · Stochastic differential equations · Weak

approximation · Sparse approximation · Combination technique

Mathematics Subject Classification 65C05 (Monte Carlo methods) · 65C30 (Stochastic differential and integral equations) · 65C35 (Stochastic particle methods)

1 Introduction

In our setting, a stochastic particle system is a system of *coupled* d -dimensional stochastic differential equations (SDEs), each modeling the state of a “particle.” Such particle systems are versatile tools that can be used to model the dynamics of various complicated phenomena using relatively simple interactions, e.g., pedestrian dynamics (Helbing and Molnar 1995; Haji-Ali 2012), collective animal behavior (Erban et al. 2016; Erban and Haskovec 2012), interactions between cells (Dobramysl et al. 2016) and in some numerical methods such as ensemble Kalman filters (Pierre Del Moral and Tugaut 2016). One common goal of the simulation of these particle systems is to average some quantity of interest computed on *all* particles, e.g., the average velocity, average exit time or average number of particles in a specific region.

Under certain conditions, most importantly the exchangeability of particles and sufficient regularity of the SDE coefficients, the stochastic particle system approaches a mean-field limit as the number of particles tends to infinity (Sznitman 1991). Exchangeability of particles refers to the assumption that all permutations of the particles have the same joint distribution. In the mean-field limit, each particle follows a single McKean–Vlasov SDE where the advection and/or diffusion coefficients depend on the distribution of the solution to the SDE (Gärtner 1988). In many cases, the objective is to approximate the expected value of a quantity

✉ Abdul-Lateef Haji-Ali
abdullateef.hajiali@maths.ox.ac.uk
Raúl Tempone
raul.tempone@kaust.edu.sa

¹ Mathematical Institute, University of Oxford, Oxford, UK

² Computer, Electrical and Mathematical Sciences and Engineering, King Abdullah University of Science and Technology, Thuwal, Saudi Arabia

of interest (QoI) in the mean-field limit as the number of particles tends to infinity, subject to some error tolerance, TOL. While it is possible to approximate the expectation of these QoIs by estimating the solution to a nonlinear PDE using traditional numerical methods, such methods usually suffer from the curse of dimensionality. Indeed, the cost of these methods is usually of $\mathcal{O}\left(\text{TOL}^{-\frac{w}{d}}\right)$ for some constant $w > 1$ that depends on the particular numerical method. Using sparse numerical methods alleviates the curse of dimensionality but requires increasing regularity as the dimensionality of the state space increases. On the other hand, Monte Carlo methods do not suffer from this curse with respect to the dimensionality of the state space. This work explores different variants and extensions of the Monte Carlo method when the underlying stochastic particle system satisfies certain crucial assumptions. We theoretically show the validity of some of these assumptions in a somewhat general setting, while verifying the other assumptions numerically on a simple stochastic particle system, leaving further theoretical justification to a future work.

Generally, the SDEs that constitute a stochastic particle system cannot be solved exactly, and their solution must instead be approximated using a time-stepping scheme with a number of time steps, N . This approximation parameter and a finite number of particles, P , are the two approximation parameters that are involved in approximating a finite average of the QoI computed for all particles in the system. Then, to approximate the expectation of this average, we use a Monte Carlo method. In such a method, multiple *independent* and *identical* stochastic particle systems, approximated with the same number of time steps, N , are simulated, and the average QoI is computed from each and an overall average is then taken. Using this method, a reduction in the variance of the estimator is achieved by increasing the number of simulations of the stochastic particle system or increasing the number of particles in the system. Section 3.1 presents the Monte Carlo method more precisely in the setting of stochastic particle systems. Particle methods that are not based on Monte Carlo were also discussed in Bossy and Talay (1996, 1997). In these methods, a *single* simulation of the stochastic particle system is carried out, and only the number of particles is increased to reduce the variance.

As an improvement of Monte Carlo methods, the Multilevel Monte Carlo (MLMC) method was first introduced in Heinrich (2001) for parametric integration and in Giles (2008b) for SDEs; see Giles (2015) and references therein for an overview. MLMC improves the efficiency of the Monte Carlo method when only an approximation, controlled with a single discretization parameter, of the solution to the underlying system can be computed. The basic idea is to reduce the number of required samples on the finest, most accurate but most expensive discretization, by reducing the variability

of this approximation with a *correlated* coarser and cheaper discretization as a control variate. More details are given in Sect. 3.2 for the case of stochastic particle systems. The application of MLMC to particle systems has been investigated in many works (Bujok et al. 2013; Haji-Ali 2012; Rosin et al. 2014). The same concepts have also been applied to nested expectations (Giles 2015). More recently, a particle method applying the MLMC methodology to stochastic particle systems was also introduced in Ricketson (2015) achieving, for a linear system with a diffusion coefficient that is independent of the state variable, a work complexity of $\mathcal{O}\left(\text{TOL}^{-2}(\log(\text{TOL}^{-1}))^5\right)$.

Recently, the Multi-index Monte Carlo (MIMC) method (Haji-Ali et al. 2015a) was introduced to tackle high-dimensional problems with more than one discretization parameter. MIMC is based on the same concepts as MLMC and improves the efficiency of MLMC even further but requires mixed regularity with respect to the discretization parameters. More details are given in Sect. 3.3 for the case of stochastic particle systems. In that section, we demonstrate the improved work complexity of MIMC compared with the work complexity of MC and MLMC, when applied to a stochastic particle system. More specifically, we show that, when using a naive simulation method for the particle system with quadratic complexity, the optimal work complexity of MIMC is $\mathcal{O}\left(\text{TOL}^{-2} \log(\text{TOL}^{-1})^2\right)$ when using the Milstein time-stepping scheme and $\mathcal{O}\left(\text{TOL}^{-2} \log(\text{TOL}^{-1})^4\right)$ when using the Euler–Maruyama time-stepping scheme. Finally, in Sect. 4, we provide numerical verification for the assumptions that are made throughout the current work and the derived rates of the work complexity.

In what follows, the notation $a \lesssim b$ means that there exists a constant c that is independent of b such that $a < cb$.

2 Problem setting

Consider a system of P exchangeable stochastic differential equations (SDEs) where for $p = 1 \dots P$, we have the following equation for $X_{p|P}(t) \in \mathbb{R}^d$

$$\begin{cases} dX_{p|P}(t) = \mathcal{A}(t, X_{p|P}(t), \lambda_{X_P(t)}) dt \\ \quad + \mathcal{B}(t, X_{p|P}(t), \lambda_{X_P(t)}) dW_p(t) \\ X_{p|P}(0) = x_p^0 \end{cases} \quad (1)$$

where $X_P(t) = \{X_{q|P}(t)\}_{q=1}^P$ and for some (possibly stochastic) functions, $\mathcal{A} : [0, \infty) \times \mathbb{R}^d \times \mathbb{P}(\mathbb{R}^d) \rightarrow \mathbb{R}^d$ and $\mathcal{B} : [0, \infty) \times \mathbb{R}^d \times \mathbb{P}(\mathbb{R}^d) \rightarrow \mathbb{R}^d \times \mathbb{R}^d$ and $\mathbb{P}(\mathbb{R}^d)$ is the space of probability measures over \mathbb{R}^d . Moreover,

$$\lambda_{X_P(t)} \stackrel{\text{def}}{=} \frac{1}{P} \sum_{q=1}^P \delta_{X_{q|P}(t)} \in \mathbb{P}(\mathbb{R}^d),$$

where δ is the Dirac measure, is called the empirical measure. In this setting, $\{W_p\}_{p \geq 1}$ are mutually independent d -dimensional Wiener processes. If, moreover, $\{x_p^0\}_{p \geq 1}$ are i.i.d., then under certain conditions on the smoothness and form of \mathcal{A} and \mathcal{B} (Sznitman 1991), as $P \rightarrow \infty$ for any $p \in \mathbb{N}$, the $X_{p|\infty}$ stochastic process satisfies

$$\begin{cases} dX_{p|\infty}(t) = \mathcal{A}(t, X_{p|\infty}(t), \mu_\infty^t)dt \\ \quad + \mathcal{B}(t, X_{p|\infty}(t), \mu_\infty^t)dW_p(t) \\ X_{p|\infty}(0) = x_p^0, \end{cases} \quad (2)$$

where $\mu_\infty^t \in \mathbb{P}(\mathbb{R}^d)$ is the corresponding mean-field measure. Under some smoothness and boundedness conditions on \mathcal{A} and \mathcal{B} , the measure μ_∞^t induces a probability density function (pdf), $\rho_\infty(t, \cdot)$, that is the Radon–Nikodym derivative with respect to the Lebesgue measure. Moreover, ρ_∞ satisfies the McKean–Vlasov equation

$$\begin{aligned} \frac{\partial \rho_\infty(t, x)}{\partial t} + \sum_{p=1}^P \frac{\partial}{\partial x_p} (\mathcal{A}(t, x_p, \mu_\infty^t) \rho_\infty(t, x)) \\ - \frac{1}{2} \sum_{p=1}^P \frac{\partial^2}{\partial x_p^2} (\mathcal{B}(t, x_p, \mu_\infty^t)^2 \rho_\infty(t, x)) = 0 \end{aligned}$$

on $t \in [0, \infty)$ and $x = (x_p)_{p=1}^P \in \mathbb{R}^d$ with $\rho_\infty(0, \cdot)$ being the pdf of x_p^0 which is given and is independent of p . Due to (2) and x_p^0 being i.i.d., $\{X_{p|\infty}\}_p$ are also i.i.d.; hence, unless we want to emphasize the particular path, we drop the p -dependence in $X_{p|\infty}$ and refer to the random process X_∞ instead. In any case, we are interested in computing $E[\psi(X_\infty(T))]$ for some given function, ψ , and some final time, $T < \infty$.

