Exact Simulation of Multivariate Itô Diffusions

Jose Blanchet
Joint work with Fan Zhang

Columbia and Stanford

July 7, 2017

Table of Contents

Introduction

Consider the multivariate Itô Stochastic Differential Equations (SDE)

$$dX(t) = \mu(X(t))dt + \sigma(X(t))dW(t),$$

where $X(\cdot)$ and $\mu(\cdot)$ are d-dimensional vectors; $W(\cdot)$ is a d'-dimensional Brownian motion.

Introduction

Consider the multivariate Itô Stochastic Differential Equations (SDE)

$$dX(t) = \mu(X(t))dt + \sigma(X(t))dW(t),$$

where $X(\cdot)$ and $\mu(\cdot)$ are d-dimensional vectors; $W(\cdot)$ is a d'-dimensional Brownian motion.

Question: Is it possible to simulate $X(T)$ exactly?

Previous Work

- (Beskos & Roberts (2005)) First exact simulation algorithm of SDE in "$d = 1$", under boundedness conditions.
- (Beskos et al. (2006), Chen & Huang (2013)) Relaxed boundedness assumption but "$d = 1$" as well.

Key Assumptions on Current Generic Methods

Existing methodos require the following:

Key Assumptions on Current Generic Methods

Existing methodos require the following:

- Can apply Lamperti's transform, which maps $X(\cdot)$ into a constant diffusion matrix diffusion.

Key Assumptions on Current Generic Methods

Existing methodos require the following:

- Can apply Lamperti's transform, which maps $X(\cdot)$ into a constant diffusion matrix diffusion.
- Drift coefficient $\mu(\cdot)$ is the gradient of some functions (i.e. $\nabla v(\cdot) = \mu(\cdot)$).

Key Assumptions on Current Generic Methods

Existing methodos require the following:

- Can apply Lamperti's transform, which maps $X(\cdot)$ into a constant diffusion matrix diffusion.
- Drift coefficient $\mu(\cdot)$ is the gradient of some functions (i.e. $\nabla v(\cdot) = \mu(\cdot)$).

Question: Is it possible to simulate $X(T)$ exactly for generic diffusions?

The Good News:

Yes! It is possible.

The Bad News:

Drawback:

- The running time of our algorithm, although finite with probability one, has infinite mean.

The Bad News:

Drawback:

- The running time of our algorithm, although finite with probability one, has infinite mean.

But please don't go! This is still interesting because:

The Bad News:

Drawback:

- The running time of our algorithm, although finite with probability one, has infinite mean.

But please don't go! This is still interesting because:

- This is the first exact simulation algorithm for generic multidimensional diffusions (aren't you curious?).

The Bad News:

Drawback:

- The running time of our algorithm, although finite with probability one, has infinite mean.

But please don't go! This is still interesting because:

- This is the first exact simulation algorithm for generic multidimensional diffusions (aren't you curious?).
- This requires a novel conceptual framework, different from Lamperti's transformation.

The Bad News:

Drawback:

- The running time of our algorithm, although finite with probability one, has infinite mean.

But please don't go! This is still interesting because:

- This is the first exact simulation algorithm for generic multidimensional diffusions (aren't you curious?).
- This requires a novel conceptual framework, different from Lamperti's transformation.
- Check out interesting algorithmic ideas: Multilevel Monte Carlo, Bernoulli factories and ε-strong simulation.

Table of Contents

Model Setup

- Assume $X(\cdot)$ is $\tilde{\mathbb{P}}$-Brownian motion.
- Define

$$L(t) = \exp\left(\sum_{i=1}^{d}\int_0^t \mu_i(X(t))dX_i(t) - \frac{1}{2}\int_0^t \|\mu(X(t))\|^2 dt\right),$$

- If $\mu(\cdot)$ is Lipchitz continuous, then $L(\cdot)$ is a martingale, define

$$d\mathbb{P} = L(T)d\tilde{\mathbb{P}}$$

- There exist a $\mathbb{P}$-Brownian motion $W(\cdot)$, such that

$$X(t) = X(0) + \int_0^t \mu(X(s))ds + W(t); \quad 0 \le t \le T.$$

Direct Algorithm by Acceptance-Rejection

Direct Acceptance-Rejection

- **Step 1**: Propose $X(T)$ under measure $\tilde{\mathbb{P}}$, which is Normal distribution.
- **Step 2**: Accept the proposal with probability proportional to $L(T)$.

