Citation for published version:

Johnson, TH, Clark, SR & Jaksch, D 2010, 'Dynamical simulations of classical stochastic systems using matrix product states', *Physical Review E*, vol. 82, 036702. <https://doi.org/10.1103/PhysRevE.82.036702>

DOI:

[10.1103/PhysRevE.82.036702](https://doi.org/10.1103/PhysRevE.82.036702)

Publication date:

2010

Document Version

Early version, also known as pre-print

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Dynamical simulations of classical stochastic systems using matrix product states

T. H. Johnson,^{1,*} S. R. Clark,^{2,1,3} and D. Jaksch^{1,2,3}

¹*Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom*

²*Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543*

³*Keble College, University of Oxford, Parks Road, Oxford OX1 3PG, United Kingdom*

(Dated: June 15, 2010)

We adapt the time-evolving block decimation (TEBD) algorithm, originally devised to simulate the dynamics of 1D quantum systems, to simulate the time-evolution of non-equilibrium stochastic systems. We describe this method in detail; a system's probability distribution is represented by a matrix product state (MPS) of finite dimension and then its time-evolution is efficiently simulated by repeatedly updating and approximately re-factorizing this representation. We examine the use of MPS as an approximation method, looking at parallels between the interpretations of applying it to quantum state vectors and probability distributions. In the context of stochastic systems we consider two types of factorization for use in the TEBD algorithm: non-negative matrix factorization (NMF), which ensures that the approximate probability distribution is manifestly non-negative, and the singular value decomposition (SVD). Comparing these factorizations we find the accuracy of the SVD to be substantially greater than current NMF algorithms. We then apply TEBD to simulate the totally asymmetric simple exclusion process (TASEP) for systems of up to hundreds of lattice sites in size. Using exact analytic results for the TASEP steady state, we find that TEBD reproduces this state such that the error in calculating expectation values can be made negligible, even when severely compressing the description of the system by restricting the dimension of the MPS to be very small. Out of the steady state we show for specific observables that expectation values converge as the dimension of the MPS is increased to a moderate size.

PACS numbers: 02.50.-r, 02.70.-c, 03.67.Mn

I. INTRODUCTION

Non-equilibrium stochastic systems have attracted interest for two key reasons. Firstly, they exhibit a wealth of non-trivial phenomena, not observed in equilibrium systems, such as boundary-induced phase transitions [1] and infinite range correlations [2]. Secondly, systems far away from equilibrium have applications in describing a variety of driven diffusive processes, including vehicular traffic flow [3–5] and biological transport [6–9]. There exists a range of analytical approaches with which to address such systems, including matrix product methods [10, 11] and the Bethe ansatz [12]. However, there is no complete analytical framework for non-equilibrium systems, especially away from the steady state, which increases the importance of numerical techniques.

To simulate these systems, we adapt the time-evolving block decimation (TEBD) algorithm, a well established numerical technique for studying the time-evolution of quantum systems [13, 14]. TEBD has been used to compute the dynamics of spin chains [14–16] and the Bose Hubbard model [17, 18], which describes an array of 1D quantum systems including cold atoms in optical lattices [19]. The algorithm provides access to various dynamical properties like atomic many-body currents [20], non-linear responses to driving fields [21] and the formation and transport of quasi-particles [22, 23]. The adaption we present here simulates the time-evolution of non-equilibrium stochastic systems, and we refer to it as classical TEBD (cTEBD).

The principle method used to simulate non-equilibrium sys-

tems is dynamic Monte Carlo (DMC), first introduced by [24], and two recent examples can be found in [25, 26]. cTEBD provides a means of calculating expectation values in a fundamentally different way to DMC simulations. The latter averages over observable values calculated from many trajectories, each a possible evolution of the system, while here we evolve an approximation to the probability distribution of the system from which observable values are calculated. The motivation is to provide a complimentary method to DMC with a contrasting source of error; rather than insufficient sampling we deal with errors arising from approximating a probability distribution by a matrix product state (MPS) [27–30].

Related to TEBD and MPS, a family of numerical techniques originally devised to calculate ground state properties of quantum systems, and later applied to classical non-equilibrium systems, is the density matrix renormalisation group (DMRG) methods [31–33]. The application of DMRG methods to stochastic non-equilibrium systems in [34–36] suggest that probability distributions of many systems may be approximated well by MPS of small dimension. However, DMRG methods may only be applied to steady states of stochastic systems, not general dynamics which we address in this paper. Also, as shown in [35] for the reaction-diffusion model, they face numerical instabilities for large systems, arising from non-Hermitian diagonalization.

A key premise that motivates us to adapt the quantum TEBD algorithm is that a large class of Markovian stochastic systems evolve according to a Schrödinger-like equation [37, 38]. In this case the evolution is in imaginary time and governed by a non-Hermitian stochastic Hamiltonian. We demonstrate how this similarity between quantum and stochastic evolution can be exploited by applying cTEBD to the totally asymmetric simple exclusion process

*Electronic address: t.johnson1@physics.ox.ac.uk

FIG. 1: (Color online) The TASEP is a paradigm of stochastic non-equilibrium systems. Three Poisson processes contribute to the dynamics of the system. Particles are injected into the first site with rate α if it is empty and ejected from the N -th site with rate β if occupied. Finally, particles in the ℓ -th site hop to the $(\ell + 1)$ -th site with a rate γ if that site is empty, resulting in flow from left to right.

(TASEP) [39, 40]. The TASEP, shown schematically in Fig. 1, is the quintessential driven particle hopping system. Such systems share the following features: particles occupy a 1D lattice; the boundary sites are connected to particle reservoirs, so particles are injected and ejected there; and a hopping process causes particles to move through the lattice (see [41, 42] for reviews). Simple particle hopping systems have drawn much attention, like the Ising and Heisenberg models of equilibrium statistical physics [43], because they are archetypal of non-equilibrium physics. Importantly, the TASEP has a Hamiltonian that consists of single-site and two-site nearest-neighbor terms only [10, 44]. Using this, we will outline how the cTEBD algorithm approximately time-evolves a probability distribution by the repeated application of two-site nearest-neighbor operators, each of which affects only the matrices of the MPS associated with those two sites. After each two-site operation the algorithm ensures that the description of the system remains compressed by using a matrix factorization, returning it to an MPS. We have adapted the TEBD to allow for a range of factorizations to be used as part of the algorithm, and choosing which factorization is the key element in determining the computational cost and accuracy of cTEBD.

The two candidates we focus on are the singular value decomposition (SVD) [45] and the non-negative matrix factorization (NMF) [46, 47], which by restricting the matrices of the MPS to be non-negative (called an sMPS) allows for an information theoretic interpretation, as recently investigated in [48]. We compare the performance of these two factorizations and find that while SVD algorithms are well developed, stable and efficient, current NMF algorithms are unable to identify accurate factorizations, due to the non-convexity of the task.

As a result we focus on the SVD-based cTEBD algorithm and investigate its performance in simulating the TASEP. We look in detail at the two main sources of error: the Trotter error, which is due to approximating the evolution over each time-step by a sequence of two-site operations, and the truncation error, due to compressing the state to an MPS of small dimension. For small systems we show how these errors behave, as well as some bounds on them, and then analyze how these results scale with the system size up to hundreds of lattice sites. The results obtained promise that cTEBD could be applicable to and accurate in simulating a wide range of non-

equilibrium systems.

The structure of the paper is as follows: We review MPS, their construction and physical interpretation in Sec. (II). Next, in Sec. (III) we show how MPS can be evolved efficiently and discuss the properties of two factorization methods for use in cTEBD, namely NMF and SVD. The TASEP and its mathematical representation are introduced in Sec. (IV). In Sec. (V) we compare the performance of cTEBD with NMF and SVD, before focusing on the errors of the SVD-based algorithm for small systems. We then go on to apply cTEBD to the study of larger systems in Sec. (VI) before presenting our conclusions and outlook for future work in Sec. (VII).

II. MATRIX PRODUCT STATES

Matrix product states are essentially a special structure for approximating vectors in extremely high dimension vector spaces, which typically arise from the tensor product of smaller vector spaces. Since our use of MPS for stochastic systems is motivated by and in many ways analogous to their use in 1D quantum systems, we begin by introducing MPS first for quantum state vectors [13, 49, 50] and then in the context of probability distributions.

A. Quantum systems

To illustrate MPS in their most general form we consider a quantum system composed of N sites each with a local d -dimensional Hilbert space. An arbitrary quantum state of this system $|\psi\rangle$, normalized according to the L_2 -norm as $\| |\psi\rangle \|_2 = 1$, can be expressed as

$$|\psi\rangle = \sum_{\mathbf{i}} \psi_{\mathbf{i}} |\mathbf{i}\rangle,$$

where $\mathbf{i} = (i_1, i_2, \dots, i_N)$ is an N -tuple of local indices $i_\ell = 0, \dots, d-1$ specifying a complete configuration and $\{|\mathbf{i}\rangle = |i_1\rangle |i_2\rangle \dots |i_N\rangle\}$ are the set of orthonormal configuration states. Since $|\psi\rangle$ inhabits the tensor product vector space $(\mathbb{C}^d)^{\otimes N}$, whose dimension grows exponentially with N as d^N , an exact description of its amplitudes $\psi_{\mathbf{i}}$ rapidly becomes intractable with increasing N . To attempt to overcome this scaling, we expand $\psi_{\mathbf{i}}$ as

$$\psi_{\mathbf{i}} = A^{[1]i_1} A^{[2]i_2} \dots A^{[N]i_N}, \quad (1)$$

where for each i_ℓ the factor $A^{[\ell]i_\ell}$ is a complex matrix $\mathbb{C}^{\chi_{\ell-1} \times \chi_\ell}$, aside from the boundary terms $A^{[1]i_1}$ and $A^{[N]i_N}$ which are instead complex vectors $\mathbb{C}^{1 \times \chi_1}$ and $\mathbb{C}^{\chi_{N-1} \times 1}$, respectively [71]. The expansion in Eq. (1) is simple to visualize; each site ℓ , depending on its local configuration i_ℓ , contributes a matrix $A^{[\ell]i_\ell}$ to a site-ordered multiplication of matrices which results in the amplitude $\psi_{\mathbf{i}}$. The boundary vectors then ensure that a scalar is recovered. Such an MPS can be thought of as a tensor network, as depicted in Fig. 2. A product state, whose MPS requires dimensions $\chi_\ell = 1$ for all ℓ

FIG. 2: (Color online) An MPS for a vector $|\psi\rangle$ as a comb-like tensor network. Each site ℓ of the system has a tensor $A^{[\ell]}$ associated with it, illustrated by the shaded circles. The thick vertical leg of a tensor $A^{[\ell]}$ represents the physical index i_ℓ , while the thinner leg(s) portray the internal index (indices). The boundary tensors have a single internal leg and the others have two. The joining of tensor legs represents the contraction (or matrix multiplication) over the internal indices.

has no correlations between different sites, while the MPS describing a vector with arbitrarily strong correlations may need matrix dimensions χ_ℓ up to $X_\ell \equiv \min(d^\ell, d^{N-\ell})$, and so scale exponentially with the system size N . Thus in describing arbitrary vectors an MPS description offers no advantage.

