

Numerical Solution of Dyson Brownian Motion and a Sampling Scheme for Invariant Matrix Ensembles

Xingjie Helen Li · Govind Menon

Received: 12 June 2013 / Accepted: 26 September 2013 / Published online: 9 October 2013
© Springer Science+Business Media New York 2013

Abstract The Dyson Brownian Motion (DBM) describes the stochastic evolution of N points on the line driven by an applied potential, a Coulombic repulsion and identical, independent Brownian forcing at each point. We use an explicit tamed Euler scheme to numerically solve the Dyson Brownian motion and sample the equilibrium measure for non-quadratic potentials. The Coulomb repulsion is too singular for the SDE to satisfy the hypotheses of rigorous convergence proofs for tamed Euler schemes (Hutzenthaler et al. in Ann. Appl. Probab. 22(4):1611–1641, 2012). Nevertheless, in practice the scheme is observed to be stable for time steps of $O(1/N^2)$ and to relax exponentially fast to the equilibrium measure with a rate constant of $O(1)$ independent of N . Further, this convergence rate appears to improve with N in accordance with $O(1/N)$ relaxation of local statistics of the Dyson Brownian motion. This allows us to use the Dyson Brownian motion to sample $N \times N$ Hermitian matrices from the invariant ensembles. The computational cost of generating M independent samples is $O(MN^4)$ with a naive scheme, and $O(MN^3 \log N)$ when a fast multipole method is used to evaluate the Coulomb interaction.

Keywords Random matrix theory · Invariant matrix ensembles · Dyson Brownian motion (DBM) · Explicit tamed Euler scheme

1 Introduction

1.1 The Coulomb Gas on the Line and Dyson Brownian Motion

Our main purpose in this article is to sample the equilibrium measure of the Coulomb gas with a given (non-quadratic) potential V . The Hamiltonian for this particle system is

$$H_N(\lambda) = \frac{1}{2} \sum_{k=1}^N V(\lambda_k) - \frac{1}{N} \sum_{1 \leq j < k \leq N} \log(\lambda_j - \lambda_k). \quad (1.1)$$

X.H. Li · G. Menon (✉)
Brown University, Providence, RI, USA
e-mail: govind_menon@brown.edu

Here $\lambda = (\lambda_1, \dots, \lambda_N)$ denotes the positions of N particles on the line and $V(s)$ is a polynomial in one variable with a term of highest degree $\alpha_{2m}s^{2m}$ with $\alpha_{2m} > 0$. We find it convenient to break permutation symmetry and work with the domain $\mathcal{W}_N = \{\lambda \in \mathbb{R}^N \mid \lambda_1 < \lambda_2 < \dots < \lambda_N\}$. At fixed inverse temperature $\beta > 0$, we define $\mu_{\beta, V}(\lambda)$ as the probability measure on \mathcal{W}_N with density proportional to $e^{-\beta N H_N(\lambda)}$. The measure $\mu_{\beta, V}$ balances two competing effects: Coulombic repulsion pushes particles apart, but the potential V confines them globally.

The Gibbs measure $\mu_{\beta, V}$ for the Coulomb gas is also the unique equilibrium measure for the Dyson Brownian motion (DBM) with potential V [11]. Let $B_k(t), k = 1, \dots, N$ denote N independent standard Brownian motions (i.e. $\mathbb{E}B_k(t) = 0, \mathbb{E}(B_k(t))^2 = t$). The Dyson Brownian motion $\lambda(t) \in \mathcal{W}_N$ is the solution to the stochastic differential equation

$$\begin{aligned}
 d\lambda_k &= -\partial_{\lambda_k} H_N(\lambda) dt + \sqrt{\frac{2}{\beta N}} dB_k, \\
 &= \left(\frac{1}{N} \sum_{j:j \neq k} \frac{1}{\lambda_k - \lambda_j} - \frac{1}{2} V'(\lambda_k) \right) dt + \sqrt{\frac{2}{\beta N}} dB_k, \quad k = 1, \dots, N. \tag{1.2}
 \end{aligned}$$

This SDE is weak and strong well-posed under very general assumptions [1, Theorem 4.3.2]. The time scale in (1.2) is chosen so that a typical initial distribution relaxes to the equilibrium measure $\mu_{\beta, V}$ in time of $O(1)$. These relaxation properties underlie our simulation scheme and allow us to rationally estimate its computational cost as follows.

We use a tamed explicit Euler scheme to numerically solve (1.2). (Note that a numerical solution is necessary, except when $V(x)$ is quadratic). Our scheme requires time steps of size $O(N^{-2})$. The naive computational cost of summing all Coulomb interactions at each time step is $O(N^2)$, but this can be reduced to $O(N \log N)$ with a fast multipole method. Moreover, since all λ_k are real, the domain decomposition required for the fast multipole method is very straightforward (see for example [2, Sect. 2.1]). Thus, M samples can be obtained with a computational effort proportional to $O(MN^3 \log N)$. In the work presented here, N is sufficiently small (≤ 300) that we did not use a fast multipole method. The performance of our scheme (in particular, stability and speed of convergence to equilibrium) is not adversely affected by N . As discussed below, local convergence of DBM occurs at a rate $O(1/N)$ and this seems to *improve* the performance of the scheme as N increases. Finally, as we explain below, sampling from $\mu_{\beta, V}$ allows us to sample random Hermitian matrices from the invariant ensembles with potential V .