Kuramoto Example (Fully connected Kuramoto model for synchronized oscillators) *Throughout this work, we focus on a simple, one-dimensional example of (1). For $p = 1, 2, \dots, P$, we seek $X_{p|P}(t) \in \mathbb{R}$ that satisfies*

$$\begin{aligned} dX_{p|P}(t) = \left(\vartheta_p + \frac{1}{P} \sum_{q=1}^P \sin(X_{p|P}(t) - X_{q|P}(t)) \right) dt \\ + \sigma dW_p(t) \\ X_{p|P}(0) = x_p^0, \end{aligned} \quad (3)$$

where $\sigma \in \mathbb{R}$ is a constant and $\{\vartheta_p\}_p$ are i.i.d. and independent from the set of i.i.d. random variables $\{x_p^0\}_p$ and the Wiener processes $\{W_p\}_p$. The limiting SDE as $P \rightarrow \infty$ is

$$\begin{aligned} dX_{p|\infty}(t) = \left(\vartheta_p + \int_{-\infty}^{\infty} \sin(X_{p|\infty}(t) - y) d\mu_\infty^t(y) \right) dt \\ + \sigma dW_p(t) \\ X_{p|\infty}(0) = x_p^0. \end{aligned}$$

Note that in terms of the generic system (1), we have

$$\mathcal{A}(t, x, \mu) = \vartheta + \int_{-\infty}^{\infty} \sin(x - y) d\mu(y)$$

with ϑ a random variable and $\mathcal{B} = \sigma$ is a constant. We are interested in

$$\begin{aligned} \text{Total synchronization} = (E[\cos(X_\infty(T))])^2 \\ + (E[\sin(X_\infty(T))])^2, \end{aligned}$$

a real number between zero and one that measures the level of synchronization in the system with an infinite number of oscillators (Acebrón et al. 2005); with zero corresponding to total disorder. In this case, we need two estimators: one where we take $\psi(\cdot) = \sin(\cdot)$ and the other where we take $\psi(\cdot) = \cos(\cdot)$.

While it is computationally efficient to approximate $E[\psi(X_\infty(T))]$ by solving the McKean–Vlasov PDE, that ρ_∞ satisfies, when the state dimensionality, d , is small (cf., e.g., Haji-Ali 2012), the cost of a standard full tensor approximation increases exponentially as the dimensionality of the state space increases. On the other hand, using sparse approximation techniques to solve the PDE requires increasing regularity assumptions as the dimensionality of the state space increases. Instead, in this work, we focus on approximating the value of $E[\psi(X_\infty)]$ by simulating the SDE system in (1). Let us now define

$$\phi_P \stackrel{\text{def}}{=} \frac{1}{P} \sum_{p=1}^P \psi(X_{p|P}(T)). \quad (4)$$

Here, due to exchangeability, $\{X_{p|P}(T)\}_{p=1}^P$ are identically distributed, but they are not independent since they are taken from the same realization of the particle system. Nevertheless, we have $E[\phi_P] = E[\psi(X_{p|P}(T))]$ for any p and P . In this case, with respect to the number of particles, P , the cost of a naive calculation of ϕ_P is $\mathcal{O}(P^2)$ due to the cost of evaluating the empirical measure in (1) for every particle in the system. It is possible to take $\{X_{p|P}\}_{p=1}^P$ in (4) as i.i.d., i.e., for each $p = 1 \dots P$, $X_{p|P}$ is taken from a different independent realization of the system (1). In this case, the usual law of large numbers applies, but the cost of a naive calculation of ϕ_P is $\mathcal{O}(P^3)$. For this reason, we focus in this work on the former method of taking identically distributed but not independent $\{X_{p|P}\}_{p=1}^P$.

Following the setup in Collier et al. (2015), Haji-Ali et al. (2015b), our objective is to build a random estimator, \mathcal{A} , approximating $\phi_\infty \stackrel{\text{def}}{=} E[\psi(X_\infty(T))]$ with minimal work, i.e., we wish to satisfy the constraint

$$P[|\mathcal{A} - \phi_\infty| \geq \text{TOL}] \leq \epsilon \tag{5}$$

for a given error tolerance, TOL, and a given confidence level determined by $0 < \epsilon \ll 1$. We instead impose the following, more restrictive, two constraints:

Bias constraint: $|E[\mathcal{A} - \phi_\infty]| \leq (1 - \theta)\text{TOL}$, (6)

Statistical constraint: $P[|\mathcal{A} - E[\mathcal{A}]| \geq \theta\text{TOL}] \leq \epsilon$, (7)

for a given tolerance splitting parameter, $\theta \in (0, 1)$, possibly a function of TOL. To show that these bounds are sufficient note that

$$P[|\mathcal{A} - \phi_\infty| \geq \text{TOL}] \leq P[|\mathcal{A} - E[\mathcal{A}]| + |E[\mathcal{A}] - \phi_\infty| \geq \text{TOL}]$$

imposing (6) yields

$$P[|\mathcal{A} - \phi_\infty| \geq \text{TOL}] \leq P[|\mathcal{A} - E[\mathcal{A}]| \geq \theta\text{TOL}]$$

then imposing (7) gives (5). Next, we can use Markov inequality and impose $\text{Var}[\mathcal{A}] \leq \epsilon(\theta\text{TOL})^2$ to satisfy (7). However, by assuming (at least asymptotic) normality of the estimator, \mathcal{A} , we can get a less stringent condition on the variance as follows:

Variance constraint: $\text{Var}[\mathcal{A}] \leq \left(\frac{\theta\text{TOL}}{C_\epsilon}\right)^2$. (8)

Here, $0 < C_\epsilon$ is such that $\Phi(C_\epsilon) = 1 - \frac{\epsilon}{2}$, where Φ is the cumulative distribution function of a standard normal random variable, e.g., $C_\epsilon \approx 1.96$ for $\epsilon = 0.05$. The asymptotic normality of the estimator is usually shown using some form of the central limit theorem (CLT) or the Lindeberg–Feller theorem (see, e.g., Collier et al. 2015; Haji-Ali et al. 2015a for CLT results for the MLMC and MIMC estimators and Fig. 3-right).

As previously mentioned, we wish to approximate the values of X_∞ by using (1) with a finite number of particles, P . For a given number of particles, P , a solution to (1) is not readily available. Instead, we have to discretize the system of SDEs using, for example, the Euler–Maruyama time-stepping scheme with N time steps. For $n = 0, 1, 2, \dots, N - 1$,

$$\begin{aligned} X_{p|P}^{n+1|N} - X_{p|P}^{n|N} &= \mathcal{A}\left(X_{p|P}^{n|N}, \lambda_{X_{p|P}^{n|N}}\right) \frac{T}{N} \\ &\quad + \mathcal{B}\left(X_{p|P}^{n|N}, \lambda_{X_{p|P}^{n|N}}\right) \Delta W_p^{n|N} \\ X_{p|P}^{0|N} &= x_p^0, \end{aligned}$$

where $X_p^{n|N} = \{X_{p|P}^{n|N}\}_{p=1}^P$ and $\Delta W_p^{n|N} \sim \mathcal{N}\left(0, \frac{T}{N}\right)$ are i.i.d. For the remainder of this work, we use the notation

$$\phi_P^N \stackrel{\text{def}}{=} \frac{1}{P} \sum_{p=1}^P \psi\left(X_{p|P}^{N|N}\right).$$

At this point, we make the following assumptions:

$$\begin{aligned} \left|E\left[\phi_P^N - \psi(X_\infty)\right]\right| &\leq \left|E\left[\psi(X_{\cdot|P}^{N|N}) - \psi(X_{\cdot|P})\right]\right| \tag{P1} \\ &\quad + \left|E\left[\psi(X_{\cdot|P}) - \psi(X_\infty)\right]\right| \\ &\lesssim N^{-1} + P^{-1}, \end{aligned}$$

$$\text{Var}\left[\phi_P^N\right] \lesssim P^{-1}. \tag{P2}$$

These assumptions will be verified numerically in Sect. 4. In general, they translate to smoothness and boundedness assumptions on \mathcal{A} , \mathcal{B} and ψ . Indeed, in (P1), the weak convergence of the Euler–Maruyama method with respect to the number of time steps is a standard result shown, for example, in Kloeden and Platen (1992) by assuming 4-time differentiability of \mathcal{A} , \mathcal{B} and ψ . Showing that the constant multiplying N^{-1} is bounded for all P is straightforward by extending the standard proof of weak convergence the Euler–Maruyama method in (Kloeden and Platen 1992, Chapter 14) and assuming boundedness of the derivatives \mathcal{A} , \mathcal{B} and ψ . On the other hand, the weak convergence with respect to the number of particles, i.e., $E[\psi(X_{p|P})] \rightarrow E[\psi(X_\infty)]$, is a consequence of the propagation of chaos which is shown, without a convergence rate, in Sznitman (1991) for ψ Lipschitz, \mathcal{B} constant and \mathcal{A} of the form

$$\mathcal{A}(t, x, \mu) = \int \kappa(t, x, y)\mu(dy) \tag{9}$$

where $\kappa(t, \cdot, \cdot)$ is Lipschitz. On the other hand, for one-dimensional systems and using the results from Kolokoltsov and Troeva (2015, Theorem 3.2) we can show the weak convergence rate with respect to the number of particles and the convergence rate for the variance of ϕ_P as the following lemma shows. Below, $C(\mathbb{R})$ is the space of continuous bounded functions and $C^k(\mathbb{R})$ is the space of continuous bounded functions whose i 'th derivative is in $C(\mathbb{R})$ for $i = 1, \dots, k$.