The utility of an MPS formulation, however, is based precisely on the fact that physical states have a structure which allows for a near exact MPS description while restricting $\chi_\ell \leq \chi$ with χ much smaller than the largest X_ℓ . For quantum systems, correlations in a many-body state $|\psi\rangle$ are quantified by the entanglement entropy and this in turn has a rigorous connection to the required MPS dimension χ [13, 51]. This is shown by first splitting the system into two contiguous blocks $L^{[\ell]} = \{1, \dots, \ell\}$ and $R^{[\ell]} = \{\ell + 1, \dots, N\}$. The state is then expanded as $|\psi\rangle = \sum_{\mathbf{i}\mathbf{j}} \psi_{\mathbf{i}\mathbf{j}} |\mathbf{i}\rangle_L |\mathbf{j}\rangle_R$, where $\mathbf{i}$ and $\mathbf{j}$ label configurations of the two subsystems and ψ is now the reshaped $d^\ell \times d^{N-\ell}$ matrix of amplitudes. To expose the correlations of this state an SVD [45] is performed which factorizes $\psi = UDV^\dagger$, where U and V are unitary matrices and D is a diagonal matrix with real non-negative elements. This establishes a so-called Schmidt decomposition [52] of $|\psi\rangle$ as

$$|\psi\rangle = \sum_{\mu=1}^{X_\ell} \lambda_\mu^{[\ell]} |L_\mu^{[\ell]}\rangle |R_\mu^{[\ell]}\rangle. \quad (2)$$

Here $\{|L_\mu^{[\ell]}\rangle\}$ and $\{|R_\mu^{[\ell]}\rangle\}$ are Schmidt states derived from the columns of U and V , while $\{\lambda_\mu^{[\ell]}\}$ are the singular values (Schmidt coefficients in this context) specified by the diagonal elements of D , arranged in non-increasing order with μ . The normalization of $|\psi\rangle$ and unitarity of U and V result in

$$\sum_{\mu} (\lambda_\mu^{[\ell]})^2 = 1 \quad \text{and} \quad \langle L_\nu^{[\ell]} | L_\mu^{[\ell]} \rangle = \langle R_\nu^{[\ell]} | R_\mu^{[\ell]} \rangle = \delta_{\nu\mu}.$$

The mutual information between the subsystems $L^{[\ell]}$ and $R^{[\ell]}$ is equal to twice the entanglement entropy [53], given by

$$S_{LR} = - \sum_{\mu=1}^{X_\ell} (\lambda_\mu^{[\ell]})^2 \log [(\lambda_\mu^{[\ell]})^2].$$

The size of the entanglement entropy between a subsystem and the rest of the system for ground states has been extensively studied in the context of area laws (see [54] for a review). These have shown that S_{LR} should scale with the

boundary connecting the two parts up to logarithmic corrections at criticality, and area laws therefore severely limit the entanglement entropy physical systems can possess. For quantum spin chains with short-range interactions the area law is manifested in an extremely rapid decay of the Schmidt coefficients $\{\lambda_\mu^{[\ell]}\}$ [14]. An immediate consequence of this is that an accurate (un-normalized) approximation $|\tilde{\psi}^{[\ell]}\rangle$ of $|\psi\rangle$ can be formed by truncating the sum in Eq. (2) to the first χ terms [14] as

$$|\tilde{\psi}^{[\ell]}\rangle = \sum_{\mu=1}^{\chi} \lambda_\mu^{[\ell]} |L_\mu^{[\ell]}\rangle |R_\mu^{[\ell]}\rangle.$$

The error made in computing the expectation value of any quantum mechanical observable with $|\tilde{\psi}^{[\ell]}\rangle$ rather than $|\psi\rangle$ is bounded by the L_2 -norm of the residual [52], given by $\| |\psi\rangle - |\tilde{\psi}^{[\ell]}\rangle \|_2 = [\sum_{\mu=\chi+1}^{X_\ell} (\lambda_\mu^{[\ell]})^2]^{1/2}$ and is therefore small if the sum of the squares of the truncated Schmidt coefficients is small [14]. These observations can be directly related to an MPS, as described in Appendix A, by applying this decomposition and truncation in an iterative sequence to all contiguous bipartitions to form an approximation $|\tilde{\psi}\rangle$ [13]. This approximation is an MPS of the form of Eq. (1) but with a maximum matrix dimension χ . The errors in this case are bounded by the singular values [51] as

$$\| |\psi\rangle - |\tilde{\psi}\rangle \|_2 \leq \left[2 \sum_{\ell=1}^{N-1} \sum_{\mu=\chi+1}^{X_\ell} (\lambda_\mu^{[\ell]})^2 \right]^{1/2}. \quad (3)$$

A rapid decay of $\{\lambda_\mu^{[\ell]}\}$ with μ for each bipartition thus enables χ to be kept small while retaining a high accuracy, and results in an MPS description parameterized by $O(Nd\chi^2)$ complex numbers [14]. This near lossless compression of the information in $\{\psi_i\}$ for physically relevant vectors is responsible for the enormous success of the MPS approach used within DMRG [31, 32] and TEBD methods [13, 14].

B. Classical systems

The general MPS formalism introduced can equally be applied to stochastic classical systems. We now consider a system of N sites each with d local configurations. The state of the system is then described by a probability vector $|P\rangle = \sum_i P_i |\mathbf{i}\rangle$ with d^N real non-negative components $\{P_i\}$ corresponding to the probabilities of the system being in each of the configurations $\{\mathbf{i}\}$. Classical states $|P\rangle$ are therefore contained in the positive orthant of the vector space $(\mathbb{R}^d)^{\otimes N}$ and are normalized according to the L_1 -norm as $\| |P\rangle \|_1 = 1$. An immediate restriction to real matrices $\{A^{[\ell]i_\ell}\}$ within the MPS expansion in Eq. (1) can be made in this case. A major question which we will address in this work is whether it is necessary, or indeed desirable, to further restrict the matrices to be non-negative. To do so manifestly ensures that the sMPS approximation to $|P\rangle$ never leaves the positive orthant. Early exact analytical calculations based on an MPS formalism were sMPS [10].

For classical systems a similar link between the dimension χ required and correlations in $|P\rangle$ can be established by working exclusively with non-negative quantities. As before, we bipartition the system after site ℓ into two contiguous blocks $L^{[\ell]}$ and $R^{[\ell]}$ and then expand $|P\rangle = \sum_{\mathbf{ij}} P_{\mathbf{ij}} |\mathbf{i}\rangle_L |\mathbf{j}\rangle_R$. Following Temme and Verstraete in [48], we may draw a close analogy to the Schmidt form of a quantum state via a decomposition of P as

$$P_{\mathbf{ij}} = \sum_{\mu=1}^{X_\ell} p_\mu^{[\ell]} P_L^{[\ell]}(\mathbf{i}|\mu) P_R^{[\ell]}(\mathbf{j}|\mu), \quad (4)$$

which has the elegant interpretation of having $\{p_\mu^{[\ell]}\}$ as probabilities, again arranged in non-increasing order with μ , which mix together the conditional marginal probability distributions $P_L^{[\ell]}$ and $P_R^{[\ell]}$. This could equally be expressed as the vector decomposition

$$|P\rangle = \sum_{\mu=1}^{X_\ell} p_\mu^{[\ell]} |L_\mu^{[\ell]}\rangle |R_\mu^{[\ell]}\rangle, \quad (5)$$

where $|L_\mu^{[\ell]}\rangle$ and $|R_\mu^{[\ell]}\rangle$ have elements $\{P_L^{[\ell]}(\mathbf{i}|\mu)\}$ and $\{P_R^{[\ell]}(\mathbf{j}|\mu)\}$ respectively, similar in form to the Schmidt decomposition in Eq. (2) except that $|P\rangle$ is expanded in terms of non-negative vectors rather than orthonormal ones. In this case the mutual information is upper-bounded by the entropy of the mixing probability distribution [48], given by

$$S(\{p_\mu^{[\ell]}\}) = - \sum_{\mu=1}^{X_\ell} p_\mu^{[\ell]} \log [p_\mu^{[\ell]}].$$

Akin to the truncation of a Schmidt decomposition, an (un-normalized) approximation $|\tilde{P}^{[\ell]}\rangle$ to $|P\rangle$ can be formed by truncating the decomposition in Eq. (4) to the first χ terms, while the L_1 -norm of the residual is $\| |P\rangle - |\tilde{P}^{[\ell]}\rangle \|_1 = \sum_{\mu=\chi+1}^{X_\ell} p_\mu^{[\ell]}$. In Appendix B we show that for classical systems the error made in computing any observable is upper-bounded by the L_1 -norm error of the probability vector, analogous to the role of the L_2 -distance for quantum systems. The decomposition Eq. (4) is in general non-unique and the most significant decomposition at this bipartition is the one with the spectrum $\{p_\mu^{[\ell]}\}$ that has the smallest entropy S_C , called the entropy cost [48]. This suggests that the entropy cost S_C should be considered analogous to the entanglement entropy S_{LR} since it quantifies both the quality of approximation and correlations. However, since S_C only upper-bounds the mutual information the connection between correlations and truncation errors is substantially weaker than that seen in the quantum case. Just like for quantum systems, repeated application of the decomposition in Eq. (5) at all bipartitions can be used to build an approximate sMPS $|\tilde{P}\rangle$, and in [48] it was shown that the error is bounded by

$$\| |P\rangle - |\tilde{P}\rangle \|_1 \leq \sum_{\ell=1}^{N-1} \sum_{\mu=\chi+1}^{X_\ell} p_\mu^{[\ell]}.$$

Exact sMPS solutions for TASEP stationary states find that the entropy minimizing spectrum decays very quickly [48] thus indicating in a similar way to the quantum case that a small χ can be chosen while maintaining an accurate description.

In this work we note that each decomposition Eq. (4) is in fact precisely equivalent to an exact NMF of the matrix P , analogous to performing an SVD in the quantum case. The exact NMF can be written as $P = UDV^T$ where the columns of U and V are the normalized non-negative probability vectors in the decomposition Eq. (5) and the diagonal elements of D are the normalized probabilities $\{p_\mu^{[\ell]}\}$ [55]. As was done with the SVD, $|\tilde{P}^{[\ell]}\rangle$ is formed by truncating the inner-dimension of the NMF to χ .

While the NMF may be a more natural choice for decomposing probability vectors at a formal level, it is perfectly possible to employ SVDs on $|P\rangle$ in precisely the same way as for a quantum state. The L_2 -norm error is given by Eq. (3), and the L_2 -norm upper-bounds the more relevant L_1 -norm with a factor of the square root of the dimension of the vector space [45], hence for the approximation $|\tilde{P}\rangle$ built in this way

$$\| |P\rangle - |\tilde{P}\rangle \|_1 \leq \Lambda \equiv 2^{N/2} \left[2 \sum_{\ell=1}^{N-1} \sum_{\mu=\chi+1}^{X_\ell} (\lambda_\mu^{[\ell]})^2 \right]^{1/2}. \quad (6)$$

The SVD approach has two pleasing features: firstly, it is unique, and secondly, given $|P\rangle = \sum_{\mathbf{ij}} P_{\mathbf{ij}} |\mathbf{i}\rangle_L |\mathbf{j}\rangle_R$ and $|Q\rangle = \sum_{\mathbf{ij}} Q_{\mathbf{ij}} |\mathbf{i}\rangle_L |\mathbf{j}\rangle_R$, a well known result first proposed by Eckart and Young [56] is that $|\tilde{P}^{[\ell]}\rangle$ is the solution to the optimization problem [45]

$$\min_{\substack{Q \in \mathbb{R}^{d^\ell \times d^{N-\ell}}, \\ \text{rank}(Q) \leq \chi}} \| |P\rangle - |Q\rangle \|_2 = \| |P\rangle - |\tilde{P}^{[\ell]}\rangle \|_2 \quad (7)$$

for any $1 \leq \chi \leq X_\ell$. Thus the L_2 -norm error of a truncated SVD will always be less than that of a truncated NMF or any other factorization of the same rank. This ensures that if there exists an exact MPS of some dimension χ then there also exists an exact SVD-MPS of dimension χ or less.