1.2 Invariant Hermitian Ensembles

Let $\text{Her}(N)$ denote the space of $N \times N$ Hermitian matrices and $U(N)$ the group of $N \times N$ unitary matrices. A probability distribution \mathbb{P} on $\text{Her}(N)$ is *invariant* if for every fixed $V \in U(N)$, and every Borel set A in $\text{Her}(N)$, we have $\mathbb{P}(M \in A) = \mathbb{P}(VMV^* \in A)$. We focus on ensembles with a density $p_V(M)$ of the form

$$p_V(M) dM = \frac{1}{Z_N} e^{-N \text{tr} V(M)} dM = \frac{1}{Z_N} e^{-N \text{tr} V(M)} \prod_{k=1}^N dM_{kk} \prod_{j < k} dM_{jk}^R dM_{jk}^I. \tag{1.3}$$

Here V is the potential of (1.1). The constant Z_N is a normalizing factor that ensures $\int_{\text{Her}(N)} p_V(M) dM = 1$, and $M_{jk} = M_{jk}^R + \sqrt{-1} M_{jk}^I$ denote the real and imaginary parts of the matrix entry M_{jk} .

Our initial motivation for this work was to sample random matrices from the distribution $p_V(M)$ when V is not quadratic (i.e. the ensemble is non-Gaussian). Our interest in this problem stems from observations of universal fluctuations in the time to convergence for Hamiltonian eigenvalue algorithms [21]. Universality was observed for a variety of Wigner ensembles, but the algorithms were not tested on matrices drawn from non-Gaussian invariant ensembles. The principal difficulty in sampling from an invariant ensemble is that the entries M_{jk} , $j \leq k$, of a random matrix M are not independent (except for Gaussian ensembles). By contrast, the Wigner ensembles have independent entries and are easy to sample from. While our work arose from the desire to test eigenvalue algorithms on random matrices chosen from invariant ensembles, the problem of sampling from these ensembles is of independent interest. For example, it is well known that matrix integrals serve as generating functions for random maps [3]. Thus, it is of interest to connect sampling schemes for matrices with combinatorial algorithms for sampling random maps (e.g. [9]).

Somewhat surprisingly, despite the considerable theoretical interest in invariant ensembles [4, 8, 19], there appears to have been no prior effort to numerically generate random matrices from these ensembles. In simultaneous work, Olver *et al.* have also considered this sampling problem [22]. Their approach uses numerical methods for Riemann-Hilbert problems and a constructive scheme for sampling determinantal processes, and is completely different from the methods used here. In combination, both works suggest some interesting numerical challenges arising in random matrix theory. In particular, our work reveals the unexpected utility of Dyson Brownian motion as a test bed for numerical methods for stochastic differential equations.

1.3 Determinantal Structure

The law of invariant ensembles is most easily described through the spectral representation $M = UxU^*$, where $x = \text{diag}(x_1, \dots, x_N)$ denotes a diagonal matrix of (unordered) eigenvalues and U is a unitary matrix of eigenvectors. We then have

$$e^{-N \text{tr} V(M)} dM = e^{-N \sum_{k=1}^N V(x_k)} \Delta^2(x) dx dU. \tag{1.4}$$

Here $\Delta(x)$ is the Vandermonde determinant $\prod_{1 \leq j < k \leq N} (x_j - x_k)$, dU and dx denote the volume elements for normalized Haar measure on $U(N)$ and Lebesgue measure on \mathbb{R}^N respectively, and $H_N(x)$ is the Hamiltonian defined in (1.1). The squared Vandermonde determinant $\Delta^2(x)$ is the Jacobian of the change of variables $M \mapsto (x, U)$ [7, 19]. When included in the Hamiltonian H_N as above, it has an appealing physical interpretation: the eigenvalues of invariant ensembles repel one another according to the Coulomb law and their equilibrium measure is the same as the Coulomb gas at inverse temperature $\beta = 2$. (A strict comparison requires that we reorder the x 's in increasing order and include a permutation factor of $N!$ in the measure. We use x to denote the unordered eigenvalues and λ the ordered eigenvalues to avoid confusion).

The eigenvalues x_1, \dots, x_N constitute a determinantal point process on \mathbb{R} described explicitly as follows. Let $\pi_k(s) = s^j + \dots$ be the monic orthogonal polynomials with respect to the measure $e^{-NV(s)} ds$ on \mathbb{R} , and define normalizing constants c_k , orthonormal functions ϕ_k in $L^2(\mathbb{R})$, and an integral kernel K_N as follows:

$$c_k^2 = \int_{\mathbb{R}} \pi_k^2(s) e^{-NV(s)} ds, \quad \phi_k(s) = \frac{1}{c_k} e^{-NV(s)/2} \pi_k(s), \quad K_N(r, s) = \sum_{k=0}^{N-1} \phi_k(r) \phi_k(s). \tag{1.5}$$

Then the m -point correlation function is given by

$$\rho_m(x_1, \dots, x_m) = \frac{(N - m)!}{N!} \det(K_N(x_j, x_k)_{1 \leq j, k \leq m}). \tag{1.6}$$

(the normalization conventions in mathematics and physics appear to differ in the inclusion of the combinatorial factor of $(N - m)!/N!$ (contrast [7, p. 108] or [15, p. 209] with [19, p. 80])). The particular case $m = 1$ yields the density of eigenvalues. Under suitable assumptions on V (e.g. that V is convex), the limiting density of eigenvalues as $N \rightarrow \infty$ is described by a unique equilibrium density:

$$\lim_{N \rightarrow \infty} \frac{1}{N} \#\{x_k \in (a, b)\} = \int_a^b \rho(s) ds, \quad -\infty < a < b < \infty. \tag{1.7}$$

The equilibrium measure for a quartic potential $V(x)$ is described explicitly in (2.2) below.

