

Stable unitary integrators for the numerical implementation of continuous unitary transformations

Samuel Savitz* and Gil Refael†

Institute of Quantum Information and Matter, Department of Physics, California Institute of Technology, Pasadena, California 91125, USA

(Received 31 July 2017; published 18 September 2017)

The technique of continuous unitary transformations has recently been used to provide physical insight into a diverse array of quantum mechanical systems. However, the question of how to best numerically implement the flow equations has received little attention. The most immediately apparent approach, using standard Runge-Kutta numerical integration algorithms, suffers from both severe inefficiency due to stiffness and the loss of unitarity. After reviewing the formalism of continuous unitary transformations and Wegner’s original choice for the infinitesimal generator of the flow, we present a number of approaches to resolving these issues including a choice of generator which induces what we call the “uniform tangent decay flow” and three numerical integrators specifically designed to perform continuous unitary transformations efficiently while preserving the unitarity of flow. We conclude by applying one of the flow algorithms to a simple calculation that visually demonstrates the many-body localization transition.

DOI: [10.1103/PhysRevB.96.115129](https://doi.org/10.1103/PhysRevB.96.115129)**I. INTRODUCTION**

While the technique of continuous unitary transformations (CUTs) was already known to contemporary mathematicians [1], it was introduced to the physics community roughly simultaneously by Franz Wegner in the language of “flow equations” [2] and by Kenneth Wilson and Stanisław Głazek who recast it as a renormalization group [3,4]. Thus, the terms “Wegner flow” or “Wegner-Wilson flow” (WWF) are established in the condensed matter literature as referring to a particular kind of CUT. The applications and mathematical properties of CUTs have been reviewed by Monthus [5] and Bartlett [6]. They have also been analyzed in terms of geodesic flows [7].

Interest in CUTs has increased recently due to the discovery that they may provide physical insights when applied to important many-body problems such as the poorly understood phenomenon of many-body localization (MBL) [8,9]. The study of MBL has exploded since Basko, Aleiner, and Altshuler published their seminal diagrammatic analysis of the phenomenon in 2006 [10]. Our understanding of MBL is deeply connected to the concepts of ergodicity and thermalization in quantum systems and has been reviewed by Nandkishore and Huse [11]. Recent results of this approach include using the WWF to find ℓ -bits representations of an MBL Hamiltonian [12,13] and to identify an Anderson transition in the power-law banded random matrix model [14]. While we will not address MBL directly until Sec. VII, we have used it as a source of interesting initial Hamiltonians in Figs. 1 and 2.

Despite this growing interest, little has been said about how best to numerically implement these flows *in silico*. In Kehrein’s monograph on the application of CUTs to many-body problems [15], the standard Runge-Kutta family of adaptive numerical integrators, including the popular Dormand-Prince method [16], is described as a “good algorithm” for this purpose. Nevertheless, we have identified a number of ways in which it can be dramatically improved. In particular, our methods eliminate a severe inefficiency near the end of

the flow, once many of the off-diagonal elements have been almost entirely eliminated. A second flaw exhibited by the current flow algorithms which we resolve is the violation of the unitarity of the transformation in a step size-dependent manner, leading to subsequent errors in the final eigenvalues.

After first reviewing the formalism of CUTs, we will proceed to explain the origins of these issues in Sec. III. In Sec. IV, we will propose a choice of infinitesimal generator that leads to what we call the “uniform tangent decay flow” and which elegantly resolves the efficiency problem without having to change integrators. Then, in Sec. V, we will solve the efficiency and unitarity issues, regardless of the choice of generator, by developing two stable, geometric [17,18] integrators designed specifically for CUTs. One is accurate to first order, and the other to third order. In Sec. VI, we will present an entirely different approach to resolving these problems that was inspired by quantized state simulators [19,20] and the Trotter decomposition [21]. Finally, we will demonstrate a simple way in which CUTs can provide insight and help visualize the MBL transition before giving some concluding thoughts regarding the interpretation of, and potential future directions for, the study of CUTs.

II. FORMALISM OF CONTINUOUS UNITARY TRANSFORMATIONS**A. Notation**

We begin with a Hamiltonian \mathbf{H} represented by an $n \times n$ Hermitian matrix. For the sake of simplicity and clarity, its elements, \mathbf{H}_{ab} , will be restricted to be real, although nothing prevents what follows from being extended to complex matrices. We will refer to the diagonal elements by $D_a = \mathbf{H}_{aa}$, and the off-diagonal elements by $\mathbf{J}_{ab} = \mathbf{J}_{ba} = \mathbf{H}_{ab}$. Furthermore, we will use $\Delta_{ab} = D_a - D_b$ to represent the current energy spacing between any two distinct states $|a\rangle$ and $|b\rangle$, ignoring off-diagonal elements. $\mathbf{X}_{ab} = -\mathbf{X}_{ba} = \Delta_{ab}/2$ will be half this value.

A generic unitary transformation applied to \mathbf{H} can be represented by $\mathbf{H}' = \mathbf{U}\mathbf{H}\mathbf{U}^{-1}$ where \mathbf{U} is a unitary matrix

*Sam@Savitz.org

†Refael@Caltech.edu

($\mathbf{U}^{-1} = \mathbf{U}^\dagger$). This can also be written as

$$\mathbf{H}' = e^\eta \mathbf{H} e^{-\eta}, \quad (1)$$

where η is the anti-Hermitian ($\eta^\dagger = -\eta$) matrix logarithm of \mathbf{U} , which we will also assume to have only real elements [22].

B. The flow equations

In the limit of a small unitary rotation, we can replace η with $d\tau \eta$. \mathbf{U} converges to the identity matrix as the infinitesimal $d\tau$ approaches zero. Truncating the Taylor series expansion of the exponential,

$$\mathbf{U} = e^{d\tau \eta} = \sum_{k=0}^{\infty} \frac{d\tau^k}{k!} \eta^k, \quad (2)$$

after the second term gives $\mathbf{U} \approx 1 + d\tau \eta$, and $\mathbf{U}^{-1} \approx 1 - d\tau \eta$. Therefore, to first order,

$$\mathbf{H}' \approx (1 + d\tau \eta) \mathbf{H} (1 - d\tau \eta) \approx \mathbf{H} + d\tau [\eta, \mathbf{H}]. \quad (3)$$

Treating τ as a fictitious “flow-time” coordinate, this can be recast as the first-order, nonlinear differential equation

$$\frac{d\mathbf{H}}{d\tau} = \dot{\mathbf{H}} = [\eta, \mathbf{H}]. \quad (4)$$

The diagonalizing unitary evolves according to $\dot{\mathbf{U}} = \eta \mathbf{U}$. These differential equations are known as the “flow equations”. While \mathbf{U} begins as the identity matrix, \mathbf{H} is initialized to whatever starting Hamiltonian one desires: $\mathbf{H}(\tau = 0) = \mathbf{H}_0$. Written in terms of the elements of \mathbf{H} , the flow equations are

$$\dot{D}_a = 2 \sum_{c \neq a} \eta_{ac} \mathbf{J}_{ca}, \quad (5a)$$

$$\dot{\mathbf{X}}_{ab} = 2\eta_{ab} \mathbf{J}_{ab} + \sum_{c \neq a, b} [\eta_{ac} \mathbf{J}_{ca} - \eta_{bc} \mathbf{J}_{cb}], \quad (5b)$$

and

$$\dot{\mathbf{J}}_{ab} = -2\eta_{ab} \mathbf{X}_{ab} + \sum_{c \neq a, b} [\eta_{bc} \mathbf{J}_{ca} + \eta_{ac} \mathbf{J}_{cb}]. \quad (5c)$$

Equation (5b) follows immediately from Eq. (5a) and possesses a pleasing symmetry with Eq. (5c).