A powerful feature of the SVD is the orthogonality constraint on the singular vectors (the columns of U and V). However, this typically results in the U and V matrices having elements of mixed signs, even for non-negative P , and so does not permit any obvious information theoretic interpretation of the components of the decomposition, as exists for an NMF. This also means that when $|P\rangle$ is truncated to form $|\tilde{P}^{[\ell]}\rangle$, negative probabilities may enter our description. The maximum negativity of any element of $|\tilde{P}^{[\ell]}\rangle$ is

$$\max_{\mathbf{ij}} | -\tilde{P}_{\mathbf{ij}}^{[\ell]} | \leq \| |P\rangle - |\tilde{P}^{[\ell]}\rangle \|_\infty \leq \sum_{\mu=\chi+1}^{X_\ell} \lambda_\mu^{[\ell]}.$$

Thus while non-negativity is not conserved it is controlled by a similar criterion to the accuracy of the factorization, namely the smallness of the truncated singular values. This bound is loose, for example in the worst case limit of truncation to $\chi = 1$ non-negativity is in fact preserved due to the Perron-Frobenius theorem [57] which ensures that both the left and right singular vectors associated to the largest singular value of a non-negative matrix are themselves non-negative.

III. TIME-EVOLVING BLOCK DECIMATION

We have seen that MPS can provide an accurate and efficient description of some probability distributions. Our purpose for this section is to review the TEBD algorithm, originally devised to update a quantum MPS under unitary time-evolution [13, 14, 18], but now paying specific attention to issues arising from its use on classical stochastic systems. Analogous to quantum systems, the evolution of a large class of stochastic systems, for a time t , is performed by $|P(t)\rangle = \exp(Ht)|P\rangle$ [37, 38]. The evolution in this case is stochastic and H is a stochastic Hamiltonian, the properties of which will be discussed in Sec. (IV).

A. Time-evolving the state

For anything but the smallest systems neither $|P\rangle$ nor the full stochastic evolution operator $\exp(Ht)$ can be represented exactly. Having applied an MPS to tackle the description of $|P\rangle$ we now follow the standard approach used for similar quantum evolutions [14], namely breaking it up into n small time-steps $\delta t = t/n$ as $\exp(Ht) = \prod_{j=1}^n \exp(H\delta t)$ and then splitting $\exp(H\delta t)$ into a product of two-site operators. To do the latter we first restrict ourselves to considering stochastic Hamiltonians H composed of a sum of nearest-neighbor two-site terms $\{h_{\ell,\ell+1}\}$ as

$$H = \sum_{\ell=1}^{N-1} h_{\ell,\ell+1},$$

where we have incorporated any single-site terms into $\{h_{\ell,\ell+1}\}$. The TASEP is in this restricted class [10, 44]. We can proceed to split up $\exp(H\delta t)$ via a second-order Suzuki-Trotter expansion [58] as

$$e^{H\delta t} = \left(\prod_{\ell=1}^{N-1} e^{h_{\ell,\ell+1}\delta t/2} \right) \left(\prod_{\ell=N-1}^1 e^{h_{\ell,\ell+1}\delta t/2} \right) + O(\delta t^3), \quad (8)$$

which represents an approximation because overlapping terms, e.g. $h_{\ell,\ell+1}$ and $h_{\ell+1,\ell+2}$, do not generally commute. This expansion therefore applies a sequence of two-site nearest-neighbor stochastic operators $S_\ell = \exp(h_{\ell,\ell+1}\delta t/2)$ sweeping across pairs of sites, as depicted in Fig. 3(a). A consequence of this expansion is that the approximation of $\exp(Ht)$ remains stochastic and so preserves the L_1 -norm and non-negativity of a probability vector $|P\rangle$. In Appendix C we show that for stochastic systems, the corresponding L_1 -norm error incurred by using this expansion, which we call the Trotter error, scales at worst as $\mathcal{E}_{ST} \sim t\delta t^2$.

B. Applying a two-site stochastic operator

The core of the TEBD algorithm is a procedure for applying a two-site nearest-neighbor operator to an MPS and performing a systematic truncation of the result back into an MPS of

FIG. 3: (Color online) (a) A circuit diagram showing the second-order Suzuki-Trotter expansion used in this work. One time-step $\exp(H\delta t)$ is approximated by applying two-site stochastic operators S_ℓ between each nearest-neighbor pair of sites, sweeping first from left to right and then back again. (b) A schematic tensor diagram of the core procedure in the TEBD algorithm [13, 14] for applying a two-site operator, described in the text. Firstly, in (i) the physical legs of operator S_ℓ are contracted with the relevant legs of the tensors $\{A^{[l]}\}$ in the MPS. This multiplies the state by S_ℓ and involves only the tensors shaded. The result of this is a two-site order-4 tensor Θ shown in (ii). Next in (iii) the Θ tensor is reshaped as a matrix, resulting in two so-called fat legs. The most computationally expensive step is (iv) where an approximate matrix factorization of Θ , of the form $\tilde{\Theta} = UDV^T$ is computed. Finally in (v) the reshaping of U and DV^T is equivalent to unravelling their legs to form the newly updated $A'^{[l]}$ and $A'^{[l+1]}$ tensors. The diagonal matrix D is absorbed into the V^T during the left to right sweep, as shown by the shading in (iv), and into the U for the opposite direction. The $A'^{[l]}$ tensor constructed from an orthogonal matrix only, automatically obeys one of the orthonormality constraints detailed in Appendix A. Which way the D is absorbed then follows from requiring that the Θ tensor for the next two-site operation has indices which correspond to orthonormal bases, as discussed in Sec. (III C).

dimension χ [13, 14]. Using this procedure then allows the sequence of operators S_ℓ making up the Trotterized approximation to $\exp(H\delta t)$ to be approximately applied to $|P\rangle$. Given an MPS of form $|P\rangle = \sum_{\mathbf{i}} A^{[1]i_1} A^{[2]i_2} \dots A^{[N]i_N} |\mathbf{i}\rangle$ the operator S_ℓ can be applied exactly to $|P\rangle$ and only affects the $A^{[l]}$ and $A^{[l+1]}$ tensors, as shown in step (i) of Fig. 3(b). This forms a new order-4 tensor associated to these sites with the contraction written out explicitly as

$$\Theta_{\nu\mu}^{i_\ell i_{\ell+1}} = \sum_{j_\ell, j_{\ell+1}=0}^{d-1} (S_\ell)_{j_\ell j_{\ell+1}}^{i_\ell i_{\ell+1}} \sum_{\zeta=1}^{\chi_\ell} A_{\nu\zeta}^{[l]j_\ell} A_{\zeta\mu}^{[l+1]j_{\ell+1}}.$$

The resulting state $S_\ell|P\rangle$, depicted in step (ii) of Fig. 3(b), no longer has a proper MPS form due to the presence of a structureless two-site Θ tensor. To establish an MPS form we first reshape the Θ tensor into a conventional matrix by combining the indices (i_ℓ, ν) and $(i_{\ell+1}, \mu)$ into so-called fat indices which represent the row and columns respectively, as depicted in step (iii) of Fig. 3(b). The resulting Θ matrix is then subject to a matrix factorization which approximately decomposes it into the form $\tilde{\Theta} = UDV^T$, where D is a $\mathbb{R}^{\chi \times \chi}$ diagonal matrix,

$U \in \mathbb{R}^{d \times \chi}$, $V \in \mathbb{R}^{d \times \chi}$ and χ is the desired MPS dimension. This is step (iv) in Fig. 3(b) and is the most computationally expensive part of the procedure, requiring a number of computations $O(d^3 \chi^3)$ for the SVD. After absorbing D into either the U or V^T matrix according to the direction of the Trotter sweep, the matrices are reshaped into order-3 tensors $A^{[\ell]}$ and $A'^{[\ell+1]}$. In this way, as shown in step (v) of Fig. 3(b), the factorization is used to impose internal structure on the original Θ tensor and by replacing Θ with the new tensors establishes a proper MPS for $S_\ell|p\rangle$. Crucially, however, the factorization also provides the means of truncating the MPS dimension to some desired maximum χ , controlling the growth of the description at the expense of becoming an approximation. We call the L_1 -norm error due to these approximate factorizations the truncation error $\mathcal{E}_t$ and we discuss the calculation of quantities which bound this error in Appendix C. In Sec. (II) we have studied the analytic properties of the SVD and NMF factorizations, while now we focus on the numerical properties of these factorizations and their accuracy when used in the two-site procedure described above.

1. Non-negative matrix factorization

There is no known algorithm for finding (non-trivial) exact NMF solutions for a given $m \times n$ non-negative matrix Θ , despite their existence [59], and so an NMF based approximation cannot currently proceed by truncation from an exact factorization. Instead an approximation $\tilde{\Theta} = WH^T$ is formed, where W and H are non-negative matrices, by solving directly the reduced rank minimization problem

$$\min_{\substack{W \in \mathbb{R}_+^{m \times \chi} \\ H \in \mathbb{R}_+^{n \times \chi}}} \{F(\Theta, WH^T)\},$$

where $F(\Theta, WH^T)$ is a cost function that quantifies the quality of the approximation to Θ . By forming diagonal matrices D_W and D_H containing the column sums of W and H any NMF can be brought into the form $\tilde{\Theta} = UDV^T$ where D is a diagonal matrix $D = D_W D_H$, satisfying $\text{tr}(D) = \sum_{ij} \Theta_{ij}$, and U and V are column stochastic [55]. Common choices for F are

$$F(\Theta, WH^T) = \begin{cases} \sqrt{\sum_{ij} |\Theta_{ij} - (WH^T)_{ij}|^2}, \\ \sum_{ij} \left\{ \Theta_{ij} \log \left[\frac{\Theta_{ij}}{(WH^T)_{ij}} \right] - \Theta_{ij} + (WH^T)_{ij} \right\}, \\ \sum_{ij} |\Theta_{ij} - (WH^T)_{ij}|, \end{cases} \quad (9)$$

which are the L_2 -norm of the residual reshaped as a vector $\|\Theta - \tilde{\Theta}\|_2$, the generalized Kullback-Leibler (KL) divergence [72], and the L_1 -norm $\|\Theta - \tilde{\Theta}\|_1$, respectively [60]. A crucial issue for any minimization algorithm based on these cost functions is that although they are convex for W and H separately they are not convex for W and H simultaneously, meaning the problem can have many local minima. A common strategy for minimization is to alternate the optimization between W and H while the other is fixed, however such algorithms can, at best, only guarantee convergence to a local minimum. Furthermore there is no guarantee that the global minimum, if located, is unique. The non-uniqueness of an NMF

is evident given that any $\mathbb{R}_+^{X \times X}$ non-negative monomial matrix M can also form an alternative NMF as $\tilde{\Theta} = WMM^{-1}H^T$. In Appendix D we outline some popular and simple NMF algorithms. From numerical test we have found, as expected, that the L_1 -norm minimizing NMF algorithms consistently deliver the smallest L_1 -norm errors and so we concentrate on this cost function.