1.4 Dyson Brownian Motion as a Sampling Method

Our task is to sample a random matrix M distributed according to (1.3). This is equivalent to sampling x and U distributed according to (1.4). Now it is immediate from (1.4) that the laws of x and U are independent. Stewart derived an efficient method for sampling random matrices from the orthogonal group $O(N)$ distributed according to the Haar measure [23]. In essence, he used the fact that the QR decomposition of a GOE matrix yields a matrix Q that is uniformly distributed on $O(N)$. Some care is needed when extending this scheme to $U(N)$, and Mezzadri presents an efficient sampling scheme in [20]. Thus, our task is only to sample the eigenvalues according to the law (1.6). Since the m -point correlation function (1.6) is identical to that of the Coulomb gas at $\beta = 2$, it is sufficient to solve Dyson Brownian motion until the equilibrium is approached. This strategy is practical because of certain fundamental convergence properties of Dyson Brownian motion reviewed below.

The convergence to $\mu_{\beta, V}$ of an arbitrary initial law under DBM has been carefully studied in random matrix theory. The recent emphasis has been on the use of DBM to establish universal fluctuations when the initial law of the eigenvalues is determined by a Wigner ensemble (see in particular [12, Sect. 2] and the references therein). However, some of the basic techniques used to establish convergence of DBM do not use the fact that $V(x) = x^2/2$ in an essential way and apply to all initial laws (including invariant ensembles). These general properties guide our numerics. Assume $\lambda(t)$ solves (1.2), and denote the distribution of $\lambda(t)$ by $f_t(\lambda)\mu_{\beta, V}$. Then $f_t\mu_{\beta, V}$ satisfies the Fokker-Planck equation associated to (1.2), and f_t satisfies the evolution equation

$$\partial_t f_t = \frac{1}{\beta N} \sum_{k=1}^N \partial_{\lambda_k}^2 f_t - \sum_{k=1}^N \partial_{\lambda_k} H(\lambda) \partial_{\lambda_k} f_t. \tag{1.8}$$

Let us assume that $V(x)$ is a convex function. Then a very minor modification of the argument in [12, Sect. 2] shows that there is a constant $\theta > 0$ depending only on β and V such that

$$\int_{\mathbb{R}} |f_t - 1| d\mu_{\beta, V} \leq C e^{-\theta t}, \tag{1.9}$$

provided f_0 has finite relative entropy (relative to 1) and defines a finite Dirichlet form.

To summarize: under mild assumptions on the initial distribution, the law $f_t \mu_{\beta, V}$ converges exponentially to $\mu_{\beta, V}$ with a rate constant of $O(1)$. One of the key recent developments in random matrix theory is a rigorous proof that the local statistics of λ relax to equilibrium much faster, at a time scale $O(1/N)$ (This is termed Dyson’s conjecture in [12]). The $O(1/N)$ relaxation rate underlies the universality of fluctuations in random matrix theory. Our view is that the relaxation rates of DBM are not just of fundamental theoretical importance, but that they may be exploited in a practical sampling algorithm. We have only used the $O(1)$ convergence in this work, but it is plausible that a more sophisticated sampling scheme could be devised using the $O(1/N)$ convergence, reducing the computational effort to $O(MN^2 \log N)$.

1.5 The Numerical Method

In order to convert DBM into an effective sampling algorithm, we need a fast numerical method for the SDE (1.2). Sampling can then proceed by parallel runs with independently chosen initial data. An implicit solver for the SDE (1.2) is too expensive for this process, and it is necessary to use an explicit scheme. But here we run into an interesting complication: the well-posedness of standard explicit Euler schemes for SDE require that the vector field $\nabla_\lambda H(\lambda)$ be globally Lipschitz continuous [14, 18]. Moreover, counterexamples show that the Lipschitz assumption is necessary [16]. Now note that the vector field for DBM is *not* Lipschitz because of the singular nature of the Coulomb interaction. While the counterexamples in [16] do not include DBM, we have found numerically that the standard explicit Euler scheme for DBM fails to converge. This failure turns out to depend on the growth of V' rather than the strong repulsion of eigenvalues and we were able to resolve it using the following tamed explicit Euler scheme that is a modification of the scheme suggested by Kloeden and his co-authors [17].

Let Δt denote the time step, $T > 0$ a fixed time of $O(1)$, and $\lambda^n = (\lambda_1^n, \lambda_2^n, \dots, \lambda_N^n)$ the discretized approximation to the solution $\lambda(t)$ to (1.2) at the time $t = n\Delta t$. Assume also that ΔB_k^n are independent standard (mean zero, variance one) normal random variables for $k = 1, \dots, N, n = 0, \dots, T/\Delta t$. We use the following numerical scheme:

$$\lambda_k^{n+1} = \lambda_k^n + \left(\frac{1}{N} \sum_{j \neq k} \frac{1}{\lambda_k^n - \lambda_j^n} - \frac{V'(\lambda_k^n)}{2 + \Delta t |V'(\lambda_k^n)|} \right) \Delta t + \sqrt{\frac{2\Delta t}{\beta N}} \Delta B_k^n, \tag{1.10}$$

$$\lambda_k^0 = \xi_k. \tag{1.11}$$

The index k denotes a coordinate in \mathcal{W}_N and runs from 1 to N . The index n is the time step and runs from 0 to $T/\Delta t$. The initial data $\lambda_k^0 = \xi_k$ are assumed given for the purpose of solving the SDE. In practice, they will be chosen at random as described below.