C. The two-state limit

The proper interpretation of the generator η can be elucidated by recasting the flow in angular terms. In order to clarify the intuitive picture, we will temporarily restrict our attention to the $n = 2$ case of two-by-two matrices. Equivalently, we can assume that η has only one pair of nonzero elements $\eta_{ab} = -\eta_{ba}$. Either of these restrictions allows us to neglect the summed terms in Eqs. (5b) and (5c). We will refer to this approximation as the “two-state limit”, and it will be recurrently useful. We will henceforth sometimes lighten and decapitalize the upper case letters denoting matrices and drop the ab subscripts when they do not enhance clarity [23].

Without further loss of generality, we can now write

$$\mathbf{H} = \begin{pmatrix} x & j \\ j & -x \end{pmatrix} = j\sigma_1 + x\sigma_3, \quad (6a)$$

and

$$\eta = \begin{pmatrix} 0 & \eta \\ -\eta & 0 \end{pmatrix} = i\eta\sigma_2, \quad (6b)$$

where the σ 's denote the standard Pauli matrices. Conveniently, these matrices will retain this form throughout the flow, so we can understand their evolution as a first-order differential equation for $x(\tau)$ and $j(\tau)$.

The angular interpretation

For each pair of distinct states, we can define an angle θ to equal $\text{atan2}(j, x)$, i.e., the complex phase of the value $z = x + ij$ [24]. A positive radius $r = \sqrt{x^2 + j^2}$ is also given by its magnitude, $|z|$. In these polar coordinates, $x = r \cos \theta$, and $j = r \sin \theta$. Note that the condition of \mathbf{H} being successfully diagonalized is, barring degeneracies, equivalent to the requirement that all of the θ 's satisfy $\sin \theta = 0$ by being an integer multiple of π .

Finally, differentiating θ according to the chain rule gives

$$\dot{\theta} = \frac{x\dot{j} - j\dot{x}}{r^2} = -2\eta. \quad (7)$$

Thus, η_{ab} is clearly proportional to the rate of rotation between the $|a\rangle$ and $|b\rangle$ axes of the Hamiltonian's basis. The negative two factor in Eq. (7) is identical in origin to the phenomenon of angle doubling in spinor homomorphisms [25]. The complication outside of the two-state limit comes, of course, from the sum terms in Eqs. (5b) and (5c), which we are neglecting. They capture how this rotation between the $|a\rangle$ and $|b\rangle$ axes affects the couplings between those two states and any distinct third state, $|c\rangle$, namely \mathbf{J}_{ac} and \mathbf{J}_{bc} .

D. Choice of generator

In order to finish specifying the flow, we must choose a formula for η . Most often η_{ab} is a function of only \mathbf{X}_{ab} , \mathbf{J}_{ab} , and, in the case of the Toda flow, the sign $\text{sgn}(b - a)$. Defining η purely in terms of the “current” value of \mathbf{H} allows us to think about a given CUT as a time-independent, first-order differential equation flowing over the space of all Hermitian matrices. Wegner's original choice of generator was $\eta = [\mathbf{H}_{\text{Diag}}, \mathbf{H}]$, where \mathbf{H}_{Diag} is the diagonal part of \mathbf{H} . Expressed in terms of the matrix elements, this specifies that $\eta = \delta j = r^2 \sin 2\theta$. Recall that $\delta = 2x$ is the difference of the two diagonal elements.

As Monthus recounts [5], a number of alternative choices of generator have since been presented. The flows reviewed in her paper are summarized in Table I. In conjunction with Eq. (7), the polar forms of these generators make it apparent why they must induce evolution towards a diagonalized Hamiltonian.

E. Matrix metrics and the convergence towards diagonalization

In order to show more rigorously that these flows ultimately lead to the diagonalization of the Hamiltonians to which they are applied, it will be useful to introduce some metrics which can be applied to the Hamiltonian during the course of the

TABLE I. The best-known flow generators, and our tangent flow, expressed in terms of both the matrix elements and their polar form.

Flow	η_{ab}
WWF [2]	$\delta j = r^2 \sin 2\theta$
White [26]	$j/\delta = \tan(\theta)/2$
Sign [27]	$\text{sgn}(x) j = \text{sgn}(x) r \sin \theta$
Toda [28]	$\text{sgn}(b-a) j = \text{sgn}(b-a) r \sin \theta$
Tangent	$\delta j/(x^2 + j^2) = \sin 2\theta$

flow: The Frobenius norm of a matrix, $\|\cdot\|_F$, is defined as

$$\|\mathbf{H}\|_F = \sqrt{\sum_{a,b} \mathbf{H}_{ab}^2} = \sqrt{I_2^D + I_2^J}, \quad (8)$$

where [5]

$$I_2^D = \sum_a D_a^2, \quad (9a)$$

and

$$I_2^J = \sum_{a \neq b} \mathbf{J}_{ab}^2. \quad (9b)$$

Continuing to decompose terms, we find that

$$I_2^D = \frac{I_2^\Delta + (\text{Tr } \mathbf{H})^2}{n}, \quad (10)$$

given

$$I_2^\Delta = \sum_{a < b} \Delta_{ab}^2 = n^2 \sigma_D^2, \quad (11)$$

where

$$\sigma_D^2 = \frac{1}{n} \sum_a (D_a - \mu)^2 \quad (12)$$

is the population variance of the diagonal elements, and $\mu = \text{Tr } \mathbf{H}/n$ is their mean. Both the Frobenius norm and the trace are invariant under unitary transformations. Therefore, the effect of the flow is to transfer weight directly between I_2^J and I_2^D , or equivalently, between I_2^J and $n\sigma_D^2$.

We can now immediately calculate that

$$I_2^D = 4 \sum_{a \neq b} \eta_{ab} \mathbf{X}_{ab} \mathbf{J}_{ab} = 2 \sum_{a \neq b} \eta_{ab} \mathbf{R}_{ab}^2 \sin 2\Theta_{ab}. \quad (13)$$

Thus, a choice of generator which ensures that the sign of η always matches that of $xj \propto \sin 2\theta$ must cause I_2^D to evolve in a nondecreasing manner. It is easy to verify that all of the generators in Table I, with the exception of the Toda [29], satisfy this condition.

In fact, the original WWF generator is given by the steepest descent of I_2^J : [6]

$$\eta \propto - \left. \frac{dI_2^J}{d\eta} \right|_{\eta=0}. \quad (14)$$

In other words, each independent component of η is set in proportion to how quickly it causes I_2^J to decrease. Given a fixed Frobenius norm $\|\eta\|_F$, no alternative choice of η leads to a faster decrease than the WWF.

As noted above, the flow towards diagonalization decreases I_2^J and increases σ_D . In other words, as j decays towards zero, the corresponding level spacing δ increases: The eigenvalues appear to repel each other. This behavior is related in origin to the Dyson Brownian motion of eigenvalues in random matrix theory [30,31].

One can monitor the progression of a flow towards diagonalization using the metric ρ which we will define in terms of Eqs. (9a) and (11) as

$$\rho = \sqrt{\frac{2I_2^J}{I_2^\Delta + 2I_2^J}} = \sqrt{\frac{2I_2^J}{n^2\sigma_D^2 + 2I_2^J}}. \quad (15)$$

The Hamiltonian is diagonalized when ρ reaches zero.