2. Singular value decomposition

For any real $m \times n$ matrix Θ there exists an SVD of the form $\Theta = UDV^T$, where U and V are $\mathbb{R}^{m \times X}$ and $\mathbb{R}^{n \times X}$ orthogonal matrices respectively, and $X = \min(m, n)$ is the maximum possible rank of Θ [45]. In this context D is a $\mathbb{R}^{X \times X}$ diagonal matrix of non-negative singular values $\{\lambda_\mu\}$ arranged in non-increasing order with μ , and the columns of U and V are the singular vectors. In contrast to the NMF, the SVD is a unique decomposition up to the signs of its left and right singular vectors. These properties, as well as the Eckart-Young theorem in Eq. (7), not only make the SVD very mathematically appealing but have also aided in the construction of efficient, accurate and numerically robust algorithms for their computation [45]. The L_2 -norm minimizing rank- χ approximation $\tilde{\Theta}$ is formed by truncating the inner dimension in the SVD from X to χ .

C. Optimality of factorizations within cTEBD

For the calculation of arbitrary expectation values, we would like to minimize the distance between $|P\rangle$ and its cTEBD approximation $|Q\rangle$. Within cTEBD it is the local Θ tensor that is factorized, and discussions in the previous section explain how the algorithms will attempt to solve for the optimality of the approximation $\tilde{\Theta}$ to the Θ tensor. Here, we examine the optimality of this factorization for the distance between $|P\rangle$ and $|Q\rangle$.

A unique property of the L_2 -norm is that it is unitarily invariant and so we are free to evaluate it in any orthonormal basis. For this reason the SVD truncation $\tilde{\Theta}$ of Θ , which minimizes the corresponding local error $\|\Theta - \tilde{\Theta}\|_2$, also minimizes the global error $\|P - Q\|_2$ between the corresponding vectors, so long as the indices of $\tilde{\Theta}$ and Θ correspond to an orthonormal basis. In Appendix A we discuss the orthonormality constraints that the tensors $\{A^{[\ell]}\}$ need to obey for this to be the case. In the context of classical systems, where inherently non-unitary local stochastic operators are applied [73], maintaining such orthonormal properties of the tensors $\{A^{[\ell]}\}$ is by no means guaranteed. Yet it turns out that by applying the two-site operators in a sweeping sequence across the system the relevant bases for each pair are always orthonormal and the SVD truncation is optimal for each operation [61]. This is described in detail in Fig. 4. The second-order Suzuki-Trotter expansion introduced earlier is based precisely on such a sweep and is therefore an ideal choice for implementing generic non-unitary evolution on an MPS if the L_2 -distance is the cost function to be minimized.

FIG. 4: (Color online) Following [61] we illustrate the orthonormality structure of an MPS. Here the boxes correspond to the sets of basis states for the left and right subsystems for different bipartitions (see Appendix A for details). If these states form an orthonormal basis then the box is shaded, and correspondingly it is hashed if they do not. (a) This is the situation obtained by following the procedure described in Fig. 3 up until $S_{\ell-1}$ has just been applied and S_ℓ is next to be applied. All splittings to the left of the one between sites $(\ell-1, \ell)$ have left states which are orthonormal, while all those to the right have their right states orthonormal. The orthogonality centre is said to be located after site $\ell-1$. If the two-site stochastic operator S_ℓ is applied to sites $(\ell, \ell+1)$ then we can be assured that all the indices of the Θ tensor are orthonormal since subsystems $L^{[\ell-1]}$ and $R^{[\ell+1]}$ have orthonormal bases. (b) This is the situation after the application of S_ℓ . While the action of this non-unitary operator destroys the orthonormality of $R^{[\ell-1]}$, the unitarity of the SVD decomposition establishes it in $L^{[\ell]}$. This results in the orthogonality centre moving one site to the right. If we applied another two-site operator to sites $(\ell+1, \ell+2)$ then we are again assured of the orthonormality of the tensor indices since subsystems $L^{[\ell]}$ and $R^{[\ell+2]}$ have orthonormal bases in (b). Finally, analogous changes in the orthonormality structure occur for two-site operations moving to the left.

For cTEBD based on sMPS and NMFs, orthogonality is less useful since in general it cannot be imposed on an sMPS due to its non-negativity, and moreover the L_1 -norm minimized by the NMF is not unitarily invariant. Despite the lack of optimality the L_1 -norm error of an NMF of the local Θ tensor $\tilde{\Theta}$ can be shown to upper-bound the full L_1 -norm error of the vectors. To see this we use the fact that the full L_1 -norm between the exact vector $|P\rangle$ and the approximate one $|Q\rangle$ simplifies in matrix product form because the sMPS of the two states differ only for the two sites $(\ell, \ell+1)$ on which an operator is applied,

$$\begin{aligned} & \| |P\rangle - |Q\rangle \|_1 \\ &= \sum_{\mathbf{i}} \left| A^{[1]i_1} \dots A^{[\ell-1]i_{\ell-1}} (\Theta^{i_\ell i_{\ell+1}} - \tilde{\Theta}^{i_\ell i_{\ell+1}}) A^{[\ell+2]i_{\ell+2}} \dots A^{[N]i_N} \right|, \\ &\leq \sum_{\mathbf{i}} A^{[1]i_1} \dots A^{[\ell-1]i_{\ell-1}} |\Theta^{i_\ell i_{\ell+1}} - \tilde{\Theta}^{i_\ell i_{\ell+1}}| A^{[\ell+2]i_{\ell+2}} \dots A^{[N]i_N}, \\ &= C^{[1]} \dots C^{[\ell-1]} \left\{ \sum_{i_\ell i_{\ell+1}} |\Theta^{i_\ell i_{\ell+1}} - \tilde{\Theta}^{i_\ell i_{\ell+1}}| \right\} C^{[\ell+2]} \dots C^{[N]}, \\ &\leq \sum_{\mu_{\ell-1} \mu_{\ell+1}} \sum_{i_\ell i_{\ell+1}} |\Theta_{\mu_{\ell-1} \mu_{\ell+1}}^{i_\ell i_{\ell+1}} - \tilde{\Theta}_{\mu_{\ell-1} \mu_{\ell+1}}^{i_\ell i_{\ell+1}}|. \end{aligned}$$

Here we have used the non-negativity of tensors $\{A^{[\ell]}\}$ along with the fact that the left $C^{[1]} \dots C^{[\ell-1]}$ and right $C^{[\ell+2]} \dots C^{[N]}$ products, where $C^{[\ell]} = \sum_{i_\ell} A^{[\ell]i_\ell}$, form a row and column vector, respectively, composed of non-negative elements ≤ 1 . The last line is then the L_1 -norm of the local tensor approxi-

mation minimized by the NMF, reshaped as a vector.

IV. THE TASEP

Having introduced cTEBD, the rest of this paper deals with analyzing the performance of cTEBD on a test system, the TASEP, which we describe in Fig. 1. For a detailed review of the TASEP we refer the reader to [39, 40], but for our purpose it is sufficient to note that the steady state of the system has been solved analytically [10], while numerical simulations are needed to calculate arbitrary expectation values away from steady state. Like many of the particle hopping systems, the TASEP is modeled as a chain of N lattice sites, each of which can be in one of d configurations, and thus can be described by the formalism introduced in Sec. (II B). For the TASEP $d = 2$ and the local configurations $i_\ell = 0, 1$ correspond to empty and occupied sites respectively. As for a large class of stochastic systems with configurations connected by Poisson processes, the probabilities of being in each complete configuration $\{P_i(t)\}$ evolve according to a master equation of the form

$$\frac{\partial P_i(t)}{\partial t} = \sum_{i' \neq i} (P_{i'}(t) R_{i' \rightarrow i} - P_i(t) R_{i \rightarrow i'}), \quad (10)$$

where $R_{i \rightarrow i'}$ is the Poisson rate governing the transition from $\mathbf{i}$ to $\mathbf{i}'$ [11]. The first term on the right hand side gives the rate of transitions into configuration $\mathbf{i}$ and the second gives the rate of leaving it.

We may rewrite such a master equation Eq. (10) as a stochastic Schrödinger equation [37, 38] of the form

$$\frac{\partial}{\partial t} |P(t)\rangle = H |P(t)\rangle,$$

where the stochastic Hamiltonian H is defined by

$$\langle \mathbf{i} | H | \mathbf{i}' \rangle = R_{i' \rightarrow i} \text{ for } \mathbf{i} \neq \mathbf{i}', \quad (11a)$$

$$\langle \mathbf{i} | H | \mathbf{i} \rangle = - \sum_{i' \neq i} R_{i \rightarrow i'}. \quad (11b)$$

The probability vector describing stochastic systems therefore undergoes non-unitary evolution via a Schrödinger equation in imaginary time with a non-Hermitian stochastic Hamiltonian. From Eqs. (11) a stochastic Hamiltonian has non-negative off-diagonal elements, a consequence of the non-negativity of transition rates, and non-positive diagonal elements to ensure the conservation of probability

$$\sum_{i'} \langle i' | H | \mathbf{i} \rangle = 0.$$

It has been shown in [62] that there always exists a stationary state, an eigenvector of H corresponding to a zero eigenvalue, and this vector is unique if H is strongly connected, i.e. all configurations are accessible, which is true of the TASEP. This implies ergodicity; after long times it will arrive at a distribution independent of its initial state. The Hamiltonian for

FIG. 5: (Color online) (a) The accuracy of approximations to the stationary state $|P\rangle$ of the 10-site TASEP with $\alpha = \beta = 1$. The L_1 -norm error is plotted against χ for the SVD ($\circ$) and L_1 -norm minimizing NMF ($\diamond$) performed at the central bipartition. For each SVD the L_1 -norm error due to the sum of all negative elements is shown ($\times$) along with the upper-bound to the L_1 -norm error (dashed line) derived from the exact L_2 -norm error. (b) Errors during the evolution of the 10-site TASEP with $\alpha = 0.25$ and $\beta = 0.75$. The NMF and SVD algorithms are compared for $\chi = 3$ and $\delta t = 10^{-2}$ where for each NMF four random restarts have been performed. Initially, the system was in the state describing a uniform distribution over all configurations. With decimation occurring at the central bipartition only, we plot the upper-bound to the L_1 -norm error incurred during each time-step, discussed in Appendix C, using the SVD (smooth dashed line) and NMF (jagged solid line), and the actual L_1 -norm error for the SVD ($\square$) and NMF ($\circ$) at time t . (c) The same as (b) but with decimation occurring at all contiguous bipartitions.

the TASEP can be written as

$$H = h_1 + h_N + \sum_{\ell=1}^{N-1} h_{\ell,\ell+1}, \quad (12)$$

where the single-site terms h_1 and h_N describe the input and output of particles and the two-site nearest-neighbor terms $\{h_{\ell,\ell+1}\}$ describe the hopping. The Hamiltonian terms can be calculated using Eqs. (11) and are given in [44]. So evolution is generated by a stochastic operator of the form considered in Sec. (III A).

For simplicity we take time t to be in units of the inverse of the hopping rate γ^{-1} and set $\gamma = 1$. Two observables that we will be particularly interested in are the density at site ℓ , ρ_ℓ , which is equal to the local configuration i_ℓ , and the current between site ℓ and $\ell + 1$, J_ℓ , equal to 1 if $i_\ell = 1, i_{\ell+1} = 0$, and 0 otherwise.