The main points to note are: (i) this scheme is explicit; (ii) the vector field $V'(x)$ has been “tamed”; and (iii) the Coulomb interaction is left unchanged. Since the scheme (1.10) can be rewritten as

$$\lambda_k^{n+1} = \lambda_k^n + \left(\frac{1}{N} \sum_{j \neq k} \frac{1}{\lambda_k^n - \lambda_j^n} - \frac{1}{2} V'(\lambda_k^n) \right) \Delta t + \sqrt{\frac{2\Delta t}{\beta N}} \Delta B_k^n + \frac{V'(\lambda_k^n) |V'(\lambda_k^n)|}{4 + 2\Delta t |V'(\lambda_k^n)|} (\Delta t)^2, \tag{1.12}$$

it differs from the standard explicit Euler scheme

$$\lambda_k^{n+1} = \lambda_k^n + \left(\frac{1}{N} \sum_{j \neq k} \frac{1}{\lambda_k^n - \lambda_j^n} - \frac{1}{2} V'(\lambda_k^n) \right) \Delta t + \sqrt{\frac{2\Delta t}{\beta N}} \Delta B_k^n,$$

by only a second-order correction. Since the explicit Euler scheme is consistent with the SDE, so is the tamed scheme. Thus, stability of the scheme suffices to establish convergence. However, as a consequence of the Coulomb interaction, it does not satisfy the assumptions of the main theorem in [17]. It is not clear to us at present, if the techniques used to establish well-posedness of the SDE (1.2) can also be used to establish stability of the numerical scheme.

2 Numerical Experiments

2.1 Equilibrium Measures for Quartic Ensembles

In this section we present the result of numerical experiments with the scheme (1.10) for the quartic potentials

$$V(x) = \frac{qx^2}{2} + \frac{gx^4}{4}, \quad q, g \geq 0. \tag{2.1}$$

The quartic ensembles are the simplest non-Gaussian invariant ensembles and provide important insights into general phenomena for invariant ensembles. For example, quartic ensembles were used to enumerate random quadrangulations [3], and to establish the first universality theorems for non-Gaussian ensembles [4]. Thus, we expect that they have similar utility for a numerical study.

In all that follows, we assume $\beta = 2$ so that the samples λ correspond to the eigenvalues of a Hermitian matrix. If $g = 0$ we obtain a Gaussian ensemble. For $q > 0$ we may always rescale so that $q = 1$. When $q = 1$ and $0 \leq g < \infty$, the equilibrium density for the quartic ensemble is given explicitly by [3]

$$\rho(s) = \frac{1}{2\pi} (1 + 2ga^2 + gs^2) \sqrt{4a^2 - s^2} \mathbf{1}_{|s| \leq 2a}, \quad a = \sqrt{\frac{\sqrt{1 + 12g} - 1}{6g}}. \tag{2.2}$$

In the limit $g = 0$ this density reduces to Wigner’s semicircle law. When $q = 0$ the potential $V(x)$ is purely quartic and has equilibrium density

$$\rho(s) = \frac{1}{2\pi} (2ga^2 + gs^2) \sqrt{4a^2 - s^2} \mathbf{1}_{|s| \leq 2a}, \quad a = (3g)^{-1/4}. \tag{2.3}$$

2.2 Description of Numerical Experiments

We fix a potential (i.e. the parameters g and q in (2.1)) and perform M independent trials. Each trial takes the following form:

- (1) Choose an initial vector λ^0 at random independent of other trials.
- (2) Solve (1.10) for λ^n for each $n = 1, \dots, T/\Delta t$.

The initial vector λ^0 is typically chosen to be the vector of eigenvalues of a Gaussian Hermitian matrix. This choice is not entirely necessary, since the universality results show that the spectrum of a random matrix from a Wigner ensemble that satisfies the fourth-moment condition $\mathbb{E}(|M_{jk}|^4) \leq C < \infty$ will relax in time of $O(1)$ to the equilibrium density $\mu_{\beta, V}$ for DBM [12]. However, we find that there is a distinct difference in the time to equilibration between a vector λ^0 that is obtained as the eigenvalues of a random matrix, or a vector λ^0 consisting of independent entries $\lambda_k^0, k = 1, \dots, N$.

The typical end time is taken to be $T = O(1)$ and all simulations are observed to equilibrate when the time step is chosen to be $\Delta t = 1/N^2$. Thus, each run involves N^2 evaluations of the vector field, and each such evaluation costs $O(N^2)$, with a total computational cost of $O(N^4)$. As we have remarked above, this can be reduced to $O(N^3 \log N)$ with a fast multipole method. There is a further computational cost associated to the initialization, since we must compute the eigenvalues of the initial matrix. However, this cost is of lower order if the initial matrix is chosen from a tridiagonal ensemble whose eigenvalues have the law $e^{-\beta \tilde{H}_N(x)}$ where \tilde{H}_N is the Hamiltonian corresponding to $V(x) = qx^2$ (for example, the β -ensembles of Dumitriu and Edelman [10]). The total cost for M trials is $O(MN^4)$.