F. Fixed points, viscous pendula, the Gudermannian function, and the renormalization group

While all stable fixed points of the WWF must be diagonal, an unstable fixed point can have nonzero \mathbf{J}_{ab} if the corresponding \mathbf{X}_{ab} is zero, i.e., the relevant diagonal elements are equal. Our uniform tangent decay flow introduced in Sec. IV also behaves this way, but the other options vary in how they respond to this situation: Both White's flow and the sign flow are ill defined at such points, and the Toda flow does not have it as a fixed point at all.

In the two-state limit, the dynamics and fixed points of the WWF and the tangent flow are identical to those of a viscous pendulum. By "viscous", we mean that the equations of motion are first order and inertial effects are negligible. Like the pendulum, the two-state solutions to both of these flows can be analytically expressed in terms of the Gudermannian function, $\text{gd } x = \sin^{-1}(\tanh x)$ [32].

For the WWF, the density of the pendulum's bob, and therefore the speed of its progression along the Gudermannian trajectory, increases like r^2 . This faster resolution of higher-energy couplings leads to the renormalization group interpretation of the WWF [3,4]. Note, however, that the arbitrarily slow "tipping" dynamics around the unstable fixed points makes this analogy somewhat rough and inexact. The tangent flow, on the other hand, holds the density of each bob constant, and so the renormalization interpretation is not even approximately applicable. White's flow [26] avoids this slow tipping behavior, but introduces its own idiosyncrasies, discussed in Sec. IV A.

III. NUMERICAL IMPLEMENTATION ISSUES

A. The Dormand-Prince method

A typical first attempt at numerically implementing CUTs is to apply an adaptive Runge-Kutta integrator such as the well-known Dormand-Prince method [16] to Eq. (4). We will review the Dormand-Prince method only by explaining that this general-purpose fifth-order algorithm iteratively employs a handful of derivative evaluations in order to propagate \mathbf{H} forward in flow time by a small step size h while estimating its own error so as to adjust h' for the next step. Calling it fifth order means that the calculation error theoretically decreases like h^5 as the number of steps increases. Due to the potentially scarce nature of the matrix of errors \mathbf{E} , we suggest using the

entry-wise infinity norm, $\|\cdot\|_\infty^{\text{entry}}$, in order to reduce \mathbf{E} to a single scalar representing the overall error magnitude:

$$\|\mathbf{E}\|_\infty^{\text{entry}} = \max_{a,b} |\mathbf{E}_{ab}|. \quad (16)$$

B. Stiffness

Upon plugging Wegner’s original choice of generator, $\eta = r^2 \sin 2\theta$, into the Dormand-Prince method, one finds that the fictitious flow time τ must grow surprisingly large before the matrix can be called diagonalized and the flow completed. This is not an artifact of the implementation; the r^2 factor in the generator causes those off-diagonal elements with small radii to decay very slowly. While one might hope that the integrator would accelerate and begin taking larger steps in this regime, this cannot occur because it would lead to over-correction and oscillatory instability in those elements with large radii. Thus, the integrator is forced to take small steps indefinitely. This general class of inefficiency is well-known in the numerical analysis literature and is referred to as “stiffness”. The most established approach for remediating stiffness is called “stabilization”. Stabilized integrators are often implicit, meaning that they require an equation to be numerically solved during the course of each step [33,34].

One can resolve this problem by using a different flow generator, but as we will show in Sec. IV, none of the previously proposed options fit the bill. We will therefore construct a generator which resolves this issue. Furthermore, in Secs. V and VI, we will introduce three integrators which are able to sidestep this issue, even with Wegner’s original choice of generator, through two very different mechanisms.

C. Unitarity

A second flaw which becomes apparent when applying the Dormand-Prince method to CUTs is the slow loss of the unitary similarity relation between \mathbf{H} and \mathbf{H}_0 . This leads to errors in the calculated eigenvalues. To some extent, such fluctuation is inevitable due to floating-point rounding errors, but the Dormand-Prince method deviates even in the absence of rounding in a step size-dependent manner.

In order to understand this source of error and how to mitigate it, consider the simple system given by $\dot{x} = -y$ and $\dot{y} = x$. These differential equations clearly induce uniform, counterclockwise circular motion about the origin with a constant radius $\sqrt{x^2 + y^2}$. However, when applied to this system, most numerical integrators will cause the radius to eventually either converge to zero or diverge to infinity in an exponentially spiraling fashion. Like the family of exponential integrators [35] which were constructed for differential equations with approximately exponential behavior, one could easily write a far superior, specialized integrator for this system, or even a perturbation on it, which uses a rotation matrix in its step propagator in order to avoid this issue. This broadly applicable technique of exactly solving for the contributions from individual terms in a differential equation separately before combining them is known as “operator splitting” [36].

The integrators in Secs. V and VI extend this idea to CUTs and therefore preserve the unitarity similarity, up to rounding error. In general, integrators which restrict their evolution to some exact submanifold in phase-space are called

“geometric integrators”. Because the invariant in our case is unitary similarity, our integrators will fall into the subcategory of “unitary integrators”.

IV. THE UNIFORM TANGENT DECAY FLOW

A. White’s flow

We will now provide a possible resolution to the issue of numerical instability and stiffness by constructing an infinitesimal flow generator. Out of the established flows in Table I, only White’s flow [26] lacks the stiffening r or r^2 factors. Returning to the two-state limit in Sec. IIC, i.e., neglecting the difficult summed terms in Eqs. (5b) and (5c), we can see that this flow generator causes each off-diagonal element to attempt to decay towards zero with the same characteristic time constant [37]. At first glance, one might think that the fact that I_2^J decays in exact proportion with $e^{-2\tau}$ resolves the inefficiency issue, but correctly implementing White’s flow with adaptive time steps is also extremely slow.

The severe slowdown in the numerical implementation of White’s flow arises at each attempted level crossing, when δ becomes *very* close to zero, and θ likewise approaches $\pm\pi/2$. This causes the generator $\eta = \tan(\theta)/2$ to begin to diverge in magnitude, the Hamiltonian to rotate extremely quickly with respect to τ , and the two diagonal elements to repel each other, so as to ultimately prevent δ from ever actually changing sign. Only once the Hamiltonian has rotated sufficiently far so as to actually interchange the two relevant axes is this “avoided crossing” completed, with δ retreating back up from zero. Through this somewhat bizarre and very slow to correctly simulate process, White’s flow generally prevents level crossings and maintains the numerical order of the diagonal elements of the Hamiltonian.

B. Derivation of the tangent flow

What is needed is some kind of “softening” of White’s generator when θ equals $\pm\pi/2$. Let us attempt to induce the uniform exponential decay not of each j , as is the case with White’s flow, but instead of each $\tan\theta = j/x$. The standard differential quotient rule tells us that, in the two-state limit,

$$\frac{d \tan \theta}{\tau} = \frac{x \dot{j} - j \dot{x}}{x^2} = -2\eta \frac{(x^2 + j^2)}{x^2} = -2\eta(\tan^2 \theta + 1). \quad (17)$$

This means that our η must be proportional to

$$\frac{2 \tan \theta}{\tan^2 \theta + 1} = \sin 2\theta = \frac{\delta j}{x^2 + j^2} = \frac{2}{x/j + j/x}. \quad (18)$$

Taking this as our choice of generator leads to what we call the “uniform tangent decay flow” which conveniently avoids both the stiffness issues exhibited by Wegner’s flow and the strange level crossing behavior of White’s flow, although level crossings still lead to some amount of reduction in step size. The tangent flow’s exponential approach towards diagonalization is demonstrated for a generic MBL Hamiltonian in Fig. 1.