V. SMALL SYSTEM ANALYSIS

A. Comparison of the NMF and SVD for the TASEP

A physically relevant comparison of the NMF and SVD can be made by considering the probability distribution of the TASEP stationary state. The L_1 -norm errors for the SVD and NMF approximations to this probability distribution bipartitioned at centre of the system are presented in Fig. 5(a) for $\alpha = \beta = 1$ and $N = 10$. This shows that the SVD is exact up to machine precision once $\chi > 5$ indicating that the stationary state has weak enough correlations to permit a considerable amount of compression. The NMF fails to identify an accurate solution as χ is increased, due to issues with local minima, and gives an error which levels off above 10^{-3} . Note also that negativity of the approximation found from the SVD only occurs for $\chi = 2$ and is several orders of magnitude smaller

than the overall L_1 -norm error. Finally, the upper-bound to the L_1 -norm error for the SVD truncation, computed from the L_2 -norm error, is also plotted and seen to provide a reasonably tight upper-bound.

To move beyond a single matrix factorization we have also tested how the NMF and SVD behave when used repeatedly to simulate time-evolution. We applied cTEBD to the TASEP model, in which the two sets of 5 sites on either side of the central bipartition of a 10-site system were merged so only this single bipartition was considered. A calculation of the time-evolution was performed up to a time $t = N$, using both factorizations with $\chi = 3$. At this level of truncation the SVD performed roughly the same as the NMF in the stationary state, as seen in Fig. 5(a). The results for this time-evolution are shown in Fig. 5(b). Initially, the system is in a $\chi = 1$ state but the NMF fails to find accurate decompositions at small times, while the SVD succeeds. At later times the difference in errors at each time-step reduce to an order of magnitude. To examine the optimality issue explained in Sec. (III C) we have repeated the calculation, this time using a full 10-site matrix product decimation. We found that while both NMF and SVD errors are slightly worse when decimation occurs at each bipartition, this is not more serious for the NMF than the SVD algorithm, as demonstrated in Fig. 5(c). Rather than optimality, the problem faced using current NMF algorithms is the plateau in accuracy seen for increasing χ . This not only prevents any improvements, in contrast to the SVD, but causes unnecessary accumulation of error. Also to support the use of the SVD, it is found that even with $\chi = 3$ the evolution never produces any negative probabilities over this time. These results strongly suggest that the potential for negativity when using the SVD within the cTEBD framework does not present a serious obstacle. For this reason we shall, from this point on, focus on SVD-based decimation.

FIG. 6: (Color online) The errors in simulating the 15-site TASEP with $\alpha = 0.3$ and $\beta = 0.6$, using cTEBD with $\delta t = 10^{-3}$, as compared to an exact simulation. (a) Using $\chi = 5$ we plot the distances between the approximate and exact probability vectors, $\| |P\rangle - |Q\rangle \|_1$ and $\| |P\rangle - |Q\rangle \|_2$ along with the errors $|\langle O \rangle^P - \langle O \rangle^Q|$ for each observable O . The upper-bound $\mathcal{E}_{\text{tr}}^+$ is also plotted and of the solid lines this takes the largest values, with the L_1 and then L_2 -norm errors below. (b) A zoom in on the small time part of (a) along with the L_1 -norm error due to the Trotter expansion $\mathcal{E}_{\text{ST}}$ calculated using a $\chi = 100$ simulation. (c) The behavior of some of these distance measures with increasing χ at $t = N$.

B. Behavior of cTEBD errors

We now examine in detail the behavior of the different types of errors of SVD-based cTEBD for small systems, in preparation for studying the behavior of the same errors for larger systems in the next section. As a demonstration of the relationship between different error measures we show in Fig. 6(a) the L_1 and L_2 -norm errors of a simulation of the initially empty 15-site TASEP over a time $t = 20N$. Also in the figure we have plotted the expectation value errors of three observables: an observable r whose value at each configuration was randomly generated from a uniform distribution between 0 and 1; the density at the middle site ρ_8 ; and the current leaving the middle site J_8 . The random observable scales with the L_2 -norm error, because it is uncorrelated with the MPS approximation method, as expected from the discussion of observable error scaling in Appendix B. The current scales slightly above this, and the density lies even closer to the L_1 -norm error, indicating that there is some correlation between cTEBD's errors and these two observables. The L_1 -distance lies close to the maximum factor of $2^{N/2}$ above the L_2 -distance and, as must be the case, all errors are bounded by the L_1 -distance.

We may isolate the Trotter error by using a χ large enough such that the truncation error is negligible. Also, it is always possible to efficiently calculate $\mathcal{E}_{\text{tr}}^+$, introduced in Appendix C, which upper-bounds the cumulative L_1 -norm error due to truncation. The interplay between truncation and Trotter errors can be clearly seen in Fig. 6(b), which shows that the L_1 -distance follows the Trotter error, until a time at which truncation errors become more significant, as indicated by the increasing bound. $\mathcal{E}_{\text{tr}}^+$ is an overestimate of $\mathcal{E}_{\text{tr}}$ and after times of a few N it plateaus, as can be seen in Fig. 6(a). As it is cumulative, the final value of $\mathcal{E}_{\text{tr}}^+$ can be used to bound the L_1 -distance between the cTEBD and exact probability distributions, and so all errors, up to this time.

Another consequence of $\mathcal{E}_{\text{tr}}^+$ being cumulative is that it has less relevance for much longer times, when the system ap-

proaches steady state, especially since the errors do not increase monotonically in time. Consider the evolution of the TASEP shown in Fig. 6(a). Initially the system is in a specific configuration, in this case the empty lattice, which can be exactly represented by an MPS with $\chi = 1$ and so there is zero error. In time, correlations are introduced into the system and the transient errors grow to a peak located at $t \approx N$, before decreasing as the steady state is approached, suggesting that the steady state is particularly compressible. This shows that even though a low χ may not be sufficient to describe the transient behavior of the system accurately, it could still give accurate results in the steady state. This is because the TASEP is ergodic [62] and the Trotter approximation to the evolution operator in Eq. (8) appears to preserve this property. So through whichever states the low χ approximation to the transient behavior takes the system, whether this be erroneous or not, the approximation to the steady state will be the same.

For small systems we can use cTEBD to calculate any observable of the TASEP precisely, because by increasing χ , while still severely compressing the system, we can restrict the L_1 -norm error to small values. This is shown, for a time $t = N$ near the peak in errors, in Fig. 6(c). As χ is increased, truncation errors shrink until the errors are dominated by the Trotter error. Even though both the L_1 and L_2 -norm errors seem to decrease monotonically with χ , from Fig. 6(c) it is clear this is not necessarily the case for observable errors. This non-monotonic behavior arises because the correlation between errors in the probability vector and an observable's values may vary with χ .

The behavior of the errors in time, shown in Fig. 6(a), suggests that the steady state is more compressible than the transient states. To confirm this we evolved the TASEP up to $t = 100N$ by which time the system was effectively in the steady state. For this state, we show in Fig. 7(a) that errors can be driven down to the Trotter error for χ even as small as 4 when using a time-step $\delta t = 10^{-1}$. To obtain more accurate results, δt can be decreased, and the results presented in

FIG. 7: (Color online) The steady state of the TASEP. (a) For the 15-site TASEP with $\alpha = 0.3$ and $\beta = 0.6$ we have plotted the convergence of error measures to the Trotter error as χ is increased. This is shown for $\delta t = 10^{-1}$ ($\circ$) and $\delta t = 10^{-3}$ ($\diamond$). For both observables O we have plotted $|\langle O \rangle^P - \langle O \rangle^Q|$. (b) For the same system as (a) this shows the Trotter errors, calculated from the $\chi = 20$ simulations, plotted against δt on a log-log axis. To guide the eye, we have connected the points with lines. They show a δt^2 dependence. (c) The steady state singular value spectra of the 20-site TASEP, corresponding to bipartite splittings through the middle of the system. Above $\mu = 11$ the singular values are integer zero.

Fig. 7(b) confirm that the scaling of $\mathcal{E}_{ST}$ is with δt^2 . The success of the cTEBD algorithm in describing the TASEP steady states accurately with an MPS of small χ stems from the rapid decay of the singular value spectra of the bipartite splittings of the probability distribution. In [48] an exact sMPS for every N -site TASEP steady state was obtained with a dimension $\chi = N + 1$. From our discussion in Sec. (II B) we know an exact SVD-MPS must exist of a dimension less than or equal to that of an exact sMPS, and this can be constructed explicitly by repeatedly performing SVDs on the matrices of the sMPS given in [48]. By doing this we have plotted, in Fig. 7(c), the spectra of singular values for the 20-site TASEP for a few values of α and β . The spectra decay super-exponentially before becoming integer zero after $\mu = N/2 + 1$. Numerical tests suggest that this is the matrix dimension needed for an exact MPS. How these properties extend to larger systems is investigated in the next section.

VI. SCALABILITY

The calculation of the exact L_1 -distance is intractable for large systems. To overcome this we use three approaches to determine the accuracy of calculated expectation values: Firstly, in the steady state we compare expectation values calculated using cTEBD directly with exact analytical results; Secondly, away from steady state, we use $\mathcal{E}_{tr}^+$ to upper-bound the L_1 -distance; Finally, we show for specific observables that as χ is increased each expectation value converges.

With the first approach in mind, we simulated the TASEP on a 200-site lattice for a total time of $100N$, enough to reach the steady state, which, from our analysis of small systems, is when we expect the most accurate results to be obtained. In Fig. 8(a) we have plotted a comparison between cTEBD and analytical expectation values for the density at a site close to the exiting boundary of the system. Even for such a large system the error is Trotter limited for a χ of 7. This observable is actually unusual. For the vast majority of expectation values

considered, e.g. current, density and arbitrary two-site correlations, the Trotter error was reached for $\chi = 1$, signifying the extreme compressibility of these states, also responsible for the success of mean field theory on these systems. It is also worth noting that the non-monotonic behavior of errors with χ , observed for smaller systems, is present for larger systems.

To understand this compressibility of the steady states it is instructive to calculate the bound on the L_1 -norm Λ , introduced in Eq. (6), using the method for calculating singular spectra from the exact stationary sMPS [48], discussed in Sec. (V B). This gives us an upper-bound to the L_1 -distance that may be achieved by an SVD-MPS constructed in the way described in Appendix A. Although this construction is not exactly what is performed by cTEBD, the accuracy of both depend on the same properties. We have calculated Λ analytically and plotted the results in 8(b) for several system sizes. Except for some super-exponential decay when $\chi \approx N/2 + 1$ it exhibits a nearly perfect exponential decay in χ and exponential increase in N of the form $\Lambda = a^N b^{-\chi}$, with $a = 1.0770$ and $b = 11.7797$ respectively. Hence the MPS dimension needed to describe the steady state with errors less than Λ is given by $\chi = N \log_b(a) - \log_b \Lambda$, as plotted in Fig. 8(c). We find that accurate MPS approximations to the steady states of large systems can be obtained using a χ of order 10, due to the rapid decay of the singular value spectra of the steady states, in line with the cTEBD results we have obtained.