For M parallel trials, we denote by $\lambda_k^{n,j}$ the k -th coordinate of the solution to equation (1.10) at time step n for the j -th trial. The empirical distribution function is defined to be

$$F_{n,M}(\lambda) = \frac{1}{NM} \sum_{j=1}^M \sum_{k=1}^N \mathbf{1}_{\lambda_k^{n,j} \leq \lambda}, \tag{2.4}$$

and the exact distribution function and equilibrium distribution function are denoted

$$F_N(\lambda) = \frac{1}{N} \int_{-\infty}^{\lambda} K_N(s, s) ds, \quad F_{\infty}(\lambda) = \int_{-\infty}^{\lambda} \rho(s) ds. \tag{2.5}$$

We use the Kolmogorov-Smirnov (KS) statistic

$$D_{n,M} = \sup_{\lambda} |F_{n,M}(\lambda) - F_N(\lambda)| \tag{2.6}$$

to measure the distance between the empirical measure $F_{n,M}$ and the exact distribution F_N . For large enough N (in practice, $N \geq 30$) we replace F_N with F_{∞} in the formula since the error $\|F_N - F_{\infty}\|_{\infty}$ is much smaller than $D_{n,M}$. This is discussed below.

The above comparison only contrasts 1-point statistics of the empirical and exact distributions of $x(t)$. In order to study convergence of a statistic that truly reflects the determinantal nature of the process, we compare the exact probability $A(\theta)$ that there are no points $x_k(t)$ in the interval $(-\theta, \theta)$ with the empirical measure of $A(\theta)$ (our notation follows [7, Sect. 5.4]). The gap probability $A(\theta)$ is given by the Fredholm determinant

$$A(\theta) = \det(I - K_N), \tag{2.7}$$

with kernel K_N acting on $L^2(-\theta, \theta)$. We compute this Fredholm determinant using a numerical procedure introduced by Bornemann [6].

There are some numerical subtleties in these calculations. While equation (1.6) provides a complete prescription for the exact distribution of all correlation functions once V is given, in order to compute F_N we must determine π_k, c_k and K_N accurately. This is a delicate problem, since c_k grows rapidly. For the choices of g and q tested here, we used the orthogonal polynomial generator developed by Gautschi along with a standardized C++ implementation developed by Yang *et al.* [13, 24]. We find that K_N can be computed accurately only for $N \leq 30$. However, as N increases, both the density $K_N(\lambda, \lambda)$ and distribution function $F_N(\lambda)$ converge exponentially fast to the density $\rho(\lambda)$ and distribution $F_{\infty}(\lambda)$ respectively, and the convergence is uniform in λ . The convergence of densities $K_N(\lambda, \lambda)$ to ρ is illustrated in Fig. 1 for different potentials. For example, the L^{∞} distance between the finite and limiting distributions shown for $N = 30$ is only $1.024e-3$. For the same reason,