Note that this generator is essentially the same as that of the WWF, but without the r^2 factor which quadratically

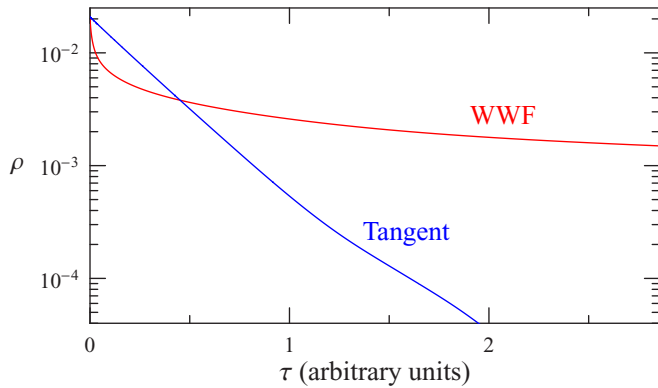


FIG. 1. In general, the tangent flow exponentially approaches diagonalization with respect to flow time, while the WWF slows down. This particular flow was initialized with a half-filled ten-site MBL Hamiltonian with $W = 1$ as described in Sec. VII. The metric ρ is defined in Eq. (15). Using the Dormand-Prince integrator, calculation time is roughly proportional to flow time, because the step sizes are bounded, although the relative flow-time scales shown here are arbitrary, reflecting our freedom to arbitrarily rescale the WWF flow time, which is inversely proportional to the square of the energy scale.

biases those off-diagonal elements with large radii to decay more rapidly. This prevents the usual renormalization group interpretation discussed in Sec. II F from being applicable to the tangent flow.

V. STABLE UNITARY INTEGRATORS

A. Introduction

As mentioned in Sec. III C, the first integrators that we will introduce for numerically implementing CUTs were inspired by the family of exponential integrators [35], which we will now briefly motivate: Given an “almost linear” differential equation such as $\dot{y} = \alpha y + \epsilon(y)$ where $\epsilon(y)$ can be regarded as a small perturbation on the first term, one can use our ability to analytically integrate the unperturbed system ($y(\tau) = y(0)e^{\alpha\tau}$) in order to take steps which are exactly correct in that limit and then “add the perturbation back in”. This frequently allows for highly accelerated performance when compared to the usual linear integrators. The widely applicable approach of exactly solving for the contributions from multiple parts of the system individually before combining them into a final step operation is known as “operator splitting” [36].

Furthermore, our integrators are in the class of geometric integrators. This means that they exactly preserve a geometric property, in this case the unitary similarity to \mathbf{H}_0 , up to rounding error [17,18]. Symplectic integrators, for instance, take steps that are proper symplectomorphisms, i.e., they preserve the symplectic two-form, and are a classic illustration of geometric integrators [38]. While the general idea of developing unitary integrators has been explored before [39–42], our particularly simple approach to stabilizing these integrators specifically to allow them to efficiently handle the stiffness of the WWF, is, to our knowledge, new.

B. First order

1. Integrator step

The crux of the idea is to, instead of starting with Eq. (4), go back to Eq. (1). Thus, each of the integrators’ steps will be governed by the formula

$$\mathbf{H}(\tau + h) = e^{h\eta(\tau)} \mathbf{H}(\tau) e^{-h\eta(\tau)}. \quad (19)$$

While the problem of actually computing matrix exponentials is notoriously rife with pitfalls [43], our situation is particularly painless because $\|h\eta\|$ can be assumed to be small. We therefore need to evaluate only a handful of terms of the infinite series

$$e^{h\eta} = \sum_{k=0}^{\infty} \frac{h^k}{k!} \eta^k \quad (20)$$

before the remaining terms fall below the rounding limit.

In fact, it is not necessary to make use of Eq. (20). Instead, the order (1, 1) Padé approximant to the exponential function,

$$e^{h\eta} \approx \frac{2 + h\eta}{2 - h\eta}, \quad (21)$$

can be faster to calculate and is perfectly unitary and accurate up to the order of our method. We have therefore used it in place of the exponential Taylor series. This approximation is also known as the Cayley transform [40,44]. Here, and for the subsequent third-order integrator, the choice of matrix exponential algorithm can be critical to performance, and experimentation is advised, taking the specific demands and available implementations into account.

2. Stabilization

This algorithm already preserves unitarity, but it is known in the literature and is still susceptible to the instability that results from the stiffness of the WWF. To resolve this issue, we will take advantage of our ability to exactly integrate the flow differential equations in the two-state limit from Sec. II C. First, we analytically solve Eq. (7), namely $\dot{\theta} = -2\eta$, with $\theta(\tau)$ properly initialized as in our flow. Then we evaluate $\theta(\tau + h)$ and use the result to calculate an effective value, constant for the duration of the step, to use as a corrected generator according to

$$\eta_h = \frac{\theta(\tau) - \theta(\tau + h)}{2h}. \quad (22)$$

This correction is applied to each independent component of the generator separately, as if it were the only active rotation.

In the case of the WWF, this gives

$$\begin{aligned} \eta_h &= \frac{\theta - \tan^{-1}\left(\frac{\tan\theta}{e^{4h^2\eta}}\right)}{2h} \\ &= \frac{\tan^{-1}\left(\frac{j}{x}\right) - \tan^{-1}\left(\frac{j}{x e^{4(x^2+j^2)h}}\right)}{2h}. \end{aligned} \quad (23)$$

The large-radius off-diagonal element instability is prevented because this corrected generator blocks the flow from mistakenly rotating “too far” and changing the sign of j , in the two-state limit. This allows the step sizes to increase as the

flow progresses towards diagonalization and slows down due to the small-radius off-diagonal elements. On the other hand, the unstable Dormand-Prince integrator's step size is bounded from above by the large-radius elements.

For our uniform tangent decay flow, we obtain

$$\eta_h = \frac{\theta - \tan^{-1}\left(\frac{\tan\theta}{e^{4h}}\right)}{2h} = \frac{\tan^{-1}\left(\frac{j}{x}\right) - \tan^{-1}\left(\frac{j}{xe^{4h}}\right)}{2h}. \quad (24)$$

This correction is not truly necessary because the flow itself is already not stiff, but it can still improve performance.

3. Adaptive step sizes

In order to develop this into an adaptive integrator, it is necessary for it to estimate its own error at each step in order to adjust its next step size. Because we can assume that each step is a true unitary rotation, all of the nonrounding errors can be attributed to the change of η during that step and is on the order of $h^2\dot{\eta}$. The error is caused by our failure to predict how η changes over the course of the step, and we wish to limit this relative to the "rotation rate", $\|\eta\|_F$. Therefore, we can calculate the subsequent step size according to

$$h' = \frac{\epsilon h}{n} \frac{\|\eta_h(\tau)\|_F}{\|\eta_0(\tau+h) - \eta'_h(\tau)\|_\infty^{\text{entry}}}, \quad (25)$$

where ϵ is an adjustable tolerance parameter, η_0 is the uncorrected flow generator, and η'_h is our approximation's prediction for the value of η at the end of the step. For the WWF this is given by

$$\eta'_h = \frac{\delta jr^2 e^{4r^2 h}}{x^2 + j^2 e^{8r^2 h}}, \quad (26)$$

and the prediction for the tangent flow is the same, after dropping the three r^2 factors. However, due to its uniform exponential decay, one may find constant step sizes to be a superior choice when implementing the tangent flow.