A potential limitation to the scalability of the cTEBD algorithm using the SVD is the exponential scaling with N of the prefactor bounding the L_1 -norm via the L_2 -norm, which appears in Eq. (6). This indicates that in order to maintain a given bound on the L_1 -norm the simulation must be performed with exponentially increasing L_2 -norm accuracy with N . The extremely rapid decay of the singular spectrum for TASEP stationary states indicates that an L_1 -norm bound can be maintained by only a very moderate increase in χ . However, we note that in order to exploit this property it may be necessary to use higher precision arithmetic than the standard 64-bit double precision. Indeed, for even moderately large

FIG. 8: (Color online) (a) The convergence of the cTEBD density expectation value $\langle \rho_{197} \rangle^\chi$ to the analytical result $\langle \rho_{197} \rangle$ for $N = 200$. The system is in the steady state for $\alpha = 0.349$, $\beta = 0.537$, having used $\delta t = 5 \times 10^{-3}$. (b) Λ plotted for different χ and N , from the $\alpha = 0.3$, $\beta = 0.6$ steady state, calculated using the exact results from [48]. (c) For the same system as in (b) this shows the MPS dimension χ needed to ensure an SVD-MPS exists with L_1 -distance less than Λ , as a function of N .

FIG. 9: (Color online) (a) The dependence of $\mathcal{E}_{\text{tr}}^+$ on N and χ , corresponding to the upper and lower x-axes respectively. We used $\alpha = 0.3$ and $\beta = 0.6$, taking the initial state to be the empty lattice, and setting $\delta t = 5 \times 10^{-3}$. Each simulation was ran up to $t = 10N$ by which time the value of $\mathcal{E}_{\text{tr}}^+$ had plateaued. (b) The evolution of the current expectation values, calculated with $\chi = 100$, for $N = 50$, $\alpha = 0.25$ and $\beta = 0.25$. The errors are bounded by $\mathcal{E}_{\text{tr}}^+$ which plateaus at 6×10^{-6} . (c) The errors between the expectation values of the products of the densities of three sequential sites $\{O = \rho_\ell \rho_{\ell+1} \rho_{\ell+2}\}$ and the $\chi = 100$ result, showing a convergence to within 10^{-5} as χ is increased to 50. We used $\alpha = 0.25$ and $\beta = 0.25$ and calculated expectation values at $t = N$. The time-step used was $\delta t = 5 \times 10^{-3}$, and at $t = 0$ the lattice was empty.

systems, above 50 sites or so, all but the first few singular values have a ratio to the largest that is below double machine precision. TEBD will not be able to reproduce these values. This is a similar problem to what was experienced in [35] for the Reaction-Diffusion model, but in our case numerical instabilities do not arise from this imprecision.

When away from the steady state, we cannot calculate the exact spectrum or expectation values. However, we can calculate $\mathcal{E}_{\text{tr}}^+$ and we have done this for a range of N and χ to produce Fig. 9(a). The results show that to ensure a small error, a larger χ is needed for larger N . We can also deduce, for instance, that the error in calculating the expectation value of any observable for $N = 50$ using $\chi = 75$ will not exceed 10^{-4} . In light of this we have plotted in Fig. 9(b) the time-evolved current profile of an initially empty 50-site TASEP for a time $t \approx 5N$ using a χ of 100. The system fills up as particles injected into the first site hop across the system, and steady state values are almost reached by $t = 5N$. From $\mathcal{E}_{\text{tr}}^+$ we can be sure

that all values are correct to a factor of 10^{-5} , despite evolving through the transient behavior where the largest errors are expected to appear. Note that in this argument we have ignored the Trotter errors, which we have found to be much smaller than the truncation errors, and are easily controlled by reducing δt .

Importantly, $\mathcal{E}_{\text{tr}}^+$ is an upper-bound, and so it is sufficient, but not necessary for it to take small values. As we saw in Fig. 6(a) this upper-bound is often a gross overestimation of the actual L_1 -norm error as it is the cumulative sum of upper-bounds to the L_1 -norm error in each two-site operation, as detailed in Appendix C. Motivated by this, we have used cTEBD to simulate the non-equilibrium dynamics of larger systems, for which $\mathcal{E}_{\text{tr}}^+$ alone could not guarantee accurate results. We considered how quickly expectation values converge with χ , and calculated this convergence for systems at $t = N$, which in the investigation of small systems was found to be a time close to that with the largest error. As shown in Fig. 9(c), a numerically calculated expectation value for $N = 100$ converges

such that the $\chi = 50$ and 100 results differ only by a fraction of 10^{-5} . So convergence occurs such that only a moderate χ is needed to accurately calculate the expectation value of this observable, even for transient states. This observable for which the results are plotted could, for example, be used to indicate the presence of a traffic jam with the TASEP used to model a single carriageway. Similar results were obtained for a range of observables, indicating that the cTEBD algorithm is able to accurately simulate the transient dynamics of the TASEP for these system sizes.

VII. CONCLUSIONS

In this paper we have investigated in detail the application of the TEBD algorithm to classical stochastic systems. A very reasonable approach to describe probability distributions is to use an sMPS [48] which is manifestly non-negative. We have shown that cTEBD can produce approximations to the dynamics of stochastic systems within the sMPS class by modifying only the factorization method, namely switching it to NMF. Unfortunately, current NMF algorithms find it difficult to locate accurate low rank solutions, due to the non-convexity of the factorization. However, issues of non-optimality due to the non-orthogonality of the MPS do not seem to cause problems, so if future algorithms overcome the former problem NMF-based cTEBD could become a practical option. The main approach explored here was to relax the constraint to be explicitly non-negative by using instead a general real MPS and applying the SVD method for the evolution in an essentially identical fashion to quantum problems. We showed that potential issues involving the proliferation of negativity and that the L_2 rather than L_1 -distance is minimized turn out not to present serious obstacles to the SVD-based algorithm for the system sizes considered.

With the TASEP as our test system, we demonstrated the accuracy and applicability of cTEBD for non-equilibrium systems. We focused our investigation on the behavior of errors with truncation χ and system size N . An interesting feature was the lack of a monotonic dependence of expectation value errors on χ , though in all cases considered, with system sizes up to the low hundreds, results converged. For the steady state, comparisons with analytically calculated expectation values were excellent and we obtained negligible errors for χ less than 10. Away from steady state we introduced an upper-bound to the L_1 -distance, and thus all errors, which could be efficiently calculated as part of the algorithm when simulating any system, not just the TASEP. This was used to ensure that some of our 50-site simulations had negligible errors. Even for large systems for which this bound could not guarantee small errors, we found that expectation values converged for a large range of observables.

This work suggests that cTEBD is a viable candidate for studying the time-evolution of a large class of non-equilibrium stochastic systems whose dynamics are governed by a stochastic Hamiltonian consisting of at most two-site nearest-neighbor terms. Since cTEBD is easily adjusted to include many rates with time and space dependence as well

as being extendable to network geometries with a bounded tree-width [63, 64], it could prove to be a powerful simulation option. The accuracy of cTEBD shown here also indicates that there is great scope to adapt other more sophisticated MPS/tensor network methods to classical stochastic systems and extend simulations to higher dimensions. These include the use of matrix product operators to express the evolution for longer ranged processes [65–68] and renormalization inspired variational techniques, which have already been used with great success for both stationary and dynamical calculations of quantum systems [49, 50]. Future work will investigate the performance of cTEBD and its extensions to a wider range of systems and geometries.

Acknowledgments

SRC and DJ thank the National Research Foundation and the Ministry of Education of Singapore for support. DJ acknowledges support from the ESF program EuroQUAM (EP-SRC grant EP/E041612/1), the EPSRC (UK) through the QIP IRC (GR/S82176/01), and the European Commission under the Marie Curie programme through QIPEST. TJ thanks Vlatko Vedral for useful discussions.

Appendix A: Construction of an MPS via repeated SVDs

In this section we follow Vidal [13] and show how through the repeated use of SVDs (Schmidt decompositions) an MPS of any vector $|\psi\rangle$ can, in principle, be found. To do this we first compute the Schmidt decomposition of $|\psi\rangle$ for every contiguous bipartition into blocks $L^{[\ell]}$ and $R^{[\ell]}$, as depicted in Fig. 10. Starting from the left boundary, each left Schmidt state $|L_{\mu_\ell}^{[\ell]}\rangle$ is iteratively expanded in terms of the local basis of the rightmost site of the block $|i_\ell\rangle$ and the left Schmidt states $\{|L_{\mu_{\ell-1}}^{[\ell-1]}\rangle\}$ of the neighboring splitting one site further to the

FIG. 10: (Color online) The sequence of $N - 1$ contiguous bipartitions (according to the labeling imposed) of the system in which the Schmidt decompositions are computed.

left. This gives the following set of expansions

$$\begin{aligned}
|L_{\mu_1}^{[1]}\rangle &= \sum_{i_1=0}^{d-1} A_{\mu_1}^{[1]i_1} |i_1\rangle, \\
|L_{\mu_2}^{[2]}\rangle &= \sum_{i_2=0}^{d-1} \sum_{\mu_1=1}^{X_1} A_{\mu_1\mu_2}^{[2]i_2} |L_{\mu_1}^{[1]}\rangle |i_2\rangle, \\
&\vdots \\
|L_{\mu_\ell}^{[\ell]}\rangle &= \sum_{i_\ell=0}^{d-1} \sum_{\mu_{\ell-1}=1}^{X_{\ell-1}} A_{\mu_{\ell-1}\mu_\ell}^{[\ell]i_\ell} |L_{\mu_{\ell-1}}^{[\ell-1]}\rangle |i_\ell\rangle, \\
&\vdots \\
|L_{\mu_N}^{[N]}\rangle &= |\psi\rangle = \sum_{i_N=0}^{d-1} \sum_{\mu_{N-1}=1}^{X_{N-1}} A_{\mu_{N-1}}^{[N]i_N} |L_{\mu_{N-1}}^{[N-1]}\rangle |i_N\rangle. \quad (A1)
\end{aligned}$$

The expansion coefficients can be seen to define the tensors $\{A^{[\ell]}\}$ of an MPS representation of $|\psi\rangle$ by taking the last expansion of $|\psi\rangle$ in Eq. (A1) and inserting all the others into it until no Schmidt states remain

$$|\psi\rangle = \sum_{\mathbf{i}} \left\{ \sum_{\mu_1=1}^{X_1} \sum_{\mu_2=1}^{X_2} \cdots \sum_{\mu_{N-1}=1}^{X_{N-1}} A_{\mu_1}^{[1]i_1} A_{\mu_1\mu_2}^{[2]i_2} \cdots A_{\mu_{N-1}}^{[N]i_N} \right\} |\mathbf{i}\rangle.$$

The term in the parentheses is identical to the matrix product expansion of the amplitudes $\psi_{\mathbf{i}}$ given in Eq. (1), but with all the summations written out explicitly. Truncating each summation (equivalent to truncating the Schmidt decomposition or SVD) to χ terms results in $|\tilde{\psi}\rangle$, an approximate MPS of dimension χ .