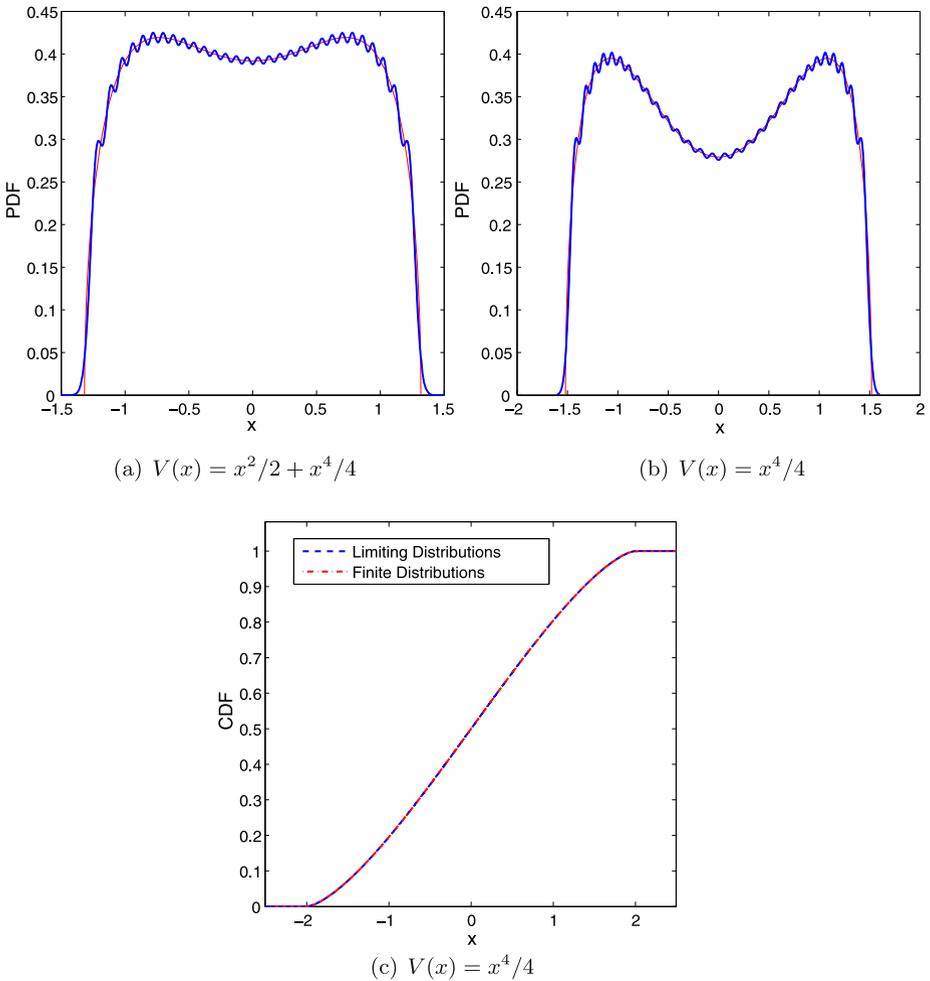


Fig. 1 Comparison of finite laws for $N = 30$ with limiting laws. **(a, b)** Comparison of probability density functions. **(c)** Cumulative-distribution functions

we restrict our comparison of $A(\theta)$ with the empirical measure to $N \leq 30$ (Note that our sampling scheme is not restricted to $N \leq 30$: the point is that it is numerically subtle to compute *exact* formulas in random matrix theory! See [5, 6] for a lively discussion of this issue).

2.3 Results

We first present a visual comparison of the empirical distributions for a single trial with a large matrix ($N = 300$). In these experiments, the time T is $O(N)$. This is a much longer time scale than required, but it serves to establish that the scheme is stable. The limiting spectral distributions F_∞ is compared with the empirical distribution for two different potentials in Fig. 2.

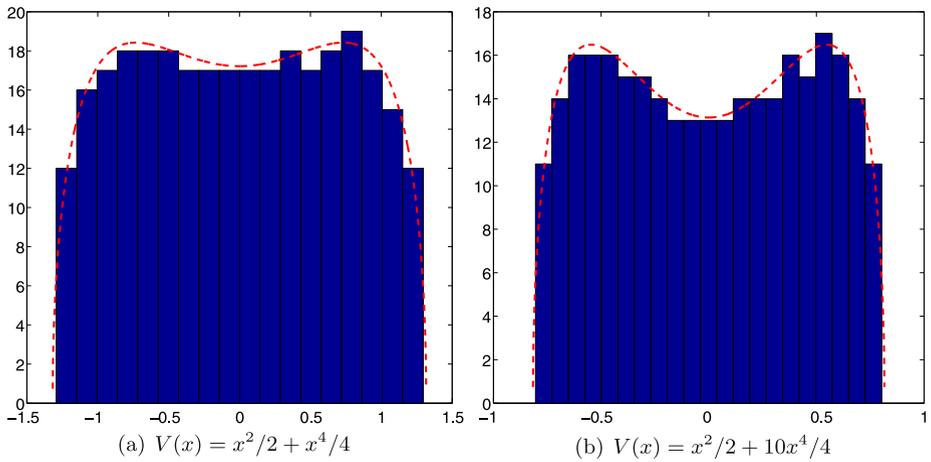


Fig. 2 Histogram of empirical distribution after one long run. A histogram of the empirical distribution of λ_k is compared with the theoretical limiting law (*dashed line*). The system size is $N = 300$, the maximal time is $T = 1200$, the time step is $\Delta t = 1/N^2$

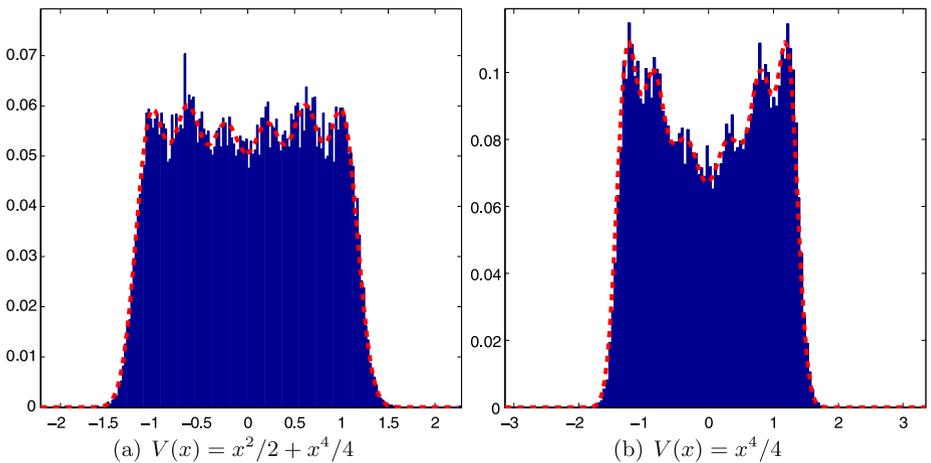
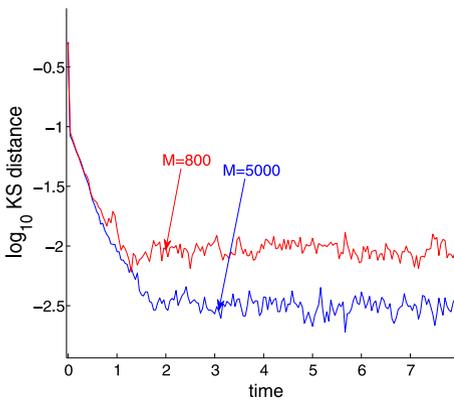
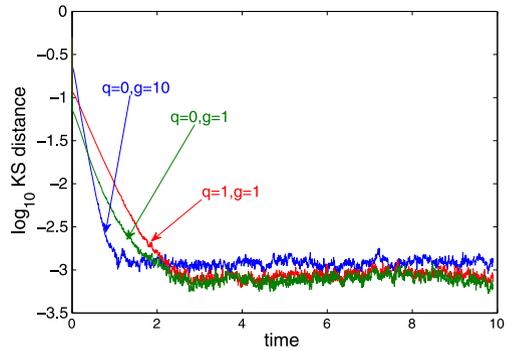