With most adaptive integrators, it is standard to limit the ratio of each pair of consecutive step sizes so as to prevent them from changing too rapidly. For example, one can require that their ratio lie between one half and two. Furthermore, if h' is less than some fixed fraction of h , such as $3h/4$, one should usually repeat the step with the new, more conservative value.

C. Third order

1. Introduction

While the above integrator performs fairly well, it is only first order and so cannot compete with the higher-order Runge-Kutta algorithms' efficiency during the early stages of the flow. Higher-order algorithms provide better-than-linear returns on precision when decreasing the step size, and so it was desirable for us to devise a third-order extension which retains the nice stability and unitarity properties. Our approach falls into the class of multidervative methods, which means we calculate not only η at each step, but also its first two derivatives. Using these to construct an effective, stabilized generator requires taking multiple types of corrections into account, as explained below.

The techniques in this section are applicable to most well-behaved choices of generator. For simplicity's sake, however, we will focus on the classic WWF and the uniform tangent decay flow from Sec. IV.

2. Integrator step

We begin with the basic WWF equations $\eta = [\mathbf{H}_{\text{Diag.}}, \mathbf{H}]$ and $\dot{\mathbf{H}} = [\dot{\eta}, \mathbf{H}]$. These immediately allow us to calculate that

$$\dot{\eta} = [\dot{\mathbf{H}}_{\text{Diag.}}, \mathbf{H}] + [\mathbf{H}_{\text{Diag.}}, \dot{\mathbf{H}}], \quad (27a)$$

$$\ddot{\mathbf{H}} = [\dot{\eta}, \mathbf{H}] + [\eta, \dot{\mathbf{H}}], \quad (27b)$$

and

$$\ddot{\eta} = [\ddot{\mathbf{H}}_{\text{Diag.}}, \mathbf{H}] + 2[\dot{\mathbf{H}}_{\text{Diag.}}, \dot{\mathbf{H}}] + [\mathbf{H}_{\text{Diag.}}, \ddot{\mathbf{H}}]. \quad (27c)$$

In terms of the elements, $\eta = \delta j$, $\dot{\eta} = \delta \dot{j} + \delta j \dot{}$, and $\ddot{\eta} = \delta \ddot{j} + 2\delta \dot{j} \dot{} + \delta j \ddot{}$. In the two-state limit from Sec. IIC, we could easily calculate a higher-order unitary step operator according to $e^{h\zeta(h)}$, where

$$\zeta(h) = \frac{1}{h} \int_0^h \eta + \dot{\eta}\tau + \frac{\ddot{\eta}}{2}\tau^2 d\tau = \eta + \frac{\dot{\eta}}{2}h + \frac{\ddot{\eta}}{6}h^2 \quad (28)$$

is an effective generator corrected by the derivatives.

However, for larger systems, the noncommutativity of η and its derivatives must also be taken into account according to the Magnus expansion [40,45]: While $\zeta(0) = \zeta_0 = \eta$, and the derivative with respect to the step size h , $\zeta'_0 = \zeta'(0) = \dot{\eta}/2$,

$$\zeta''_0 = \frac{2\ddot{\eta} - [\eta, \dot{\eta}]}{6}. \quad (29)$$

$\zeta(h)$ can now be correctly expressed as a second-order Maclaurin series using these corrected coefficients:

$$\zeta(h) = \zeta_0 + \zeta'_0 h + \frac{\zeta''_0}{2} h^2. \quad (30)$$

In order to finish upgrading to a third-order algorithm, the degree of the Padé approximant must also be increased to (2,2):

$$e^{h\zeta} \approx \frac{12 + 6h\zeta + h^2\zeta^2}{12 - 6h\zeta + h^2\zeta^2}. \quad (31)$$

3. Stabilization

This algorithm is already third order, but without stabilization, the multidervative polynomial extrapolation makes it highly unstable. We suggest stabilizing it through the observation that η decays asymptotically like $e^{-4r^2\tau}$. Thus, we make the ansatz that

$$h\zeta_{r_0}(h) = \int_0^h \frac{(c_0 + c_1\tau + \frac{c_2}{2}\tau^2)}{e^{4r_0^2\tau}} d\tau, \quad (32)$$

where r_0 is the radius at the beginning of the step. Evaluating this integral analytically gives

$$h\zeta_{r_0}(h) = \frac{1}{64r_0^6} \left[16c_0r_0^4 + 4c_1r_0^2 + c_2 - \frac{16c_0r_0^4 + 4c_1r_0^2 + c_2 + (16c_1r_0^4 + 4c_2r_0^2)h + 8c_2r_0^4h^2}{e^{4r_0^2h}} \right]. \quad (33)$$

Matching terms in the corresponding Maclaurin series,

$$\zeta_{r_0}(h) \approx c_0 + \frac{-4c_0r_0^2 + c_1}{2}h + \frac{16c_0r_0^4 - 8c_1r_0^2 + c_2}{6}h^2, \quad (34)$$

with Eq. (30) allowing us to find the three c coefficients to be

$$c_0 = \zeta_0, \quad (35a)$$

$$c_1 = 4r_0^2\zeta_0 + 2\zeta_0', \quad (35b)$$

and

$$c_2 = 16r_0^4\zeta_0 + 16r_0^2\zeta_0' + 3\zeta_0''. \quad (35c)$$

Plugging the above results into Eq. (33) ultimately results in an expression for the stabilized ζ_{r_0} in terms of η and its derivatives:

$$\begin{aligned} h\zeta_{r_0}(h) = & \frac{1}{128r_0^6} \left\{ 96r_0^4\eta + 24r_0^2\dot{\eta} + 2\ddot{\eta} - [\eta, \dot{\eta}] \right. \\ & - e^{-4r_0^2h} [96r_0^4\eta + 24r_0^2\dot{\eta} + 2\ddot{\eta} - [\eta, \dot{\eta}]] \\ & + (256r_0^6\eta + 96r_0^4\dot{\eta} + 8r_0^2\ddot{\eta} - 4r_0^2[\eta, \dot{\eta}])h \\ & \left. + (256r_0^8\eta + 128r_0^6\dot{\eta} + 16r_0^4\ddot{\eta} - 8r_0^4[\eta, \dot{\eta}])h^2 \right\}. \end{aligned} \quad (36)$$

When the exponent $4r_0^2h$ is small, care must be taken to avoid catastrophic cancellation during the evaluation of these corrected expressions. In this limit, an uncorrected Taylor series such as Eq. (30) can be used instead.

4. Adaptive step sizes

The error of each step is now on the order of $h^4\ddot{\eta}$. We suggest simply adding a cube root to Eq. (25) and adjusting the step sizes according to

$$h' = h \left(\frac{\epsilon}{n} \frac{\|\zeta(\tau)\|_F}{\|\eta_0(\tau+h) - \eta_h'(\tau)\|_\infty^{\text{entry}}} \right)^{1/3}, \quad (37)$$

where the predicted generator

$$\eta_h' = \frac{2\eta + (8\eta r_0^2 + 2\dot{\eta})h + (16\eta r_0^4 + 8\dot{\eta}r_0^2 + \ddot{\eta})h^2}{2e^{4r_0^2h}} \quad (38)$$

is the integrand of Eq. (32) when $\tau = h$ without the Magnus correction commutator term in Eq. (29) [46].