Lemma 2.1 (Weak and variance convergence rates w.r.t. number of particles) *Consider (1) and (2) with $d = 1$, strictly positive $\mathcal{B}(t, \cdot, \mu) = \mathcal{B}(\cdot) \in C^3(\mathbb{R})$ and \mathcal{A} as in (9) with*

$\kappa(t, x, \cdot) \in C^2(\mathbb{R})$, $\frac{\partial \kappa(t, x, \cdot)}{\partial x} \in C(\mathbb{R})$ and $\kappa(t, \cdot, y) \in C^2(\mathbb{R})$ where the norms are assumed to be uniform with respect to the arguments, x and y , respectively. If, moreover, $\psi \in C^2(\mathbb{R})$, then

$$|E[\psi(X_{\cdot|P}) - E[\psi(X_\infty)]]| \lesssim P^{-1}, \tag{10}$$

$$\text{Var} \left[\frac{1}{P} \sum_{p=1}^P \psi(X_{p|P}) \right] \lesssim P^{-1}. \tag{11}$$

Proof The system in this lemma is a special case of the system in Kolokoltsov and Troeva (2015, Theorem 3.2). From there and given the assumptions of the current lemma, (10) immediately follows. Moreover, from the same reference, we can further conclude that

$$|E[\psi(X_{p|P})\psi(X_{q|P})] - E[\psi(X_\infty)]^2| \lesssim P^{-1}$$

for $1 \leq p \neq q \leq P$. Using this, we can show (11) since

$$\begin{aligned} \text{Var} \left[\frac{1}{P} \sum_{p=1}^P \psi(X_{p|P}) \right] &= \frac{1}{P} \text{Var}[\psi(X_{\cdot|P})] \\ &\quad + \frac{1}{P^2} \sum_{p=1}^P \sum_{q=1, p \neq q}^P \text{Cov}[\psi(X_{p|P}), \psi(X_{q|P})] \end{aligned}$$

and

$$\begin{aligned} \text{Cov}[X_{p|P}, X_{q|P}] &= E[\psi(X_{p|P})\psi(X_{q|P})] - E[\psi(X_{\cdot|P})]^2 \\ &= E[\psi(X_{p|P})\psi(X_{q|P})] - E[\psi(X_\infty)]^2 \\ &\quad - (E[\psi(X_{\cdot|P})] - E[\psi(X_\infty)])^2 \\ &\quad - 2E[\psi(X_\infty)](E[\psi(X_{\cdot|P})] \\ &\quad - E[\psi(X_\infty)]) \\ &\lesssim P^{-1}. \end{aligned}$$

□

From here, the rate of convergence for the variance of ϕ_P^N can be shown by noting that

$$\text{Var}[\phi_P^N] \leq |\text{Var}[\phi_P^N] - \text{Var}[\phi_P]| + \text{Var}[\phi_P]$$

and noting that $\text{Var}[\phi_P] \lesssim P^{-1}$, then showing that the first term is $|\text{Var}[\phi_P^N] - \text{Var}[\phi_P]| \lesssim N^{-1}P^{-1}$ because of the weak convergence with respect to the number of time steps.

Finally, as mentioned above, with a naive method, the total cost to compute a single sample of ϕ_P^N is $\mathcal{O}(NP^2)$. The quadratic power of P can be reduced by using, for example, a multipole algorithm (Carrier et al. 1988; Greengard and Rokhlin 1987). In general, we consider the work required

to compute one sample of ϕ_P^N as $\mathcal{O}(NP^{\gamma_p})$ for a positive constant, $\gamma_p \geq 1$.

3 Monte Carlo methods

In this section, we study different Monte Carlo methods that can be used to estimate the previous quantity, ϕ_∞ . In the following, we use the notation $\omega_{p:P}^{(m)} \stackrel{\text{def}}{=} (\omega_q^{(m)})_{q=p}^P$ where, for each q , $\omega_q^{(m)}$ denotes the m 'th sample of the set of underlying random variables that are used in calculating $X_{q|P}^{N|N}$, i.e., the Wiener path, W_q , the initial condition, x_q^0 , and any random variables that are used in \mathcal{A} or \mathcal{B} . Moreover, we sometimes write $\phi_P^N(\omega_{1:P}^{(m)})$ to emphasize the dependence of the m 'th sample of ϕ_P^N on the underlying random variables.

3.1 Monte Carlo (MC)

The first estimator that we look at is a Monte Carlo estimator. For a given number of samples, M , number of particles, P , and number of time steps, N , we can write the MC estimator as follows:

$$\mathcal{A}_{\text{MC}}(M, P, N) = \frac{1}{M} \sum_{m=1}^M \phi_P^N(\omega_{1:P}^{(m)}).$$

Here,

$$\begin{aligned} E[\mathcal{A}_{\text{MC}}(M, P, N)] &= E[\phi_P^N] = \frac{1}{P} \sum_{p=1}^P E[\psi(X_{p|P}^{N|N})] \\ &= E[\psi(X_{\cdot|P}^{N|N})], \end{aligned}$$

$$\text{and } \text{Var}[\mathcal{A}_{\text{MC}}(M, P, N)] = \frac{\text{Var}[\phi_P^N]}{M},$$

while the total work is $\text{Work}[\mathcal{A}_{\text{MC}}(M, P, N)] = MN P^{\gamma_p}$.

Hence, due to (P1), we must have $P = \mathcal{O}(\text{TOL}^{-1})$ and $N = \mathcal{O}(\text{TOL}^{-1})$ to satisfy (6), and, due to (P2), we must have $M = \mathcal{O}(\text{TOL}^{-1})$ to satisfy (8). Based on these choices, the total work to compute \mathcal{A}_{MC} is

$$\text{Work}[\mathcal{A}_{\text{MC}}] = \mathcal{O}(\text{TOL}^{-2-\gamma_p}).$$

Kuramoto Example Using a naive calculation method of ϕ_P^N (i.e., $\gamma_p = 2$) gives a work complexity of $\mathcal{O}(\text{TOL}^{-4})$. See also Table 1 for the work complexities for different common values of γ_p .

3.2 Multilevel Monte Carlo (MLMC)

For a given $L \in \mathbb{N}$, define two hierarchies, $\{N_\ell\}_{\ell=0}^L$ and $\{P_\ell\}_{\ell=0}^L$, satisfying $P_{\ell-1} \leq P_\ell$ and $N_{\ell-1} \leq N_\ell$ for all ℓ . Then, we can write the MLMC estimator as follows:

$$\mathcal{A}_{\text{MLMC}}(L) = \sum_{\ell=0}^L \frac{1}{M_\ell} \sum_{m=1}^{M_\ell} \left(\phi_{P_\ell}^{N_\ell} - \phi_{P_{\ell-1}}^{N_{\ell-1}} \right) \left(\omega_{1:P_\ell}^{(\ell,m)} \right), \quad (12)$$

where we later choose the function $\phi_{P_{\ell-1}}^{N_{\ell-1}}(\cdot)$ such that $\phi_{P_{\ell-1}}^{N_{\ell-1}}(\cdot) = 0$ and $\mathbb{E} \left[\phi_{P_{\ell-1}}^{N_{\ell-1}} \right] = \mathbb{E} \left[\phi_{P_{\ell-1}}^{N_{\ell-1}} \right]$, so that $\mathbb{E}[\mathcal{A}_{\text{MLMC}}] = \mathbb{E} \left[\phi_{P_L}^{N_L} \right]$ due to the telescopic sum. For MLMC to have better work complexity than that of Monte Carlo, $\phi_{P_\ell}^{N_\ell}(\omega_{1:P_\ell}^{(\ell,m)})$ and $\phi_{P_{\ell-1}}^{N_{\ell-1}}(\omega_{1:P_\ell}^{(\ell,m)})$ must be correlated for every ℓ and m , so that their difference has a smaller variance than either $\phi_{P_\ell}^{N_\ell}(\omega_{1:P_\ell}^{(\ell,m)})$ or $\phi_{P_{\ell-1}}^{N_{\ell-1}}(\omega_{1:P_\ell}^{(\ell,m)})$ for all $\ell > 0$.