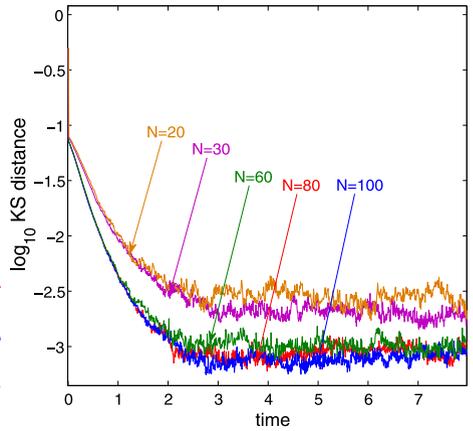
Fig. 3 Oscillations for small N . A histograms of the empirical distribution of λ_k is compared with the exact law for two different quartic potentials. The number of particles is $N = 6$, the maximal time is $T = 24$, the time step is $\Delta t = 1/N^4$ and $M = 5000$ trials are performed

For small N the density $\rho_N(s)$ oscillates and it is interesting to test if the scheme is sensitive to these oscillations. Of course, for small N it is necessary to use a large number of trials to obtain a large enough data set. Figure 3 presents the results of such a numerical experiment. We see that the scheme captures the oscillations of the density quite well. We measured the rate of convergence to equilibrium in the KS distance (2.6) for many different ensembles (i.e. potentials V), systems sizes N , number of samples M , and time steps Δt . A representative sample of the results is provided in Figs. 4, 5, and 6. We draw the reader’s attention to the robust exponential decay to the equilibrium measure with a rate of $O(1)$ in

Fig. 4 Exponential decay of the KS distance. The KS distance defined in (2.6) is plotted as a function of time for several ensembles. In these simulations, $N = 100$, the maximal time is $T = 10$, the time step is $\Delta t = 1/2N^2$, and $M = 1000$ trials were performed. Exponential decay to the equilibrium distribution is seen in a time of $O(1)$



(a) Variation with M



(b) Variation with N

Fig. 5 Exponential decay of the KS distance for the empirical measure. The KS distance defined in (2.6) is plotted as a function of time for the ensemble with $V(x) = x^4/4$ for different values of N . (a) $N = 6$, maximal time $T = 8$ and time step $\Delta t = 1/N^4$. (b) Maximal time $T = 8$, time step $\Delta t = 1/4N^2$ and $M = 1000$ trials

all these figures. Exponential convergence to equilibrium with a rate of $O(1)$ is also seen for the gap probability $A(\theta)$ in Fig. 6.

3 Conclusion

We have used a tamed explicit Euler scheme for Dyson Brownian motion to sample the equilibrium measure of the Coulomb gas on the line with a non-Gaussian potential V . Despite the singular nature of the Coulomb interaction and the nonlinear growth of the confining potential V , our scheme was always observed to be stable and convergent for time steps of $O(1/N^2)$, and to reach equilibrium in a time of $O(1)$. At present, these rates refer only to our numerical experiments with quartic potentials, but the stable performance in this case, as well as the theoretical bound (1.9) suggest that these rates may hold for more general convex potentials. The simulation results agree with the exact laws for both finite and limiting cases. The KS distances between the empirical and exact distributions decay exponentially with respect to simulation time. We view this as a demonstration of the practical utility of tamed

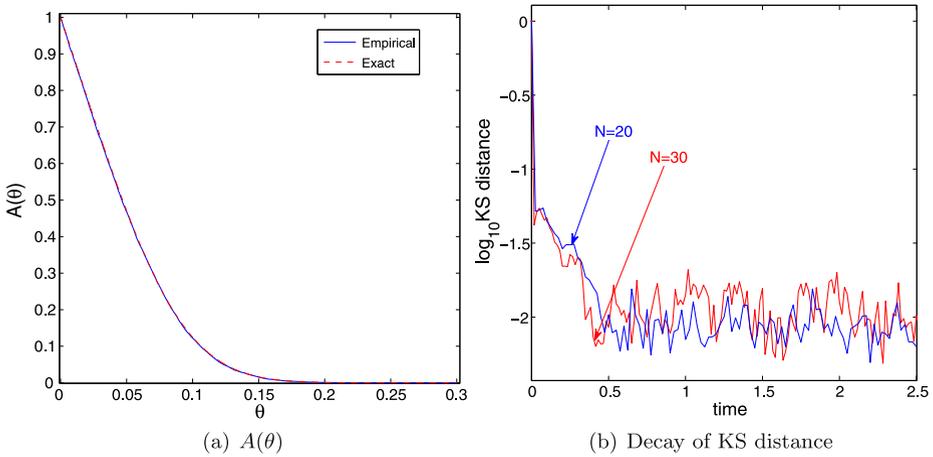


Fig. 6 Exponential decay of the KS distance for gap probabilities. $A(\theta)$ for the quartic potential $V(x) = x^4/4$. **(a)** System size $N = 20$, maximal time $T = 2.5$ and time step $\Delta t = 1/4N^2$. **(b)** Maximal time $T = 2.5$, time step $\Delta t = 1/4N^2$ and $M = 10000$ trials

explicit Euler schemes in situations where theoretical convergence results are not known to apply.