5. Alternative flows

This third-order integrator can be easily extended to alternative choices of generator by modifying Eqs. (4), (27a), (27c), and the decay ansatz in Eq. (32). However, the efficiency of the multiderivative approach may be negatively impacted.

For example, to implement the uniform tangent decay flow from Sec. IV, we calculate that $\eta = \sin 2\theta$, $\dot{\eta} = 2\dot{\theta} \cos 2\theta$, and $\ddot{\eta} = 2\ddot{\theta} \cos(2\theta) - 4\dot{\theta}^2 \sin(2\theta)$. These can be expressed in terms of the elements of the Hamiltonian matrix using $\theta = \tan^{-1} j/x$, but the resulting expressions are somewhat lengthy and will be omitted. Additionally, one must modify the decay ansatz by setting r and r_0 to unity in Eqs. (32) through (38).

Note that simulating the tangent flow to third-order precision is a computationally expensive process for many physically inspired initial Hamiltonians. In particular, small-radius level crossings, where x changes sign while $j \approx 0$, involve rapid fluctuations in η which can require very small step sizes to correctly integrate.

D. Results

As shown in Fig. 2, these implementations of the WWF are in fact accurate to their predicted theoretical order. Where applicable, we recommend that the third-order integrator be used for the numerical implementation of CUTs. While we have not judged it as worthwhile, one could consider devising even higher-order extensions.

VI. QUANTIZED TROTTER INTEGRATOR

A. Introduction

Next, we will develop another stable, unitary integrator. While all of the integrators introduced in this paper involve operator splitting [36], this one takes it much further than the previous. We seek an integrator which puts ‘‘just the right amount’’ of computational effort into each pair of off-diagonal

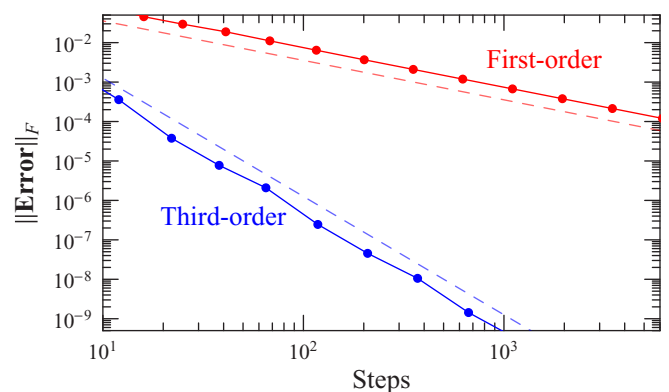


FIG. 2. The third-order integrator requires far fewer steps to converge to the true WWF than the first-order one. The dashed lines show the slope theoretically predicted by their order. These flows were initialized for a half-filled ten-site MBL Hamiltonian with $W = 1$ as explained in Sec. VII and run until $\tau = 1$. The error matrix was estimated by subtracting the resulting $\mathbf{H}(1)$ from the output of a high-precision flow calculation.

elements, in proportion to the rate of rotation between them as set by the generator of the flow.

One lesser-known class of integrators which function along these lines is known as the quantized state system simulators. Instead of basing the construction of the algorithm on the slicing of time into discrete steps as is typically done, these methods quantize the configuration space and then calculate when the system would be expected to switch from one state “cell” to another [19,20]. We derived our specialized integrator with this idea and the Trotter decomposition [21] in mind, while seeking to retain the unitary exactness of the integrators in the previous section. However, due to the noncommutativity of the rotation generators, our partitioning of configuration space into cells is not invariant over the course of the flow.

B. Integrator step

We begin by decomposing η into a real linear combination of the $(n^2 - n)/2$ generators of the $\mathfrak{o}(n)$ Lie algebra. We will denote these generators as $\mathbf{O}_{ab} = \delta_{ac}\delta_{bd} - \delta_{ad}\delta_{bc}$ for all $1 \leq a < b \leq n$. Note that $e^{\theta\mathbf{O}_{ab}}$ represents a rotation of θ radians between the $|a\rangle$ and $|b\rangle$ axes, holding all others fixed.

A trivial decomposition is given by

$$\eta = \sum_{a < b} \eta_{ab} \mathbf{O}_{ab}. \tag{39}$$

Referring back to Eq. (19), the unitary operator which we seek to implement is $e^{h\eta}$. The specific form of the Trotter decomposition which we will base our construction on is the first three factors of the Zassenhaus formula [47]:

$$e^{\epsilon(\mathbf{A}+\mathbf{B})} = e^{\epsilon\mathbf{A}} e^{\epsilon\mathbf{B}} e^{\frac{\epsilon^2}{2}[\mathbf{B},\mathbf{A}]} \dots \tag{40}$$

From this we can immediately write that, for small h ,

$$e^{h\eta} \approx \prod_{a < b} e^{h\eta_{ab} \mathbf{O}_{ab}}, \tag{41}$$

but we can do better. Instead of performing each rotation in the proper proportion at every step, we will instead fix the magnitude of each rotation to be some small angle ι , and perform the rotations in an interleaved manner where each occurs with the proper “frequency” with respect to τ . The choice of a fixed angle is convenient not only because it makes the numerical implementation more efficient, but also because it bounds the magnitude of the dominant error introduced by the Trotter decomposition as given by the final $e^{\epsilon^2[\mathbf{B},\mathbf{A}]/2}$ factor in Eq. (40).

More formally, we can imagine integrating η into an accumulator Σ , which is initially the zero matrix. Again using the entry-wise infinity norm defined in Eq. (16), when $\|\Sigma\|_{\infty}^{\text{entry}}$ grows to reach ι , we can generally say that this was caused by a single pair of elements $\Sigma_{ab} = -\Sigma_{ba} = \pm\iota$. This triggers a Jacobi rotation [48] of magnitude $\mp\iota$ between the $|a\rangle$ and $|b\rangle$ axes. Next, Σ_{ab} and Σ_{ba} are reset to zero and the integration continues until the entry-wise infinity norm again reaches ι . See Fig. 3 for a schematic illustration.

Note that the step sizes are generally distributed in a Poissonian fashion at any given point in the flow. The instantaneous step rate is roughly proportional to the entry-wise one norm of




FIG. 3. In the Trotter integrator, each pair of off-diagonal elements in the accumulator Σ grows at the rate indicated by the corresponding element of η . One can calculate a “hitting time” for each of them when that $\sigma(\tau)$ will reach $\pm\iota$. To perform each step, we find which element will hit next, jump forward to that flow time, perform a Jacobi rotation of magnitude $\mp\iota$ between the appropriate basis axes, and reset that σ to zero. Here, a Jacobi rotation of magnitude $-\iota$ between the $|1\rangle$ and $|4\rangle$ axes is indicated in red, which will then require recalculating all of the shown generators and expected hitting times except for the \mathbf{O}_{23} element.

the generator:

$$\frac{\|\eta\|_1^{\text{entry}}}{\iota} = \frac{1}{\iota} \sum_{a < b} |\eta_{ab}|. \tag{42}$$

C. Stabilization

At the end of the flow, some care must be taken to ensure that ι does not cause us to rotate “too far”: j can be set to zero as in the usual Jacobi eigenvalue algorithm, but it should not be caused to change sign. Alternatively, if we are willing to give up the efficiency boost allowed by the use of a fixed ι , we can ignore this consideration and instead always rotate by a corrected ι_θ value calculated in a manner similar to the corrected generator η_h in Sec. V B 2. For both the WWF and the uniform tangent decay flow from Sec. IV, analytically solving the two-state differential equation gives us a corrected Jacobi rotation angle of

$$\begin{aligned} \iota_\theta &= \frac{\tan^{-1}(\tan \theta e^{4\iota \csc 2\theta}) - \theta}{2} \\ &= \frac{\tan^{-1}\left(\frac{j}{x} e^{2\iota \frac{x^2+j^2}{xj}}\right) - \tan^{-1}\left(\frac{j}{x}\right)}{2}. \end{aligned} \tag{43}$$

This strategy appears to reduce performance slightly.