Given two discretization levels, N_ℓ and $N_{\ell-1}$, with the same number of particles, P , we can generate a sample of $\phi_P^{N_{\ell-1}}(\omega_{1:P}^{(\ell,m)})$ that is correlated with $\phi_P^{N_\ell}(\omega_{1:P}^{(\ell,m)})$ by taking

$$\phi_P^{N_{\ell-1}}(\omega_{1:P}^{(\ell,m)}) = \phi_P^{N_\ell}(\omega_{1:P}^{(\ell,m)}).$$

That is, we use the same samples of the initial values, $\{x_p^0\}_{p \geq 1}$, the same Wiener paths, $\{W_p\}_{p=1}^P$, and, in case they are random as in (3), the same samples of the advection and diffusion coefficients, \mathcal{A} and \mathcal{B} , respectively. We can improve the correlation by using an antithetic sampler as detailed in Giles and Szpruch (2014) or by using a higher-order scheme like the Milstein scheme (Giles 2008a). In the Kuramoto example, the Euler–Maruyama and the Milstein schemes are equivalent since the diffusion coefficient is constant.

On the other hand, given two different sizes of the particle system, P_ℓ and $P_{\ell-1}$, with the same discretization level, N , we can generate a sample of $\phi_{P_{\ell-1}}^N(\omega_{1:P_\ell}^{(\ell,m)})$ that is correlated with $\phi_{P_\ell}^N(\omega_{1:P_\ell}^{(\ell,m)})$ by taking

$$\phi_{P_{\ell-1}}^N(\omega_{1:P_\ell}^{(\ell,m)}) = \bar{\phi}_{P_{\ell-1}}^N(\omega_{1:P_\ell}^{(\ell,m)}) \stackrel{\text{def}}{=} \phi_{P_{\ell-1}}^N(\omega_{1:P_{\ell-1}}^{(\ell,m)}). \quad (13)$$

In other words, we use the same $P_{\ell-1}$ sets of random variables out of the total P_ℓ sets of random variables to run an independent simulation of the stochastic system with $P_{\ell-1}$ particles.

We also consider another estimator that is more correlated with $\phi_{P_\ell}^N(\omega_{1:P_\ell}^{(\ell,m)})$. The ‘‘antithetic’’ estimator was first independently introduced in Haji-Ali (2012, Chapter 5) and Bujok et al. (2013) and subsequently used in other works on particle systems (Rosin et al. 2014) and nested simulations (Giles 2015). In this work, we call this estimator a ‘‘partitioning’’ estimator to clearly distinguish it from the anti-

thetic estimator in Giles and Szpruch (2014). We assume that $P_\ell = \beta_p P_{\ell-1}$ for all ℓ and some positive integer β_p and take

$$\begin{aligned} \phi_{P_{\ell-1}}^N(\omega_{1:P_\ell}^{(\ell,m)}) &= \widehat{\phi}_{P_{\ell-1}}^N(\omega_{1:P_\ell}^{(\ell,m)}) \\ &\stackrel{\text{def}}{=} \frac{1}{\beta_p} \sum_{i=1}^{\beta_p} \phi_{P_{\ell-1}}^N \left(\omega_{((i-1)P_{\ell-1}+1) : i P_{\ell-1}}^{(\ell,m)} \right). \end{aligned} \quad (14)$$

That is, we split the underlying P_ℓ sets of random variables into β_p identically distributed and independent groups, each of size $P_{\ell-1}$, and independently simulate β_p particle systems, each of size $P_{\ell-1}$. Finally, for each particle system, we compute the quantity of interest and take the average of the β_p quantities.

In the following subsections, we look at different settings in which either P_ℓ or N_ℓ depends on ℓ while the other parameter is constant for all ℓ . We begin by recalling the optimal convergence rates of MLMC when applied to a generic random variable, Y , with a trivial generalization to the case when there are two discretization parameters: one that is a function of the level, ℓ , and the other, \tilde{L} , that is fixed for all levels.

Theorem 3.1 (Optimal MLMC complexity) *Let $Y_{\tilde{L},\ell}$ be an approximation of the random variable, Y , for every $(\tilde{L}, \ell) \in \mathbb{N}^2$. Denote by $Y^{(\ell,m)}$ a sample of Y and denote its corresponding approximation by $Y_{\tilde{L},\ell}^{(\ell,m)}$, where we assume that the samples $\{Y^{(\ell,m)}\}_{\ell,m}$ are mutually independent. Consider the MLMC estimator*

$$\mathcal{A}_{\text{MLMC}}(\tilde{L}, L) = \sum_{\ell=0}^L \frac{1}{M_\ell} \sum_{m=1}^{M_\ell} \left(Y_{\tilde{L},\ell}^{(\ell,m)} - Y_{\tilde{L},\ell-1}^{(\ell,m)} \right)$$

with $Y_{\tilde{L},-1}^{\ell,m} = 0$ and for $\beta, w, \gamma, s, \tilde{\beta}, \tilde{w}, \tilde{\gamma}, \tilde{c} > 0$ where $s \leq 2w$, and assume the following:

1. $|\mathbb{E}[Y - Y_{\tilde{L},\ell}]| \lesssim \tilde{\beta}^{-\tilde{w}\tilde{L}} + \beta^{-w\ell}$
2. $\text{Var}[Y_{\tilde{L},\ell} - Y_{\tilde{L},\ell-1}] \lesssim \tilde{\beta}^{-\tilde{c}\tilde{L}} \beta^{-s\ell}$
3. $\text{Work}[Y_{\tilde{L},\ell} - Y_{\tilde{L},\ell-1}] \lesssim \tilde{\beta}^{\tilde{\gamma}\tilde{L}} \beta^{\gamma\ell}$.

Then, for any $\text{TOL} < e^{-1}$, there exist \tilde{L}, L and a sequence of $\{M_\ell\}_{\ell=0}^L$ such that

$$\mathbb{P}[|\mathcal{A}_{\text{MLMC}}(\tilde{L}, L) - Y| \geq \text{TOL}] \leq \epsilon \quad (15)$$

and

$$\begin{aligned} \text{Work}[\mathcal{A}_{\text{MLMC}}(\tilde{L}, L)] &\stackrel{\text{def}}{=} \sum_{\ell=0}^L M_\ell \text{Work}[Y_{\tilde{L},\ell} - Y_{\tilde{L},\ell-1}] \\ &\lesssim \begin{cases} \text{TOL}^{-2-\frac{\tilde{\gamma}-\tilde{c}}{w}} & \text{if } s > \gamma \\ \text{TOL}^{-2-\frac{\tilde{\gamma}-\tilde{c}}{w}} \log(\text{TOL}^{-1})^2 & \text{if } s = \gamma \\ \text{TOL}^{-2-\frac{\tilde{\gamma}-\tilde{c}}{w}-\frac{\gamma-s}{w}} & \text{if } s < \gamma. \end{cases} \end{aligned} \tag{16}$$

Proof The proof can be straightforwardly derived from the proof of Cliffe et al. (2011, Theorem 1), we sketch here the main steps. First, we split constraint (15) to a bias and variance constraints similar to (6) to (8), respectively. Then, since $E[\mathcal{A}_{\text{MLMC}}(\tilde{L}, L)] = E[Y_{\tilde{L},L}]$, given the first assumption of the theorem and imposing the bias constraint yield $\tilde{L} = \mathcal{O}\left(\frac{1}{\tilde{w} \log(\beta)} \log(\text{TOL}^{-1})\right)$ and $L = \mathcal{O}\left(\frac{1}{w \log(\beta)} \log(\text{TOL}^{-1})\right)$. The assumptions on the variance and work then give:

$$\begin{aligned} \text{Var}[Y_{\tilde{L},\ell} - Y_{\tilde{L},\ell-1}] &\lesssim \text{TOL}^{\frac{\tilde{c}}{w}} \beta^{-s\ell}, \\ \text{Work}[Y_{\tilde{L},\ell} - Y_{\tilde{L},\ell-1}] &\lesssim \text{TOL}^{-\frac{\tilde{\gamma}}{w}} \beta^{\gamma\ell}. \end{aligned}$$

Then

$$\begin{aligned} \text{Var}[\mathcal{A}_{\text{MLMC}}(\tilde{L}, L)] &= \sum_{\ell=0}^L M_\ell^{-1} \text{Var}[Y_{\tilde{L},\ell} - Y_{\tilde{L},\ell-1}] \lesssim \text{TOL}^{\frac{\tilde{c}}{w}} \sum_{\ell=0}^L M_\ell^{-1} \beta^{-s\ell}, \end{aligned}$$

due to mutual independence of $\{Y^{(\ell,m)}\}_{\ell,m}$. Moreover,

$$\begin{aligned} \text{Work}[\mathcal{A}_{\text{MLMC}}(\tilde{L}, L)] &= \sum_{\ell=0}^L M_\ell \text{Work}[Y_{\tilde{L},\ell} - Y_{\tilde{L},\ell-1}] \\ &\lesssim \text{TOL}^{-\frac{\tilde{\gamma}}{w}} \sum_{\ell=0}^L M_\ell^{-1} \beta^{\gamma\ell} \end{aligned}$$