Our method can be combined with sampling techniques for matrices from the unitary group to sample a random matrix from the invariant ensemble with potential V . We find that the cost of M trials is $O(MN^3)$, though further acceleration may be provided by the fast local relaxation of the Dyson Brownian motion. To the best of our knowledge, our work and the simultaneous work of Olver et al. [22], present the first systematic sampling schemes for invariant ensembles.

Acknowledgements We thank Percy Deift for several stimulating discussions on numerical computations in random matrix theory. We thank Xiu Yang for assistance with the methods of [24]. This work has been supported in part by NSF grant DMS 07-48482.

References

1. Anderson, G.W., Guionnet, A., Zeitouni, O.: An Introduction to Random Matrices. Cambridge Studies in Advanced Mathematics, vol. 118. Cambridge University Press, Cambridge (2010)
2. Beatson, R., Greengard, L.: A short course on fast multipole methods. In: Wavelets, Multilevel Methods and Elliptic PDEs, Leicester, 1996. Numer. Math. Sci. Comput., pp. 1–37. Oxford Univ. Press, New York (1997)
3. Bessis, D., Itzykson, C., Zuber, J.B.: Quantum field theory techniques in graphical enumeration. Adv. Appl. Math. **1**(2), 109–157 (1980)
4. Bleher, P., Its, A.: Semiclassical asymptotics of orthogonal polynomials, Riemann-Hilbert problem, and universality in the matrix model. Ann. Math. **150**(1), 185–266 (1999)
5. Bornemann, F.: On the numerical evaluation of distributions in random matrix theory: a review. Markov Process. Relat. Fields **16**(4), 803–866 (2010)
6. Bornemann, F.: On the numerical evaluation of Fredholm determinants. Math. Comput. **79**(270), 871–915 (2010)
7. Deift, P.: Orthogonal Polynomials and Random Matrices: A Riemann-Hilbert Approach. Courant Lecture Notes in Mathematics, vol. 3. New York University Courant Institute of Mathematical Sciences, New York (1999)
8. Deift, P., Gioev, D.: Random Matrix Theory: Invariant Ensembles and Universality. Courant Lecture Notes in Mathematics, vol. 18. Courant Institute of Mathematical Sciences, New York (2009)

9. Duchon, P., Flajolet, P., Louchard, G., Schaeffer, G.: Boltzmann samplers for the random generation of combinatorial structures. *Comb. Probab. Comput.* **13**(4–5), 577–625 (2004)
10. Dumitriu, I., Edelman, A.: Matrix models for beta ensembles. *J. Math. Phys.* **43**(11), 5830–5847 (2002)
11. Dyson, F.J.: A Brownian-motion model for the eigenvalues of a random matrix. *J. Math. Phys.* **3**, 1191–1198 (1962)
12. Erdős, L., Yau, H.-T.: Universality of local spectral statistics of random matrices. *Bull. Am. Math. Soc.* **49**(3), 377–414 (2012)
13. Gautschi, W.: Orthogonal polynomials: applications and computation. *Acta Numer.* **5**, 45–119 (1996)
14. Higham, D.J.: An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Rev.* **43**(3), 525–546 (2001) (electronic)
15. Hough, J.B., Krishnapur, M., Peres, Y., Virág, B.: Zeros of Gaussian Analytic Functions and Determinantal Point Processes. University Lecture Series, vol. 51. American Mathematical Society, Providence (2009)
16. Hutzenthaler, M., Jentzen, A., Kloeden, P.E.: Strong and weak divergence in finite time of Euler’s method for stochastic differential equations with non-globally Lipschitz continuous coefficients. *Proc. R. Soc., Math. Phys. Eng. Sci.* **467**(2130), 1563–1576 (2011)
17. Hutzenthaler, M., Jentzen, A., Kloeden, P.E.: Strong convergence of an explicit numerical method for SDEs with nonglobally Lipschitz continuous coefficients. *Ann. Appl. Probab.* **22**(4), 1611–1641 (2012)
18. Kloeden, P.E., Platen, E.: Numerical Solution of Stochastic Differential Equations. Applications of Mathematics (New York), vol. 23. Springer, Berlin (1992)
19. Mehta, M.L.: Random Matrices, 2nd edn. Academic Press, Boston (1991)
20. Mezzadri, F.: How to generate random matrices from the classical compact groups. *Not. Am. Math. Soc.* **54**(5), 592–604 (2007)
21. Pfrang, C.W., Deift, P., Menon, G.: How long does it take to compute the eigenvalues of a random symmetric matrix? [arXiv:1203.4635](https://arxiv.org/abs/1203.4635) (2012)
22. Rao, N.R., Olver, S., Trogdon, T.: Sampling invariant ensembles (2013, in preparation)
23. Stewart, G.W.: The efficient generation of random orthogonal matrices with an application to condition estimators. *SIAM J. Numer. Anal.* **17**(3), 403–409 (1980)
24. Yang, X., Choi, M., Lin, G., Karniadakis, G.E.: Adaptive ANOVA decomposition of stochastic incompressible and compressible flows. *J. Comput. Phys.* **231**(4), 1587–1614 (2012)