Direct Algorithm by Acceptance-Rejection

Direct Acceptance-Rejection

- **Step 1**: Propose $X(T)$ under measure $\tilde{\mathbb{P}}$, which is Normal distribution.
- **Step 2**: Accept the proposal with probability proportional to $L(T)$.

There are two challenges in the execution of Step 2:

1. Unboundedness of $L(T)$.
2. "Intractability" of stochastic integral in $L(T)$.

Conditional on integer part of $L(T)$

- Assume we can sample $L(T)$ under $\mathbb{P}$ and $\tilde{\mathbb{P}}$.
- Suppose $U \sim \text{Unif}(0,1)$ independent of everything else,

$$\mathbb{P}(X(T) \in dx | \lfloor L(T) \rfloor = k)$$
$$= \tilde{\mathbb{P}}(X(T) \in dx | (k+1)U < L(T); \lfloor L(T) \rfloor = k)$$

Conditional on integer part of $L(T)$

- Assume we can sample $L(T)$ under $\mathbb{P}$ and $\tilde{\mathbb{P}}$.
- Suppose $U \sim \text{Unif}(0,1)$ independent of everything else,

$$\begin{aligned}&\mathbb{P}(X(T) \in dx \,|\, \lfloor L(T) \rfloor = k) \\ &\qquad = \tilde{\mathbb{P}}(X(T) \in dx \,|\, (k+1)U < L(T); \lfloor L(T) \rfloor = k)\end{aligned}$$

Acceptance-Rejection with Unbounded Likelihood Ratio

- **Step 1**: Sample $\lfloor L(T) \rfloor$ under measure $\mathbb{P}$, let $k = \lfloor L(T) \rfloor$.
- **Step 2**: Sample $U \sim \text{Unif}(0,1)$, let $u = U$.
- **Step 3**: Sample $(X(T), L(T))$ jointly from $\tilde{\mathbb{P}}$ until $\max((k+1)u, k) < L(T) < k+1$.

Observation

- We don't need to simulate $L(T)$ without any bias.
- It suffices to be able to simulate $\lfloor L(T) \rfloor$ and/or $I(a < L(T) < b)$.
- Approximate $L(T)$ by ε-strong simulation.

ε-Strong Simulation for SDEs

Theorem (Blanchet, Chen and Dong (2017))

Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ *and the following multidimensional SDE:*

$$dY(t) = \alpha(Y(t))dt + \nu(Y(t))dW(t), \quad Y(t) = y_0 \tag{1}$$

Suppose that $\alpha(\cdot)$ *and* $\nu(\cdot)$ *satisfy suitable regularity conditions. Then, given any deterministic* $\varepsilon > 0$*, we can simulate a piecewise constant process* $Y_\varepsilon(\cdot)$*, such that*

$$\sup_{t\in[0,T]} \|Y_\varepsilon(t) - Y(t)\|_2 \leq \varepsilon \quad a.s.$$

Furthermore, for any $m > 1$ *and* $0 < \varepsilon_m < \cdots < \varepsilon_1 < 1$*, we can simulate* Y_{ε_m} *conditional on* $Y_{\varepsilon_1}, \ldots, Y_{\varepsilon_{m-1}}$*.*