Finally, given L , solving for $\{M_\ell\}_{\ell=0}^L$ to minimize the work while satisfying the variance constraint gives the desired result. \square

3.2.1 MLMC hierarchy based on the number of time steps

In this setting, we take $N_\ell = (\beta_t)^\ell$ for some $\beta_t > 0$ and $P_\ell = P_L$ for all ℓ , i.e., the number of particles is a constant, P_L , on all levels. We make an extra assumption in this case, namely:

$$\text{Var}[\phi_{P_L}^{N_\ell} - \phi_{P_L}^{N_{\ell-1}}] \lesssim P_L^{-1} N_\ell^{-s_t} = P_L^{-1} (\beta_t)^{-s_t \ell}, \tag{MLMC1}$$

for some constant $s_t > 0$. The factor $(\beta_t)^{-s_t \ell}$ is the usual assumption on the variance convergence of the level difference in MLMC theory (Giles 2008b) and is a standard result for the Euler–Maruyama scheme with $s_t = 1$ and for the Milstein scheme with $s_t = 2$, (Kloeden and Platen 1992). On the other hand, the factor P_L^{-1} can be motivated from (P2), which states that the variance of each term in the difference converges at this rate.

Due to Theorem 3.1, we can conclude that the work complexity of MLMC is

$$\text{Work}[\mathcal{A}_{\text{MLMC}}] \lesssim \begin{cases} \text{TOL}^{-1-\gamma_p} & \text{if } s_t > 1 \\ \text{TOL}^{-1-\gamma_p} \log(\text{TOL}^{-1})^2 & \text{if } s_t = 1 \\ \text{TOL}^{-2-\gamma_p+s_t} & \text{if } s_t < 1. \end{cases} \tag{17}$$

Kuramoto Example In this example, using the Milstein time-stepping scheme, we have $s_t = 2$ (cf. Fig. 1), and a naive calculation method of ϕ_P^N ($\gamma_p = 2$) gives a work complexity of $\mathcal{O}(\text{TOL}^{-3})$. See also Table 1 for the work complexities for different common values of s_t and γ_p .

3.2.2 MLMC hierarchy based on the number of particles

In this setting, we take $P_\ell = (\beta_p)^\ell$ for some $\beta_p > 0$ and $N_\ell = N_L$ for all ℓ , i.e., we take the number of time steps to be a constant, N_L , on all levels. We make an extra assumption in this case:

$$\text{Var}[\phi_{P_\ell}^{N_L} - \phi_{P_{\ell-1}}^{N_L}] \lesssim P_\ell^{-s_p-1} = \beta_p^{-\ell(s_p+1)}, \tag{MLMC2}$$

for some constant $s_p \geq 0$. The factor $\beta_p^{-s_p \ell}$ is the usual assumption on the variance convergence of the level difference in MLMC theory (Giles 2008b). On the other hand, the factor P_ℓ^{-1} can be motivated from (P2), since the variance of each term in the difference is converging at this rate.

Due to Theorem 3.1, we can conclude that the work complexity of MLMC in this case is

$$\text{Work}[\mathcal{A}_{\text{MLMC}}] \lesssim \begin{cases} \text{TOL}^{-3} & \text{if } s_p + 1 > \gamma_p \\ \text{TOL}^{-3} \log(\text{TOL}^{-1})^2 & \text{if } s_p + 1 = \gamma_p \\ \text{TOL}^{-2-\gamma_p+s_p} & \text{if } s_p + 1 < \gamma_p. \end{cases} \tag{18}$$

Kuramoto Example Using a naive calculation method of ϕ_P^N ($\gamma_p = 2$), we distinguish between the two samplers:

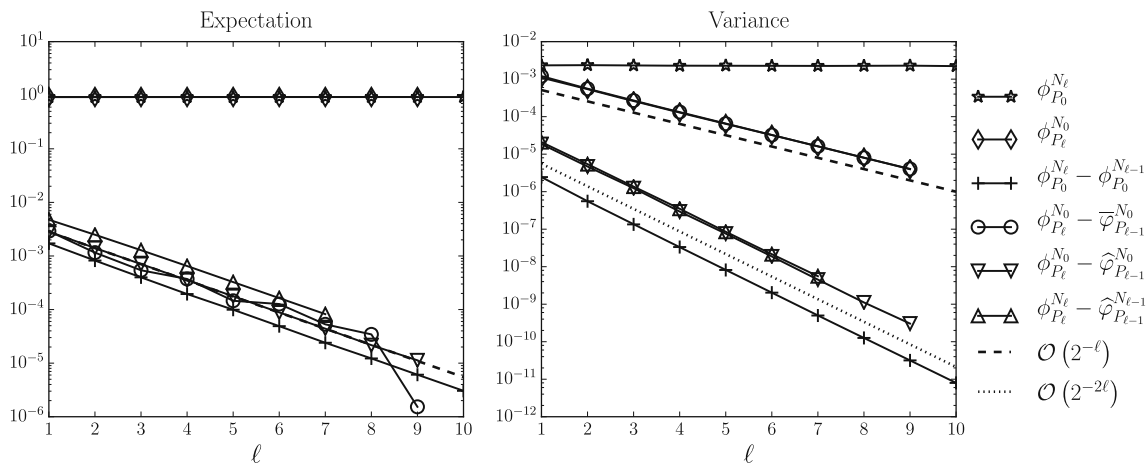


Fig. 1 This plot shows, for the Kuramoto example (3), numerical evidence for Assumption (P1) (left) and Assumptions (P2), (MLMC1)–(MLMC3) (right). Here, P_ℓ and N_ℓ are chosen according to (22). From the right plot, we can confirm that $s_t = 2$ for the Milstein method. We

can also deduce that using the $\bar{\varphi}$ sampler in (13) yields $s_p = 0$ in (MIMC2) (i.e., no variance reduction compared to $\text{Var}[\phi_{P_\ell}^{N_\ell}]$), while using the $\hat{\varphi}$ sampler in (14) yields $s_p = 1$ in (MIMC2) (i.e., $\mathcal{O}(P^{-2})$)

Table 1 The work complexity of the different methods presented in this work in common situations, encoded as (a, b) to represent $\mathcal{O}(\text{TOL}^{-a}(\log(\text{TOL}^{-1}))^b)$

Method	$s_t = 1, \gamma_p = 1$	$s_t = 1, \gamma_p = 2$	$s_t = 2, \gamma_p = 1$	$s_t = 2, \gamma_p = 2$
MC (Sect. 3.1)	(3, 0)	(4, 0)	(3, 0)	(4, 0)
MLMC (Sect. 3.2.1)	(2, 2)	(3, 2)	(2, 0)	(3, 0)
MLMC (Sect. 3.2.2)	(3, 0)	(3, 2)	(3, 0)	(3, 2)
MLMC (Sect. 3.2.3)	(2, 2)	(3, 0)	(2, 2)	(3, 0)
MIMC (Sect. 3.3)	(2, 2)	(2, 4)	(2, 0)	(2, 2)

When appropriate, we use the partitioning estimator (i.e., $s_p = 1$). In general, MIMC has always the best complexity. However, when $\gamma_p = 1$ MIMC does not offer an advantage over an appropriate MLMC method

- Using the sampler $\bar{\varphi}$ in (13), we verify numerically that $s_p = 0$ (cf. Fig. 1). Hence, the work complexity is $\mathcal{O}(\text{TOL}^{-4})$, which is the same work complexity as a Monte Carlo estimator. This should be expected since using the “correlated” samples of $\bar{\varphi}_{P_{\ell-1}}^N$ and $\phi_{P_\ell}^N$ does not reduce the variance of the difference, as Fig. 1 shows.
- Using the partitioning estimator, $\hat{\varphi}$, in (14), we verify numerically that $s_p = 1$ (cf. Fig. 1). Hence, the work complexity is $\mathcal{O}(\text{TOL}^{-3} \log(\text{TOL}^{-1})^2)$. Here the samples of $\hat{\varphi}_{P_{\ell-1}}^N$ have higher correlation to corresponding samples of $\phi_{P_\ell}^N$, thus reducing the variance of the difference. Still, using MLMC with hierarchies based on the number of times steps (fixing the number of particles) yields better work complexity. See also Table 1 for the work complexities for different common values of s_t and γ_p .