While the full product of matrices $\{A^{[\ell]i_\ell}\}$ generates $|\psi\rangle$, a partial product up to a site ℓ instead forms an MPS of a left Schmidt state $|L_{\mu_\ell}^{[\ell]}\rangle$ as

$$|L_{\mu_\ell}^{[\ell]}\rangle = \sum_{\mathbf{i}} (A^{[1]i_1} A^{[2]i_2} \cdots A^{[\ell]i_\ell})_{\mu_\ell} |\mathbf{i}\rangle_L.$$

The orthonormality of the left Schmidt states results in all the tensors $\{A^{[\ell]}\}$ derived from this construction obeying a left orthonormality property [18]

$$\sum_{i_\ell=0}^{d-1} (A^{[\ell]i_\ell})^\dagger A^{[\ell]i_\ell} = \mathbb{I}_{\chi_\ell}, \quad (A2)$$

where $\mathbb{I}_{\chi_\ell}$ is the $\chi_\ell \times \chi_\ell$ identity matrix. Notice that we could have equally performed this construction starting from the right boundary and using the right Schmidt states $\{|R_{\mu_\ell}^{[\ell]}\rangle\}$. This would result in a different set of tensors $\{A^{[\ell]}\}$ obeying a right orthonormality property

$$\sum_{i_\ell=0}^{d-1} A^{[\ell]i_\ell} (A^{[\ell]i_\ell})^\dagger = \mathbb{I}_{\chi_{\ell-1}}. \quad (A3)$$

If we were to apply the left procedure up to and including site ℓ and the right procedure up to an including site $\ell+1$ then the resulting MPS would be of the form

$$|\psi\rangle = \sum_{\mathbf{i}} A^{[1]i_1} A^{[2]i_2} \cdots A^{[\ell]i_\ell} D^{[\ell]} A^{[\ell+1]i_{\ell+1}} \cdots A^{[N]i_N} |\mathbf{i}\rangle, \quad (A4)$$

where $D^{[\ell]}$ is the diagonal matrix of singular values $\{\lambda_{\mu_\ell}^{[\ell]}\}$. This MPS has an orthogonality centre, or twist in its handedness, located precisely at the bipartition after site ℓ , since the constraints Eq. (A2) and Eq. (A3) are obeyed by the tensors to the left and right of $D^{[\ell]}$, respectively. This is crucial since we can readily extract from the MPS in Eq. (A4) an expansion of $|\psi\rangle$ in an orthonormal basis

$$|\psi\rangle = \sum_{\mu_{\ell-1}, \mu_{\ell+1}} \sum_{i_\ell, i_{\ell+1}} \psi_{\mu_{\ell-1}\mu_{\ell+1}}^{i_\ell i_{\ell+1}} |L_{\mu_{\ell-1}}^{[\ell-1]}\rangle |i_\ell\rangle |i_{\ell+1}\rangle |R_{\mu_{\ell+1}}^{[\ell+1]}\rangle,$$

by forming the two-site tensor about the orthogonality centre

$$\psi_{\mu_{\ell-1}\mu_{\ell+1}}^{i_\ell i_{\ell+1}} = \sum_{\mu_\ell} A_{\mu_{\ell-1}\mu_\ell}^{[\ell]i_\ell} \lambda_{\mu_\ell}^{[\ell]} A_{\mu_\ell\mu_{\ell+1}}^{[\ell+1]i_{\ell+1}}.$$

It is only for this pair of sites, or $(\ell-1, \ell)$ and $(\ell+1, \ell+2)$ adjacent to the orthogonality centre, that a tensor of amplitudes for $|\psi\rangle$ can be formed from components of the MPS in Eq. (A4), such that all indices correspond to orthonormal states. As discussed in Sec. (III C) moving the orthogonality centre, so as to ensure that it is always located adjacent to any pair of sites on which a two-site operator is applied, is essential for the optimality of the subsequent SVD decimation.

Appendix B: Observable errors

Here, we consider how the errors in calculating an observable using an approximate probability distribution, such as that given by a truncated MPS, should scale with the L_1 and L_2 -distances. Let $|P\rangle$ be the exact probability vector and $|Q\rangle$ be an approximate probability vector. Now consider an observable O taking a value $O_{\mathbf{i}}$ for each configuration $\mathbf{i}$. Since the expectation value is calculated as $\langle O \rangle^P = \sum_{\mathbf{i}} O_{\mathbf{i}} P_{\mathbf{i}}$, the difference between the expected values of the observable, calculated using the two probability vectors is then

$$\langle O \rangle^P - \langle O \rangle^Q = \sum_{\mathbf{i}} O_{\mathbf{i}} (P_{\mathbf{i}} - Q_{\mathbf{i}}), \quad (B1)$$

The worst case scenario for this error is that the observable values and the values of the probability error for each configuration are correlated such that each term in the sum is of the same sign. In this case the magnitude of the error can equal its upper-bound

$$\begin{aligned}
|\langle O \rangle^P - \langle O \rangle^Q| &\leq \sum_{\mathbf{i}} |O_{\mathbf{i}}| |P_{\mathbf{i}} - Q_{\mathbf{i}}|, \\
&\leq \max_{\mathbf{i}} \{O_{\mathbf{i}}\} \| |P\rangle - |Q\rangle \|_1.
\end{aligned}$$

So the error in calculating an expectation value of an observable is always bounded by the L_1 -norm error. However, for errors to scale as badly as this requires complete correlation between the approximation method and the observable. For many observables we expect no such correlation and it is instructive to consider how observable errors typically scale in this case. This typical scaling would be revealed by calculating the expected value of the magnitude of the observable

error Eq. (B1), averaging over all possible pairings of observable values O_i with errors $P_i - Q_i$. An approximation to this, which we can calculate analytically, is the average over all ways of replacing the set $\{O_i\}$ by a sample of 2^N values from itself. This is equivalent to replacing each O_i by a random variable $\hat{O}$ which takes values from the original $\{O_i\}$ with equal probability, and hence its mean and variance, $E[\hat{O}]$ and $\text{Var}[\hat{O}]$, are the mean and variance of $\{O_i\}$. This replacement is typically a good approximation for observables like current, density and other correlations. Then, normalizing $|Q\rangle$ such that $\sum_i Q_i = 1$, we make the substitution of O_i for $\hat{O}$ in Eq. (B1) and find its expected value and variance to be

$$\begin{aligned} E[\langle O \rangle^P - \langle O \rangle^Q] &= 0, \\ \text{Var}[\langle O \rangle^P - \langle O \rangle^Q] &= \text{Var}[\hat{O}] \| |P\rangle - |Q\rangle \|_2^2. \end{aligned}$$

It then follows that

$$E[|\langle O \rangle^P - \langle O \rangle^Q|] \approx \text{SD}[\hat{O}] \| |P\rangle - |Q\rangle \|_2,$$

where SD means the standard deviation. This scales with the L_2 rather than the L_1 -norm error, and we conclude that if there is no correlation between observable values and the approximation method then typically the observable error will scale with the L_2 -norm error. An example of this is shown in Fig. 6(a). Note that due to the possibility of negative probabilities the normalization $\sum_i Q_i = 1$ is not always the same as $\| |Q\rangle \|_1 = 1$.

Appendix C: The Trotter and truncation errors

In approximating the full stochastic evolution $\exp(Ht)$ by a product of n time-steps, each an identical stochastic evolution $S = (\prod_{\ell=1}^{N-1} S_\ell)(\prod_{\ell=N-1}^1 S_\ell)$, we quantify the error by the L_1 -norm of the residual

$$\begin{aligned} \mathcal{E}_{\text{ST}} &= \left\| e^{Ht} - \prod_{j=n}^1 S \right\|_1, \\ &\leq \| (e^{H\delta t} - S) \prod_{j=n-1}^1 S |P\rangle \|_1 + \| e^{H\delta t} (e^{H\delta t} - S) \prod_{j=n-2}^1 S |P\rangle \|_1 \\ &\quad \dots + \left\| \prod_{j=n}^3 e^{H\delta t} (e^{H\delta t} - S) S |P\rangle \|_1 + \left\| \prod_{j=n}^2 e^{H\delta t} (e^{H\delta t} - S) |P\rangle \|_1, \\ &\leq n \| e^{H\delta t} - S \|_1. \end{aligned}$$

Now $\| e^{H\delta t} - S \|_1 \sim \delta t^3$, while $n = t/\delta t$ so the error scales at worst as $\mathcal{E}_{\text{ST}} \sim t \delta t^2$.

For convenience let's relabel the total set of two-site stochastic evolution operators over the whole simulation as $S_K S_{K-1} \dots S_k \dots S_2 S_1 |P\rangle$ where $K = 2nN$ and the index k denotes the position in the sequence. The second type of error is the truncation error incurred when any two-site operator S is applied to some MPS $|Q\rangle$. Rather than being applied exactly, the subsequent factorization and truncation back to an MPS with a maximum inner-dimension of χ produces an MPS $|Q'\rangle$. This can be thought of as being equivalent to

implementing exactly some other, generally non-stochastic, transformation T , which depends on k , $|Q\rangle$ and χ , such that $|Q'\rangle = T|Q\rangle$. The error in doing this is then quantified by the L_1 -norm $\epsilon = \|(S - T)|Q\rangle\|_1$. Once the sequence has been performed up to K the accumulated approximation is $T_K T_{K-1} \dots T_1 |Q\rangle$. For each step k it is possible to extract from the TEBD algorithm an upper-bound to the L_1 -norm error $\epsilon_k = \|(S_k - T_k) T_{k-1} \dots T_1 |P\rangle\|_1$ caused by truncation after applying S_k to the accumulative approximate state up to point $k-1$. Specifically for the SVD, for each two-site operation we can calculate the L_2 -norm error by the Eckart Young theorem Eq. (7) and thus get an upper-bound to each L_1 -norm ϵ_k by $2^{N/2}$ times this value. This information then provides an upper-bound to the total accumulative L_1 -norm error due to truncation $\mathcal{E}_{\text{tr}}$ after K two-site operations since

$$\begin{aligned} \mathcal{E}_{\text{tr}} &= \left\| \left\{ \prod_{k=1}^K S_k - \prod_{k=1}^K T_k \right\} |P\rangle \right\|_1, \\ &\leq \|(S_K - T_K) \prod_{k=1}^{K-1} T_k |P\rangle\|_1 + \|S_K (S_{K-1} - T_{K-1}) \prod_{k=1}^{K-2} T_k |P\rangle\|_1 \\ &\quad \dots + \left\| \prod_{k=3}^K S_k (S_2 - T_2) T_1 |P\rangle \right\|_1 + \left\| \prod_{k=2}^K S_k (S_1 - T_1) |P\rangle \right\|_1, \\ &\leq \|(S_K - T_K) \prod_{k=1}^{K-1} T_k |P\rangle\|_1 + \|(S_{K-1} - T_{K-1}) \prod_{k=1}^{K-2} T_k |P\rangle\|_1 \\ &\quad \dots + \|(S_2 - T_2) T_1 |P\rangle\|_1 + \|(S_1 - T_1) |P\rangle\|_1, \\ &= \sum_{k=1}^K \epsilon_k. \end{aligned}$$

Thus the truncation error $\mathcal{E}_{\text{tr}}$ grows at worst additively during the Trotter sequence and can be monitored within the algorithm by computing the sum of the upper-bounds for each two-site operation. This gives us a conservative upper-bound $\mathcal{E}_{\text{tr}}^+ \geq \mathcal{E}_{\text{tr}}$.

Appendix D: Non-negative matrix factorization algorithms

The lack of convexity and non-uniqueness of the NMF problem for the three most common cost functions in Eq. (9) has prompted a variety of algorithms to be proposed [69]. These algorithms tackle the minimization problem directly for the required rank χ using alternating least-squares, conjugate gradients, multiplicative updates or projected gradient descent [47, 69, 70]. Here we shall briefly mention some details for the latter two approaches. The projected gradient descent method additively updates the elements of the W and H matrices as

$$\begin{aligned} W_{i\mu} &\mapsto W_{i\mu} + \eta_W \frac{\partial F}{\partial W_{i\mu}}, \\ (H^T)_{\mu j} &\mapsto (H^T)_{\mu j} + \eta_H \frac{\partial F}{\partial (H^T)_{\mu j}}, \end{aligned}$$

where η_W and η_H are appropriately chosen descent step-sizes. For the L_1 -norm the relevant derivatives are

$$\frac{\partial F}{\partial W_{\mu}} = - \sum_j \frac{\Delta_{ij}}{|\Delta_{ij}|} (H^T)_{\mu j}$$

$$\frac{\partial F}{\partial (H^T)_{\mu j}} = - \sum_i W_{i\mu} \frac{\Delta_{ij}}{|\Delta_{ij}|},$$

where $\Delta_{ij} = (\Theta - WH^T)_{ij}$ is an element of the residual matrix. Under these update rules there is no explicit preservation of the non-negativity of W and H so typically a projection of the updated matrices is made to the positive orthant at each descent step by setting any negative elements to zero. Furthermore there is no preservation of the unit L_1 -norm of Θ within the approximation WH^T so this too has to be explicitly enforced [69].