Illustration of ε-Strong Simulation

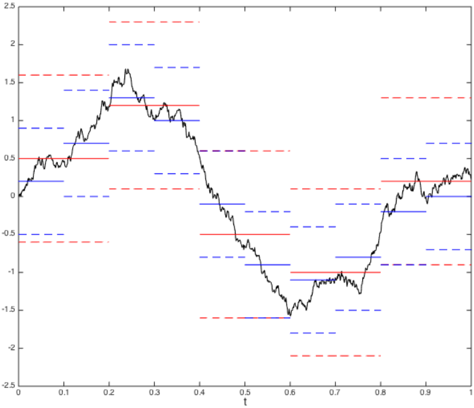

Figure: Illustation for ε-Strong Simulation in one dimension. Black Path = True; Red Line = Coarse Approximation; Red Dashed Line = Coarse Error Bound; Blue Line = Refined Approximation; Blue Dashed Line = Refined Error Bound;

Localization for SDEs: $d = 1$ for Illustration

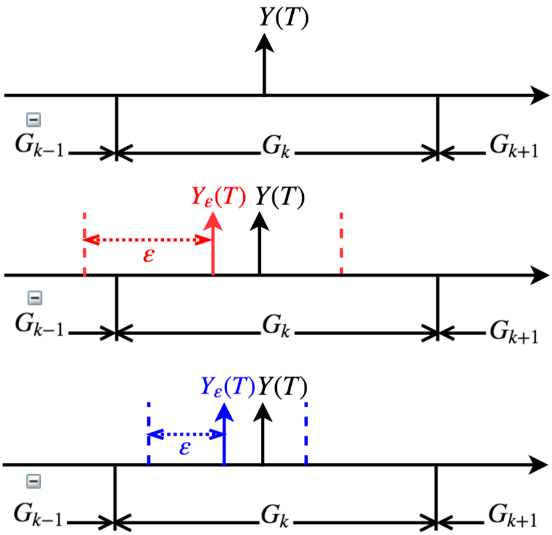

ε-Strong Simulation of $(X(\cdot), L(\cdot))$

- Under measure $\mathbb{P}$:

$$\begin{cases} dL(t) = L(t)\|\mu(X(t))\|_2^2 dt + \mu^T(X(t))dW(t), \\ dX(t) = \mu(X(t))dt + dW(t), \end{cases}$$

- Under measure $\tilde{\mathbb{P}}$:

$$\begin{cases} dL(t) = L(t)\mu^T(X(t))dX(t). \\ dX(t) = dX(t). \end{cases}$$

- ε-strong simulation is applicable to $(X(\cdot), L(\cdot))$ under both measure.
- Under measure $\tilde{\mathbb{P}}$, the first step of ε-strong simulation is sampling $X(T)$.

Algorithm

Exact Simulation of SDEs with Identity Diffusion Matrix

- **Step 1**: Sample $\lfloor L(T) \rfloor$ under measure $\mathbb{P}$ by ε-strong simulation, let $k = \lfloor L(T) \rfloor$.
- **Step 2**: Sample $U \sim \text{Unif}(0,1)$, let $u = U$.
- **Step 3**: Apply ε-strong simulation to sample $I(\max((k+1)u, k) < L(T) < k+1)$ and $X(T)$ under measure $\tilde{\mathbb{P}}$.
- **Step 4**: Output $X(T)$ if $I(\max((k+1)u, k) < L(T) < k+1) = 1$; otherwise return to Step 2.

Table of Contents

When Lamperti's Transform is not Applicable

$$dX(t) = \mu(X(t))dt + \sigma(X(t))dW(t).$$

- Girsanov's theorem can't substantially simplify the problem.
- Idea: Construct an unbiased estimator of density of $X(T)$.