D. Implementation

The key feature of this algorithm is that the computational expense is approximately proportional to the rotational “path length” of the basis during the course of the flow. This is why it is immune to the features of the WWF which cause certain other integrators to perform poorly.

This approach can be conveniently implemented by noting that η is now piecewise constant, assuming it does not depend explicitly on τ , and so Σ is piecewise linear. Each pair of

off-diagonal elements in Σ has an expected ‘‘hitting time’’ of $\tau + (\text{sgn}(\eta)\iota - \sigma)/\eta$ to reach $\pm\iota$. We can efficiently keep track of these in a bimap-like tree-based data structure, and then read off the next indicated rotation from the first leaf of the tree. After each rotation, only $2n - 3$ generator values and hitting times need to be updated.

If one desires to halt the flow at a particular time τ , it is necessary to iterate until just before that value is surpassed and then ‘‘clear out’’ the queue from the beginning to the end, rotating by only the appropriate fraction of ι each time. In terms of raw performance, we found this algorithm to be inferior to the third-order one in Sec. VC but still far superior to the Dormand-Prince method. Regardless, we feel that it is fundamentally different enough that it may prove useful in future work.

It seems likely that this algorithm could be parallelized in essentially the same manner as the Jacobi eigenvalue algorithm [49]. Additionally, we suspect that higher-order extensions are possible, such as treating η as piecewise linear instead of piecewise constant.

VII. APPLICATION TO MANY-BODY LOCALIZATION

A. The MBL model

We tested the above CUT integrators by applying them to a standard model of one-dimensional MBL on a periodic lattice [13]. The Hamiltonian can be written in terms of second-quantized spinless fermions as

$$\mathbf{H} = \sum_{k=1}^L [\mu_k n_k + V n_k n_{k+1} + (t c_k^\dagger c_{k+1} + \text{H.c.})], \quad (44)$$

where $n_k = c_k^\dagger c_k$, and each μ_k was drawn randomly from the uniform distribution over the range $[-W/2, W/2]$. The three terms correspond to on-site potentials with quenched disorder, interactions of strength $V = 1$, and hopping at the rate $t = 1$, respectively.

As W is increased, this system is believed to undergo a transition from an ergodic extended phase to a nonergodic localized one, possibly passing through an intermediate nonergodic extended phase [50]. Our understanding of these transitions is hampered by finite-size effects controlled by L . This is particularly problematic for the computational study of MBL because the dimensionality of the Hilbert space depends exponentially on L .

We set L to be ten and, as the Hamiltonian preserves the total particle number, restrict our attention to the half-filled sector of the Hilbert space containing five particles. Therefore n , the dimension of our matrices, is $\binom{10}{5} = 252$.

Before beginning the flow, we are free to pre-rotate our Hamiltonian into whatever basis we prefer. We found that the easily calculable single-particle localized basis which diagonalizes the free Hamiltonian with $V = 0$ gives a significant performance boost. Not only does this pre-rotation immediately reduce the ρ metric from Eq. (15) by almost an order of magnitude, it also smooths and accelerates the subsequent flow.

B. Level repulsion metric

In order to visualize the concept of Hilbert space percolation present in the MBL literature [51] using CUTs, we constructed a metric that measures how much ‘‘work’’ the decay of any given pair of off-diagonal elements performs towards inducing the repulsion of the values of the diagonal elements which they connect. Formally, we integrated each x^2 in the two-state limit from Sec. IIC:

$$\begin{aligned} \Xi_{ab} &= \int_0^\infty \frac{d\mathbf{X}_{ab}^2}{d\tau} \Big|_{\eta=\eta_{\mathbf{O}_{ab}}} d\tau \\ &= \int_0^\infty 2\mathbf{X}_{ab} \dot{\mathbf{X}}_{ab} \Big|_{\eta=\eta_{\mathbf{O}_{ab}}} d\tau. \end{aligned} \quad (45)$$

If we were not in the two-state limit, this would simply equal $x_\infty^2 - x_0^2$, where x_∞ is taken from the diagonalized matrix. Instead, Eq. (5b) becomes just $\dot{x} = 2\eta j$, so

$$\Xi = \int_0^\infty 4\eta x j d\tau. \quad (46)$$

For the WWF, which we will use in this section, the final integrand equals $2\eta^2$.

In some sense, Ξ_{ab} measures the total amount of repulsion which has occurred between D_a and D_b due to the decay of \mathbf{J}_{ab} throughout the course of the flow. One can expect that Ξ will be approximately sparse if and only if the system is localized. This is because localized Hamiltonians do not effectively mix most pairs of states. This has the consequence that the eigenvalues in the extended phase obey random matrix, Wigner-Dyson, Gaussian orthogonal ensemble statistics [31], but behave in a Poissonian manner in the localized phase. The intermediate phase is believed to obey some sort of power-law repulsion eigenvalue statistics [50].

Furthermore, assuming the flow succeeds in diagonalizing the Hamiltonian, one can easily show that

$$\sum_{a \neq b} \Xi_{ab} = \sum_{a \neq b} \mathbf{J}_{ab}(\tau = 0)^2 = I_2^J(\tau = 0), \quad (47)$$

the initial value of the metric defined in Eq. (9a). The various phases of localization phenomena correspond analogously to qualitatively distinct economic systems for allocating the scarce initial I_2^J resource amongst the Ξ elements.

It is generally not difficult to extend an integrator to keep track of the matrix Ξ . For example, the third-order WWF integrator from Sec. VC can be extended by symbolically calculating the $\Delta\Xi$ increment integral and then numerically evaluating the resulting expression at every step for each pair of off-diagonal elements using the same decay ansatz and numerical techniques as in Sec. VC 3.

C. Results

We now consider the distribution of the values of $\log \Xi$, treating all $(n^2 - n)/2$ pairs of off-diagonal elements as a single statistical population. As shown in Fig. 4, their distributions are roughly bell shaped. Increasing the disorder strength spreads the distribution to the left. For the strongly disordered cases, it is apparent that a large fraction of the pairs of diagonal elements repelled only trivially during the




FIG. 4. This density plot shows that the distributions of the logarithms of all of the elements of Ξ collected together tend to be roughly bell shaped, with increasing disorder causing them to shift to the left. The low values in the strongly disordered, localized cases reveal the failure of a large number of pairs of diagonal elements to significantly repel during the course of the flow, corresponding to the failure of Hilbert space percolation. The horizontal line indicates the approximate location of the MBL transition, $W_c \approx 5$ [13]. The transition is not sharp due to the strong finite-size effects of a ten-site chain. The curved line shows the medians of the distributions. This data was averaged over four disorder realizations.

course of the flow. This is compatible with the Poissonian level spacing expected in the many-body localized phase. Due to the strong finite-size effects at play in this $L = 10$ system, it is unsurprising that the phase transition is not sharp. We suspect that the nonmonotonicity in the low-disorder behavior is due to the almost preserved translation-invariance causing near degeneracies, highlighting the integrability revealed by the Bethe ansatz. One could consider also tracking the distance or energy dependence of the level repulsion behavior.