Gradient descent can also be formulated for the L_2 -norm and the KL divergence. However, by making step sizes η_W and η_H depend on the matrix element being minimized via specially chosen functions of the current matrices W and H , so-called multiplicative algorithms can be derived [47, 60]. For the KL divergence the updates to be applied are

$$W_{i\mu} \mapsto \frac{W_{i\mu}}{\sum_j (H^T)_{\mu j}} \sum_j \left\{ (H^T)_{\mu j} \frac{\Theta_{ij}}{(WH^T)_{ij}} \right\},$$

$$(H^T)_{\mu j} \mapsto \frac{(H^T)_{\mu j}}{\sum_i W_{i\mu}} \sum_i \left\{ W_{i\mu} \frac{\Theta_{ij}}{(WH^T)_{ij}} \right\}.$$

Under these transformations the KL divergence is non-increasing and is invariant if and only if W and H are at a stationary point [60].

Overall, we find that the above NMF algorithms are highly sensitive to the initial conditions. This means that to get reasonable results many random restarts are required and considerable fluctuations in the final cost function value are observed. They also often show slow convergence (if at all) and are generally much slower than the SVD. For the most relevant case of the L_1 -norm cost function we have found that using a random initial matrix driven to a KL divergence solution often provides a much higher quality result than simply applying the gradient descent to random initial matrices. Despite this, however, we have found that, for our task, the SVD is superior to current NMF algorithms not only at moderate truncations but also at identifying when a near exact low-rank solution exists [74]. This is an essential ingredient of an effective decimation algorithm.

-
- [1] J. Krug, Phys. Rev. Lett. **67**, 1882 (1991).
 - [2] G. Grinstein, D.-H. Lee and S. Sachdev, Phys. Rev. Lett. **64**, 1927 (1990).
 - [3] A. Schadschneider, Physica A **313**, 153 (2002).
 - [4] D. Helbing, Rev. Mod. Phys. **73**, 1067 (2001).
 - [5] D. Chowdhury, L. Santen and A. Schadschneider, Phys. Rep. **329**, 199 (2000).
 - [6] D. Chowdhury, A. Schadschneider and K. Nishinari, Phys. Life Rev. **2**, 318 (2005).
 - [7] Y. Aghababaie, G. I. Menon and M. Plischke, Phys. Rev. E **59**, 2578 (1999).
 - [8] T. Chou and D. Lohse, Phys. Rev. Lett. **82**, 3552 (1999).
 - [9] A. Parmeggiani, T. Franosch and E. Frey, Phys. Rev. Lett. **90**, 086601 (2003).
 - [10] B. Derrida, M. R. Evans, V. Hakim and V. Pasquier, J. Phys. A **26**, 1493 (1993).
 - [11] R. A. Blythe and M. R. Evans, J. Phys. A **40**, R333 (2007).
 - [12] G. M. Schütz, in *Phase Transitions and Critical Phenomena*, edited by C. Domb and J. L. Lebowitz (Academic Press, London, 2000), Vol. 19.
 - [13] G. Vidal, Phys. Rev. Lett. **91**, 147902 (2003).
 - [14] G. Vidal, Phys. Rev. Lett. **93**, 040502 (2004).
 - [15] D. Gobert, C. Kollath, U. Schollwöck and G. Schütz, Phys. Rev. E **71**, 036102 (2005).
 - [16] S. R. White and A. E. Feiguin, Phys. Rev. Lett. **93**, 076401 (2004).
 - [17] S. R. Clark and D. Jaksch, Phys. Rev. A **70**, 043612 (2004).
 - [18] A. J. Daley, C. Kollath, U. Schollwöck and G. Vidal, J. Stat. Mech. P04005 (2004).
 - [19] D. Jaksch, Contemporary Physics **45**, 367 (2004).
 - [20] A. J. Daley, S. R. Clark, D. Jaksch and P. Zoller, Phys. Rev. A **72**, 043618 (2005).
 - [21] S. R. Clark and D. Jaksch, New J. Phys. **8**, 160 (2006).
 - [22] M. Bruderer, A. Klein, S. R. Clark and D. Jaksch, Phys. Rev. A **76**, 011605(R) (2007).
 - [23] M. Bruderer, A. Klein, S. R. Clark and D. Jaksch, New J. Phys. **10**, 033015 (2008).
 - [24] D. T. Gillespie, J. Comput. Phys. **22**, 403 (1976).
 - [25] V. Popkov, M. Salerno and G. M. Schütz, Phys. Rev. E **78**, 011122 (2008).
 - [26] A. Lipowski and D. Lipowska, Phys. Rev. E **79**, 060102(R) (2009).
 - [27] M. Fannes, B. Nachtergaele, and R. F. Werner, Commun. Math. Phys. **144**, 443 (1992).
 - [28] S. Östlund and S. Rommer, Phys. Rev. Lett. **75**, 3537 (1995).
 - [29] S. Rommer and S. Östlund, Phys. Rev. B **55**, 2164 (1997).
 - [30] D. Perez-Garcia, F. Verstraete, M. M. Wolf and J. I. Cirac, Quant. Inf. Comput. **7**, 401 (2007).
 - [31] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).
 - [32] S. R. White, Phys. Rev. B **48**, 10345 (1993).
 - [33] U. Schollwöck, Rev. Mod. Phys. **77**, 259 (2005).
 - [34] Y. Hieida, J. Phys. Soc. Jpn. **67**, 369 (1998).
 - [35] E. Carlon, M. Henkel and U. Schollwöck, Eur. Phys. J. B **12**, 99 (1999).
 - [36] E. Carlon, M. Henkel and U. Schollwöck, Phys. Rev. E **63**, 036101 (2001).
 - [37] R. J. Glauber, J. Math. Phys. **4**, 294 (1963).
 - [38] F. C. Alcaraz, M. Droz, M. Henkel and V. Rittenberg, Ann.

- Phys. **230**, 250 (1994).
- [39] O. Golinelli and K. Mallick, J. Phys. A **39**, 12679 (2006).
 - [40] B. Derrida, Phys. Rep. **301**, 65 (1998).
 - [41] B. Schmittmann and R. K. P. Zia, in *Phase Transitions and Critical Phenomena*, edited by C. Domb and J. L. Lebowitz (Academic Press, New York, 1995), Vol. 17.
 - [42] R. Stinchcombe, Adv. Phys. **50**, 431 (2001).
 - [43] S. Sachdev, *Quantum Phase Transitions* (Cambridge University Press, Cambridge, 2001).
 - [44] N. Rajewsky, L. Santen, A. Schadschneider and M. Schreckenberg, J. Stat. Phys. **92**, 151 (2004).
 - [45] G. H. Golub and C. F. V. Loan, *Matrix Computations* (Johns Hopkins University Press, Baltimore, 1996), 3rd Ed.
 - [46] P. Paatero and U. Tapper, Environmetrics **5**, 111 (1994).
 - [47] D. D. Lee and H. S. Seung, Nature **401**, 788 (1999).
 - [48] K. Temme and F. Verstraete, Phys. Rev. Lett. **104**, 210502 (2010).
 - [49] F. Verstraete, V. Murg and J. I. Cirac, Adv. Phys. **57**, 143 (2008).
 - [50] J. I. Cirac and F. Verstraete, J. Phys. A **42**, 504004 (2009).
 - [51] F. Verstraete and J. I. Cirac, Phys. Rev. B **73**, 094423 (2006).
 - [52] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
 - [53] L. Amico, R. Fazio, A. Osterloh and V. Vedral, Rev. Mod. Phys. **80**, 517 (2008).
 - [54] J. Eisert, M. Cramer and M. B. Plenio, Rev. Mod. Phys. **82**, 277 (2010).
 - [55] E. Gaussier and C. Goutte, in *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval* (Salvador, 2005), p. 601.
 - [56] C. Eckart and G. Young, Psychometrika **1**, 211 (1936).
 - [57] E. Seneta, *Non-negative matrices and Markov chains* (Springer, New York, 2006), 2nd Ed.
 - [58] M. Suzuki, Phys. Lett. A **146**, 319 (1990); J. Math. Phys. **32**, 400 (1991).
 - [59] N. Vasiloglou, A. G. Gray and D. V. Anderson, in *Proceedings of the Ninth SIAM International Conference on Data Mining* (Sparks, Nevada, 2009), p. 673.
 - [60] D. D. Lee and H. S. Seung, Advances in Neural Information Processing Systems **13**, 556 (2000).
 - [61] M. Zwolak and G. Vidal (private communication, 2004).
 - [62] J. Keizer, J. Stat. Phys. **6**, 67 (1972).
 - [63] Y. Y. Shi, L. M. Duan and G. Vidal, Phys. Rev. A **74**, 022320 (2006).
 - [64] I. L. Markov, Y. Y. Shi, Algorithmica DOI: 10.1007/s00453-009-9312-5 (2009).
 - [65] R. Huebener, V. Nebendahl and W. Dür, New J. Phys. **12**, 025004 (2010).
 - [66] F. Fröwis, V. Nebendahl and W. Dür, preprint arXiv:1003.1047v1 (2010).
 - [67] M. C. Bañuls, M. B. Hastings, F. Verstraete and J. I. Cirac, Phys. Rev. Lett. **102**, 240603 (2009).
 - [68] E. M. Stoudenmire and S. R. White, New J. Phys. **12**, 055026 (2010).
 - [69] M. W. Berry, M. Browne, A. N. Langville, V. P. Pauca and R. J. Plemmons, Computational Statistics and Data Analysis **52**, 155 (2007).
 - [70] C.-J. Lin, Neural Computation **19**, 2756 (2007).
 - [71] It follows from Eq. (1) that an MPS is invariant under the transformation $A^{[\ell]i_\ell} \mapsto A^{[\ell]i_\ell} X$, $A^{[\ell+1]i_{\ell+1}} \mapsto X^{-1} A^{[\ell+1]i_{\ell+1}}$, where X is a non-singular matrix. The formalism of Sec. (II A) and Appendix A automatically exploits this freedom so that the tensors $\{A^{[\ell]}\}$ satisfy orthonormality conditions.
 - [72] Despite not being a distance measure the KL divergence does have an information theoretic interpretation. When $\sum_{ij} (WH^T)_{ij} = 1$, in addition to Θ , the KL divergence reduces to the relative entropy of the two probability distributions. It then quantifies the expected number of extra bits required to code samples from Θ when using a code based on WH^T , rather than using a code based on Θ .
 - [73] A similar situation occurs for imaginary time-evolution of quantum systems and the evolution of mixed states according to a quantum master equation [61].
 - [74] We note that for data-mining applications the SVD is typically used to benchmark what rank reduction would be acceptable for a subsequent NMF, precisely because it identifies a suitable rank which contains most of the data.