Construction of Density Estimator

- Suppose $r_n \searrow 0$; $B_{r_n}(x)$ = Ball with radius r_n centered at x;

$$\hat{p}_n(x) = [V(B_{r_n}(0))]^{-1} \times I(X(T) \in B_{r_n}(x)); \hat{p}_0(x) = 0.$$

- We can sample $\hat{p}_n(x)$ by ε-strong simulation.
- Due to smoothness of density

$$p(x) = \mathbb{E}\left[\sum_{n=0}^{\infty} (\hat{p}_{n+1}(x) - \hat{p}_n(x))\right].$$

- Connection to multi-level Monte Carlo (Giles (2008), Rhee-Glynn (2015))

Multilevel Monte Carlo Estimator

- For a given integer r.v. N, define

$$\Lambda_n(x) = \sum_{k=0}^{n} \frac{\hat{p}_{k+1}(x) - \hat{p}_k(x)}{\mathbb{P}(N \geq k)} \quad \text{for} \quad n \geq 0.$$

- $p(x) = \mathbb{E}[\Lambda_N(x)]$.
- We want the estimator to be nonnegative and bounded.

On Non-negativity and Boundedness

Key Lemma

For any compact set G, there exist a family of random variables $\{\Lambda_n^+(x); n \in \mathbb{N}, x \in G\}$ and computable constants m and $\{m_n; n \in \mathbb{N}\}$, such that the following properties hold:

1. $0 \leq \Lambda_n^+(x) \leq m_n < \infty$.
2. $0 \leq \mathbb{E}[\Lambda_n^+(x)] = \mathbb{E}[\Lambda_n(x)] \leq m, \quad \forall n \in \mathbb{N}$.
3. Given n and x, there is an algorithm to simulate $\Lambda_n^+(x)$.

The constants m and $\{m_n; n \in \mathbb{N}\}$ depend only on set G and bounds on $\mu(\cdot), \sigma(\cdot)$ and their derivatives of order $1, 2, 3$.

Introduction of Extra Randomness

- $\Lambda_N^+(x)$ is unbiased estimator of $p(x)$, but it is unbounded.
- We introduce an auxiliary r.v. N' coupled with $X(T)$ in the following way

$$\mathbb{P}(N' = n | X(T) = x) \propto \mathbb{P}(N = n) \times E[\Lambda_n^+(x)]$$

- Given $N' = n$ and $X(T) \in G$, the conditional density is easy to sample

$$\mathbb{P}(X(T) \in dx | N' = n, X(T) \in G) \propto \mathbb{E}[\Lambda_n^+(x)].$$

Outline of the Algorithm

1. Sample a compact set G s.t. $X(T) \in G$.
2. Sample $(N'|X(T) \in G)$.
3. Sample $(X(T)|N', X(T) \in G)$.

Step 1: Localization of $X(T)$

- Divide the space $\mathbb{R}$ into disjoint boxes.
- Apply ε-strong simulation to decide the box G to which $X(T)$ belongs.

Step 2: Simulation of N' conditional on $X(T) \in G$

$$\mathbb{P}(N' = n | X(T) \in G) = \frac{\mathbb{P}(N = n)}{\mathbb{P}(X(T) \in G)} \times \int_G \mathbb{E}[\Lambda_n^+(x)] dx.$$

- Use N as proposal for N'.
- Given $N = n$, accept the proposal with probability $(mV(G))^{-1} \int_G \mathbb{E}[\Lambda_n^+(x)] dx$.
  - Use upper bounds - recall m and m_n;

Step 3: Simulation of $X(T)$ conditional on $X(T) \in G$ and N'

$$\mathbb{P}(X(T) \in dx | N' = n, X(T) \in G) \propto \mathbb{E}[\Lambda_n^+(x)],$$

- Sample x from uniform distribution of G.
- Accept x as $X(T)$ with probability $\propto \mathbb{E}[\Lambda_n^+(x)]$
  - Sample $\Lambda_n^+(x)$ and $U \sim \text{Unif}(0, m_n)$.
  - Accept x if $U \leq \Lambda_n^+(x)$.

Conclusions

- First exact sampler for generic d-dimensional diffusions.
- Algorithm exposes the role of ε-strong simulation.
- Interesting unbiased estimators (density).
- Interplay of several debiasing techniques.
- Key Open Problem: Can we execute $Y(T) \in G$ in finite expected time?