3.2.3 MLMC hierarchy based on both the number of particles and the number of times steps

In this case, we vary both the number of particles and the number of time steps across MLMC levels. That is, we take

$P_\ell = (\beta_p)^\ell$ and $N_\ell = (\beta_t)^\ell$ for all ℓ . In this case, a reasonable assumption is

$$\text{Var}[\phi_{P_\ell}^{N_\ell} - \phi_{P_{\ell-1}}^{N_{\ell-1}}] \lesssim (\beta_p)^{-\ell} (\max((\beta_p)^{-s_p}, (\beta_t)^{-s_t}))^\ell. \tag{MLMC3}$$

The factor $\beta_p^{-\ell}$ can be motivated from (P2) since the variance of each term in the difference is converged at this rate. On the other hand, $(\max(\beta_p^{-s_p}, \beta_t^{-s_t}))^\ell$ is the larger factor of (MLMC1) and (MLMC2).

Due to Theorem 3.1 and defining

$$\begin{aligned} s &= \log(\beta_p) + \min(s_p \log(\beta_p), s_t \log(\beta_t)) \\ \gamma &= \gamma_p \log(\beta_p) + \log(\beta_t) \\ w &= \min(\log(\beta_p), \log(\beta_t)), \end{aligned}$$

we can conclude that the work complexity of MLMC is

$$\text{Work}[\mathcal{A}_{\text{MLMC}}] \lesssim \begin{cases} \text{TOL}^{-2} & \text{if } s > \gamma \\ \text{TOL}^{-2} \log(\text{TOL}^{-1})^2 & \text{if } s = \gamma \\ \text{TOL}^{-2 - \frac{\gamma - s}{w}} & \text{if } s < \gamma. \end{cases} \tag{19}$$

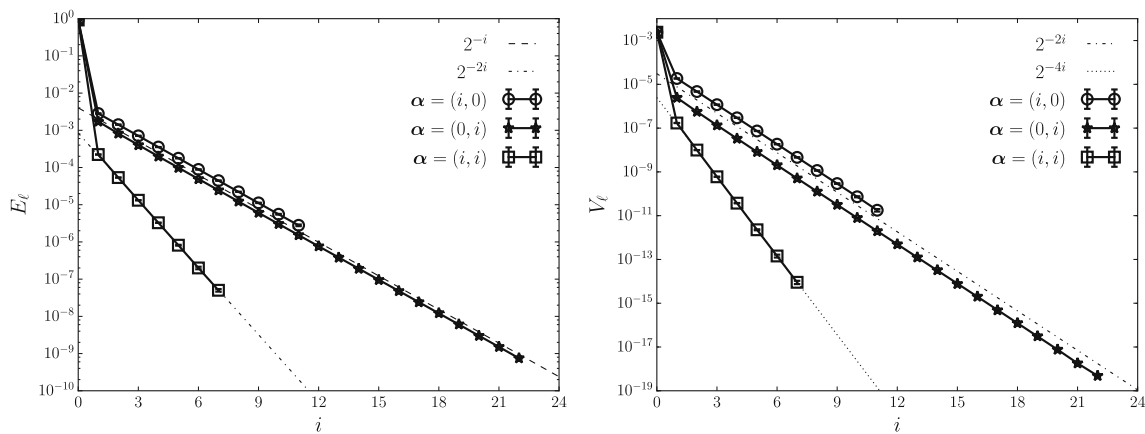


Fig. 2 This figure provides numerical evidence for (MIMC1) (left) and (MIMC2) (right) for the Kuramoto example (3) when using the Milstein scheme for time discretization (yielding $s_t = 2$) and the par-

titioning sample of particle systems (yielding $s_p = 1$). Here, P_{α_1} and N_{α_2} are chosen according to (22). When considering a mixed difference (i.e. $\alpha = (i, i)$), a higher rate of convergence is observed

Kuramoto Example We choose $\beta_p = \beta_t$ and use a naive calculation method of ϕ_P^N (yielding $\gamma_p = 2$) and the partitioning sampler (yielding $s_p = 1$). Finally, using the Milstein time-stepping scheme, we have $s_t = 2$. Refer to Fig. 1 for numerical verification. Based on these rates, we have, in (19), $s = 2 \log(\beta_p)$, $w = \log(\beta_p)$ and $\gamma = 3 \log(\beta_p)$. The MLMC work complexity in this case is $\mathcal{O}(\text{TOL}^{-3})$. See also Table 1 for the work complexities for different common values of s_t and γ_p .

3.3 Multi-index Monte Carlo (MIMC)

Following (Haji-Ali et al. 2015a), for every multi-index $\alpha = (\alpha_1, \alpha_2) \in \mathbb{N}^2$, let $P_{\alpha_1} = (\beta_p)^{\alpha_1}$ and $N_{\alpha_2} = (\beta_t)^{\alpha_2}$ and define the first-order mixed-difference operator in two dimensions as

$$\Delta \phi_{P_{\alpha_1}}^{N_{\alpha_2}}(\omega_{1:P_{\alpha_1}}) = \left(\left(\phi_{P_{\alpha_1}}^{N_{\alpha_2}} - \phi_{P_{\alpha_1-1}}^{N_{\alpha_2}} \right) - \left(\phi_{P_{\alpha_1}}^{N_{\alpha_2-1}} - \phi_{P_{\alpha_1-1}}^{N_{\alpha_2-1}} \right) \right) (\omega_{1:P_{\alpha_1}})$$
 with $\phi_{P_{\alpha_1}}^0 = 0$ and $\phi_P^{N_{\alpha_2}-1} = 0$. The MIMC estimator is then written for a given $\mathcal{I} \subset \mathbb{N}^2$ as

$$\mathcal{A}_{\text{MIMC}} = \sum_{\alpha \in \mathcal{I}} \frac{1}{M_{\alpha}} \sum_{m=1}^{M_{\alpha}} \Delta \phi_{P_{\alpha_1}}^{N_{\alpha_2}}(\omega_{1:P_{\alpha_1}}^{(\alpha, m)}) \quad (20)$$

At this point, similar to the original work on MIMC (Haji-Ali et al. 2015a), we make the following assumptions on the convergence of $\Delta \phi_{P_{\alpha_1}}^{N_{\alpha_2}}$, namely

$$E \left[\Delta \phi_{P_{\alpha_1}}^{N_{\alpha_2}} \right] \lesssim P_{\alpha_1}^{-1} N_{\alpha_2}^{-1} \quad (\text{MIMC1})$$

$$\text{Var} \left[\Delta \phi_{P_{\alpha_1}}^{N_{\alpha_2}} \right] \lesssim P_{\alpha_1}^{-s_p-1} N_{\alpha_2}^{-s_t}. \quad (\text{MIMC2})$$

Assumption (MIMC1) is motivated from (P1) by assuming that the mixed first-order difference, $\Delta \phi_{P_{\alpha_1}}^{N_{\alpha_2}}$, gives a prod-

uct of the convergence terms instead of a sum. Similarly, (MIMC2) is motivated from (MLMC1) and (MLMC2). To the best of our knowledge, there are currently no proofs of these assumptions for particle systems, but we verify them numerically for (3) in Fig. 2.

Henceforth, we will assume that $\beta_t = \beta_p$ for easier presentation. Following (Haji-Ali et al. 2015a, Lemma 2.1) and recalling the assumption on cost per sample, $\text{Work} \left[\Delta \phi_{P_{\alpha_1}}^{N_{\alpha_2}} \right] \lesssim P_{\alpha_1}^{\gamma_p} N_{\alpha_2}$, then, for every value of $L \in \mathbb{R}^+$, the optimal set can be written as

$$\mathcal{I}(L) = \left\{ \alpha \in \mathbb{N}^2 : (1 - s_p + \gamma_p)\alpha_1 + (3 - s_t)\alpha_2 \leq L \right\}, \quad (21)$$

and the optimal computational complexity of MIMC is $\mathcal{O}(\text{TOL}^{-2-2 \max(0, \zeta)} \log(\text{TOL}^{-1})^p)$, where

$$\zeta = \max \left(\frac{\gamma_p - s_p - 1}{2}, \frac{1 - s_t}{2} \right),$$

$$\xi = \min \left(\frac{2 - s_p}{\gamma_p}, 2 - s_t \right) \geq 0,$$

$$p = \begin{cases} 0 & \zeta < 0 \\ 2\mathfrak{z} & \zeta = 0 \\ 2(\mathfrak{z} - 1)(\zeta + 1) & \zeta > 0 \text{ and } \xi > 0 \\ 1 + 2(\mathfrak{z} - 1)(\zeta + 1) & \zeta > 0 \text{ and } \xi = 0 \end{cases}$$

$$\text{and } \mathfrak{z} = \begin{cases} 1 & \gamma_p - s_p - 1 \neq 1 - s_t \\ 2 & \gamma_p - s_p - 1 = 1 - s_t. \end{cases}$$

Kuramoto Example Here again, we use a naive calculation method of ϕ_P^N (yielding $\gamma_p = 2$) and the partitioning

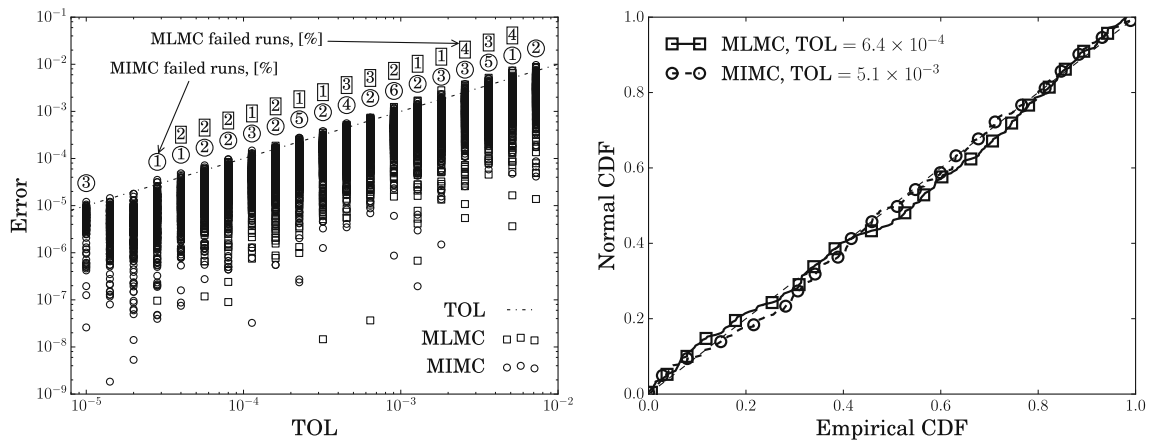


Fig. 3 In these plots, each marker represents a separate run of the MLMC or MIMC estimators (as detailed in Sects. 3.2.3 and 3.3, respectively) when applied to the Kuramoto example (3). *Left* the exact errors of the estimators, estimated using a reference approximation that was computed with a very small TOL. This plot shows that, up to the prescribed 95% confidence level, both methods approximate the quantity of interest to the same required tolerance, TOL. The upper and lower

numbers above the *linear line* represent the percentage of runs that failed to meet the prescribed tolerance, if any, for both MLMC and MIMC, respectively. *Right* A PP plot of the CDF of the value of both estimators for certain tolerances, shifted by their mean and scaled by their standard deviation showing that both estimators, when appropriately shifted and scaled, are well approximated by a standard normal random variable




Fig. 4 Maximum discretization level of the number of time steps and the number of particles for both MLMC and MIMC (as detailed in Sects. 3.2.3 and 3.3, respectively) for different tolerances (cf. (22)). Recall that, for a fixed tolerance in MIMC, $2\alpha_2 + \alpha_1$ is bounded by a constant (cf. (21))

sampler (yielding $s_p = 1$). Finally, using the Milstein time-stepping scheme, we have $s_t = 2$. Hence, $\zeta = 0$, $\mathfrak{z} = 1$ and $\text{Work}[A_{\text{MIMC}}] = \mathcal{O}\left(\text{TOL}^{-2} \log(\text{TOL}^{-1})^2\right)$. See also Table 1 for the work complexities for different common values of s_t and γ_p .

4 Numerical example

In this section, we provide numerical evidence of the assumptions and work complexities that are made in Sect. 3. This section also verifies that the constants of the work complexity

(which were not tracked) are not significant for reasonable error tolerances. The results in this section were obtained using the `mimc`lib software library (Haji-Ali 2016) and GNU `parallel` (Tange 2011).

In the results outlined below, we focus on the Kuramoto example in (3), with the following choices: $\sigma = 0.4$, $T = 1$, $x_p^0 \sim \mathcal{N}(0, 0.2)$ and $\vartheta_p \sim \mathcal{U}(-0.2, 0.2)$ for all p . We also set

$$P_\ell = 5 \times 2^\ell \quad \text{and} \quad N_\ell = 4 \times 2^\ell \quad \text{for MLMC,}$$

$$\text{and} \quad P_{\alpha_1} = 5 \times 2^{\alpha_1} \quad \text{and} \quad N_{\alpha_2} = 4 \times 2^{\alpha_2} \quad \text{for MIMC.} \tag{22}$$

Figure 1 shows the absolute expectation and variance of the level differences for the different MLMC settings that are outlined in Sect. 3.2. These figures verify Assumptions (P1), (P2) and (MLMC1)–(MLMC3) with the values $s_t = 2$ and $s_p = 0$ for the $\bar{\varphi}$ sampler in (13) or the value $s_p = 1$ for the $\hat{\varphi}$ sampler in (14). For the same parameter values, Fig. 2 provides numerical evidence for Assumptions (MIMC1) and (MIMC2) for the $\hat{\varphi}$ sampler (14).

We now compare the MLMC method (Giles 2008b) in the setting that is presented in Sect. 3.2.3 and the MIMC method (Haji-Ali et al. 2015a) that is presented in Sect. 3.3. In both methods, we use the Milstein time-stepping scheme and the partitioning sampler, $\hat{\varphi}$, in (14). Recall that in this case, we verified numerically that $\gamma_p = 2$, $s_p = 1$ and $s_t = 2$. We also use the MLMC and MIMC algorithms that were outlined in their original work and use an initial 25 samples on each level or multi-index to compute a corresponding variance estimate that is required to compute the optimal number of samples. In

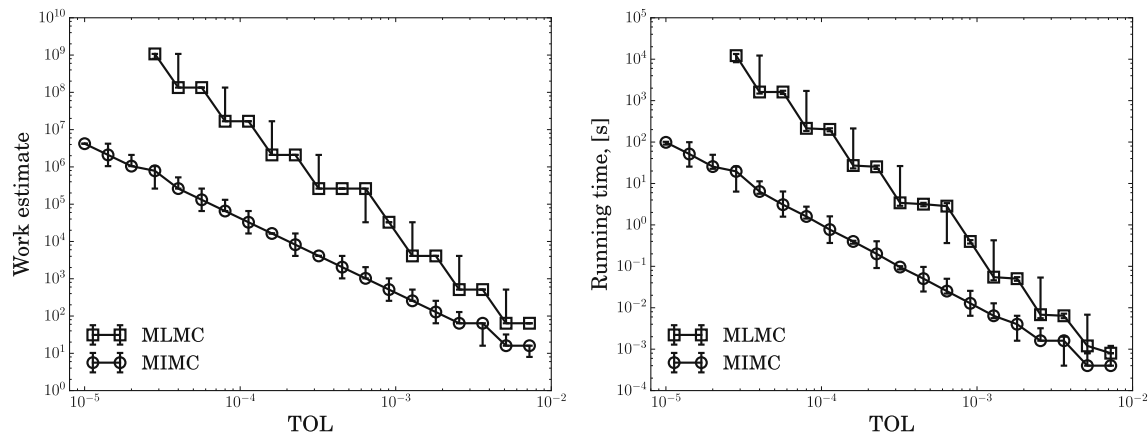


Fig. 5 Maximum work estimate (*left*) and running time (in seconds, *right*) of the samples used in MLMC and MIMC (as detailed in Sects. 3.2.3 and 3.3, respectively) when applied to the Kuramoto exam-

ple (3). These plots show the single, largest indivisible work unit in MLMC and MIMC which gives an indication of the parallelization scaling of both methods

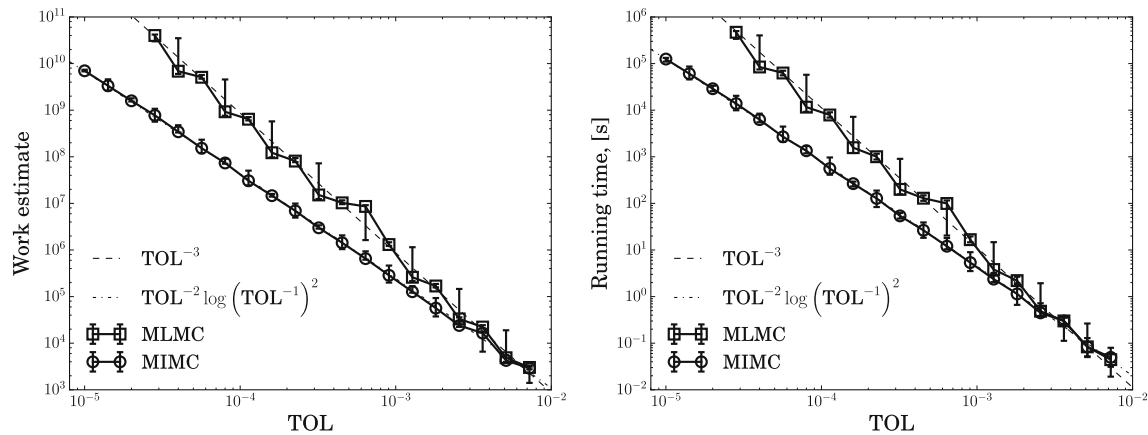


Fig. 6 Work estimate (*left*) and running time (in seconds, *right*) of MLMC and MIMC (as detailed in Sects. 3.2.3 and 3.3, respectively) when applied to the Kuramoto example (3). For sufficiently small tol-

erances, the running time closely follows the predicted theoretical rates (also plotted) and shows the performance improvement of MIMC

the following, we refer to these methods as simply “MLMC” and “MIMC.” We focus on the settings in Sects. 3.2.3 and 3.3 since checking the bias of the estimator in those settings can be done straightforwardly by checking the absolute value of the level differences in MLMC or the multi-index differences in MIMC. On the other hand, checking the bias in the settings outlined in Sects. 3.1, 3.2.1 and 3.2.2 is not as straightforward and determining the number of time steps and/or the number of particles to satisfy a certain error tolerance requires more sophisticated algorithms. This makes a fair numerical comparison with these later settings somewhat difficult.

Figure 3-left shows the exact errors of both MLMC and MIMC for different prescribed tolerances. This plot shows that both methods estimate the quantity of interest up to the same error tolerance; comparing their work complexity is thus fair. On the other hand, Fig. 3-right shows a PP plot,

i.e., a plot of the cumulative distribution function (CDF) of the MLMC and MIMC estimators, normalized by their variance and shifted by their mean, versus the CDF of a standard normal distribution. This figure shows that our assumption in Sect. 2 of the asymptotic normality of these estimators is well founded. Figure 4 shows the maximum discretization level for both the number of time steps and the number of particles for MLMC and MIMC (cf. (22)). Recall that, for a fixed tolerance in MIMC, $2\alpha_2 + \alpha_1$ is bounded by a constant (cf. (21)). Hence, Fig. 4 shows a direct implication on the results reported in Fig. 5 where we plot the maximum cost of the samples used in both MLMC and MIMC for different tolerances. This cost represents an indivisible unit of simulation for both methods, assuming we treat the simulation of the particle system as a black box. Hence, Fig. 5 shows that MIMC has better parallelization scaling, i.e., even with an

infinite number of computation nodes MIMC would still be more efficient than MLMC.

Finally, we show in Fig. 6 the cost estimates of MLMC and MIMC for different tolerances. This figure clearly shows the performance improvement of MIMC over MLMC and shows that the complexity rates that we derived in this work are reasonably accurate.

5 Conclusions

This work has shown both numerically and theoretically under certain assumptions, which could be verified numerically, the improvement of MIMC over MLMC when used to approximate a quantity of interest computed on a particle system as the number of particles goes to infinity. The application to other particle systems (or equivalently other McKean–Vlasov SDEs) is straightforward, and similar improvements are expected. The same machinery was also suggested for approximating nested expectations in Giles (2015), and the analysis here applies to that setting as well. Moreover, the same machinery, i.e., multi-index structure with respect to time steps and number of particles coupled with a partitioning estimator, could be used to create control variates to reduce the computational cost of approximating quantities of interest on stochastic particle systems with a finite number of particles.

Future work includes analyzing the optimal level separation parameters, β_p and β_t , and the behavior of the tolerance splitting parameter, θ . Another direction could be applying the MIMC method to higher-dimensional particle systems such as the crowd model in Haji-Ali (2012). On the theoretical side, the next step is to prove the assumptions that were postulated and verified numerically in this work for certain classes of particle systems, namely:

the second-order convergence with respect to the number of particles of the variance of the partitioning estimator (14) and the convergence rates for mixed differences (MIMC1) and (MIMC2).

Acknowledgements R. Tempone is a member of the KAUST Strategic Research Initiative, Center for Uncertainty Quantification in Computational Sciences and Engineering. R. Tempone received support from the KAUST CRG3 Award Ref: 2281 and the KAUST CRG4 Award Ref: 2584. The authors would like to thank Lukas Szpruch for the valuable discussions regarding the theoretical foundations of the methods.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Acebrón, J.A., Bonilla, L.L., Vicente, C.J.P., Ritort, F., Spigler, R.: The Kuramoto model: a simple paradigm for synchronization phenomena. *Rev. Mod. Phys.* **77**(1), 137 (2005)
- Bossy, M., Talay, D.: Convergence rate for the approximation of the limit law of weakly interacting particles: application to the Burgers equation. *Ann. Appl. Probab.* **6**(3), 818–861 (1996)
- Bossy, M., Talay, D.: A stochastic particle method for the McKean–Vlasov and the Burgers equation. *Math. Comput. Am. Math. Soc.* **66**(217), 157–192 (1997)
- Bujok, K., Hambly, B., Reisinger, C.: Multilevel simulation of functionals of Bernoulli random variables with application to basket credit derivatives. *Methodol. Comput. Appl. Probab.* **73**, 1–26 (2013)
- Carrier, J., Greengard, L., Rokhlin, V.: A fast adaptive multipole algorithm for particle simulations. *SIAM J. Sci. Stat. Comput.* **9**(4), 669–686 (1988)
- Cliffe, K., Giles, M., Scheichl, R., Teckentrup, A.: Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Comput. Vis. Sci.* **14**(1), 3–15 (2011)
- Collier, N., Haji-Ali, A.-L., Nobile, F., von Schwerin, E., Tempone, R.: A continuation multilevel Monte Carlo algorithm. *BIT Numer. Math.* **55**(2), 399–432 (2015)
- Dobramysl, U., Rüdiger, S., Erban, R.: Particle-based multiscale modeling of calcium puff dynamics. *Multiscale Model. Simul.* **14**(3), 997–1016 (2016)
- Erban, R., Haskovec, J.: From individual to collective behaviour of coupled velocity jump processes: a locust example. *Kinet. Relat. Models* **5**(4), 817–842 (2012)
- Erban, R., Haskovec, J., Sun, Y.: A Cucker–Smale model with noise and delay. *SIAM J. Appl. Math.* **76**(4), 1535–1557 (2016)
- Gärtner, J.: On the McKean–Vlasov limit for interacting diffusions. *Mathematische Nachrichten* **137**(1), 197–248 (1988)
- Giles, M.B.: Improved multilevel Monte Carlo convergence using the Milstein scheme. In: Keller, A., Heinrich, S., Niederreiter, H. (eds.) *Monte Carlo and Quasi-Monte Carlo Methods 2006*, pp. 343–358. Springer, Berlin (2008a)
- Giles, M.B.: Multilevel Monte Carlo path simulation. *Oper. Res.* **56**(3), 607–617 (2008b)
- Giles, M.B.: Multilevel Monte Carlo methods. *Acta Numer.* **24**, 259–328 (2015)
- Giles, M.B., Szpruch, L.: Antithetic multilevel Monte Carlo estimation for multi-dimensional SDEs without Lévy area simulation. *Ann. Appl. Probab.* **24**(4), 1585–1620 (2014)
- Greengard, L., Rokhlin, V.: A fast algorithm for particle simulations. *J. Comput. Phys.* **73**(2), 325–348 (1987)
- Haji-Ali, A.-L.: Pedestrian flow in the mean-field limit. King Abdullah University of Science and Technology (KAUST). <http://hdl.handle.net/10754/250912> (2012)
- Haji-Ali, A.-L.: mimclib. <https://github.com/StochasticNumerics/mimclib> (2016)
- Haji-Ali, A.-L., Nobile, F., Tempone, R.: Multi-index Monte Carlo: when sparsity meets sampling. *Numer. Math.* **132**, 767–806 (2015a)
- Haji-Ali, A.-L., Nobile, F., von Schwerin, E., Tempone, R.: Optimization of mesh hierarchies in multilevel Monte Carlo samplers. *Stoch. Partial Differ. Equ. Anal. Comput.* **4**, 76–112 (2015b)
- Heinrich, S.: Multilevel Monte Carlo methods. In: *Large-Scale Scientific Computing*, vol. 2179 of Lecture Notes in Computer Science. Springer, Berlin, pp. 58–67 (2001)
- Helbing, D., Molnar, P.: Social force model for pedestrian dynamics. *Phys. Rev. E* **51**(5), 4282 (1995)
- Kloeden, P.E., Platen, E.: Numerical solution of stochastic differential equations. Springer, Berlin (1992). doi:10.1007/978-3-662-12616-5

- Kolokoltsov, V., Troeva, M.: On the mean field games with common noise and the McKean–Vlasov SPDEs. ArXiv preprint [arXiv:1506.04594](https://arxiv.org/abs/1506.04594) (2015)
- Pierre Del Moral, A.K., Tugaut, J.: On the stability and the uniform propagation of chaos of a class of extended Ensemble Kalman–Bucy filters. *SIAM J. Control Optim.* **55**(1), 119–155 (2016)
- Ricketson, L.: A multilevel Monte Carlo method for a class of McKean–Vlasov processes. ArXiv preprint [arXiv:1508.02299](https://arxiv.org/abs/1508.02299) (2015)
- Rosin, M., Ricketson, L., Dimits, A., Caflisch, R., Cohen, B.: Multilevel Monte Carlo simulation of Coulomb collisions. *J. Comput. Phys.* **274**, 140–157 (2014)
- Sznitman, A.-S.: Topics in propagation of chaos. In: *Ecole d’été de probabilités de Saint-Flour XIX–1989*. Springer, pp. 165–251 (1991)
- Tange, O.: GNU parallel—the command-line power tool. *login USENIX Mag.* **36**(1), 42–47 (2011)