VIII. CONCLUSION

A. The interpretation of continuous unitary transformations

Traditionally, we think of CUTs as an inefficient way to diagonalize a matrix [43]. Instead, they occupy a similar conceptual niche as the strong-disorder renormalization group [52–54]. Furthermore, we think of their true utility as providing an intuitive picture of the diagonalization of a Hamiltonian: specifically, a causal reinterpretation of perturbation theory. Starting with an unperturbed Hamiltonian in its diagonal eigenbasis, adding a nontrivial perturbation generates off-diagonal elements. If we then apply a flow to the perturbed Hamiltonian, we can watch as these off-diagonal elements decay in magnitude while repelling the diagonal elements they connect and rotating other off-diagonal elements into each other. While truly small perturbations decay before disturbing the Hamiltonian too significantly, nonperturbative

phenomena can be schematically visualized as ultimately having significant, global consequences over the whole of the matrix.

Several flow generators have been suggested so far. As discussed in Sec. III F, some choices for the infinitesimal flow generator, including the original WWF, can be approximately understood in terms of a renormalization group [3,4]. For example, the WWF preferentially tries to flow away off-diagonal elements which connect levels that already have a large energy difference, δ . It can therefore be thought of as flowing from large energy scales to small ones.

B. Summary

In Sec. IV, we introduced the uniform tangent decay flow. It is quite similar to the original WWF, with the key difference that it does not have this bias towards handling high-energy couplings more quickly. Whether this is beneficial depends on the application. Among the advantages of the tangent flow are that it does not exhibit the stiffness of the WWF: off-diagonal elements decay in synchrony, thereby increasing the efficiency of unstabilized integrators such as the Dormand-Prince method, and it is invariant under energy rescalings. In contrast, halving the energy scale of the WWF quadruples the fictitious flow-time scale. However, the tangent flow's small-radius off-diagonal elements are quite sensitive to the motions of the larger-radius elements, and this, in turn, can slow down numerical integration.

On top of the analytic value of CUTs, they also constitute an important numerical technique [12,55,56]. Thus, it is important that we have efficient algorithms available for the implementation of CUTs *in silico*. However, the existing literature on this topic could be substantially expanded. With this motivation in mind, we have improved upon the existing integrators such as the Dormand-Prince method by presenting three integrators designed specifically for the purpose of CUTs.

In most situations, we recommend taking advantage of the stability, unitary exactness, and third-order convergence properties of the integrator developed in Sec. V C. It should not be difficult to extend it to flows other than the WWF and the tangent flow, so long as the derivatives of the generator can be calculated and the ansatz in Eq. (32) is modified appropriately. The primary advantage of its first-order predecessor in Sec. V B is its relative simplicity, although its stability might also be more robust for some flow generators. The Trotter integrator in Sec. VI does not appear to be immediately optimal for any of our concrete applications, but has intriguing conceptual, stability, and computational complexity properties, and so was included for the sake of completeness.

We took advantage of the third-order integrator and applied it to an MBL problem in Sec. VII in order to explore whether flows provide insight into the one-dimensional MBL transition. In particular, we sought to track the spread of the off-diagonal couplings through the Hamiltonian matrix in order to better understand the asymptotic scaling behavior of a perturbation at its critical intensity. The analysis revealed a clear connection between the distribution of integrated level repulsion strengths

and the ergodicity breaking transition. It also captures the integrability via the Bethe ansatz of the model at zero disorder.

C. Future directions

For future work, we suggest considering whether the unperturbed Hamiltonian provides a natural sense of “distance” between pairs of its eigenstates. If so, one can watch as the presumably initially short-range perturbation “expands” into longer-range interactions during the course of the flow. Phase transitions, including those involving the debated nonergodic extended phase [50,57], should correspond to qualitative changes in the competition between the rate of this interaction expansion, the decay of the off-diagonal elements, and the exponential growth of the number of sites at a given distance. One signature of the nonergodic extended phase should be Ξ distributions at long distances with exponentially small fractions of—but still exponentially many—pairs above some arbitrary repulsion threshold. We hope to apply this analysis to the paradigmatic models of MBL consisting of Anderson localization on a high-dimensional [58] or hyperbolic space, such as a 2^N site hypercube [59] or the Bethe lattice [60,61].

Second, all of the stable fixed points of the flows listed in Table I are diagonal. As mentioned in Sec. I, the resulting permutation of the eigenvectors can be used to construct local integrals of motion for localized Hamiltonians [13,14]. Perhaps flows with a larger set of stable fixed points would provide even more insight into the relevant physics. In particular, we are considering changing the objective function in Eq. (14): For the WWF, it is j^2 , which we suggest modifying to $x^2 j^2$. The resulting generator changes from $\eta =$

$r^2 \sin 2\theta = \delta j$ to $r^4 \sin 4\theta = 4xj(x^2 - j^2)$. The set of stable fixed points is now characterized by block diagonal matrices, possibly permuted, with constant diagonals within each block. Hopefully, the average size of the blocks is proportional to the localization length in the relevant Hilbert space. The stabilizing decay ansatz employed in the construction of the third-order integrator in Sec. V C 3 should not be difficult to generalize to this flow.

A final potential avenue for future research involves constructing the generator η out of something other than just the Hamiltonian \mathbf{H} . In particular, we imagine that a wave function ψ co-evolving according to both a CUT flow and Schrödinger evolution, $\dot{\psi} = (\eta - i\mathbf{H})\psi$, could provide useful information for flows designed to highlight localization phenomena.

ACKNOWLEDGMENTS

This work was supported by the Institute of Quantum Information and Matter, a National Science Foundation frontier center partially funded by the Gordon and Betty Moore Foundation. G.R. acknowledges the generous support of the Packard Foundation and the National Science Foundation through award DMR-1410435. Thanks to E. van Nieuwenburg, S. Kehrein, M. Bintz, C. White, A. Bourzutschky, and P. Titum for many fruitful discussions. The numerical CUT flows were implemented using double-precision floating-point matrices calculated by the open-source linear algebra library ARMADILLO [62]. Special thanks to C. Peng for GPU-accelerating the flows in Sec. V using the open-source linear algebra library ARRAYFIRE [63] and thereby helping to test their performance.

-
- [1] P. Deift, T. Nanda, and C. Tomei, *SIAM J. Num. Anal.* **20**, 1 (1983).
 - [2] F. Wegner, *Ann. Phys.* **506**, 77 (1994).
 - [3] S. D. Glazek and K. G. Wilson, *Phys. Rev. D* **48**, 5863 (1993).
 - [4] S. D. Glazek and K. G. Wilson, *Phys. Rev. D* **49**, 4214 (1994).
 - [5] C. Monthus, *J. Phys. A* **49**, 305002 (2016). (The novel generator introduced in Sec. VII of this paper truncates some terms of the Hamiltonian, preventing them from being generated by the induced flow, which is thus explicitly approximate. We have therefore not considered it as a potential resolution to the instability issue from Sec. III B.)
 - [6] B. H. Bartlett, *Flow Equations for Hamiltonians from Continuous Unitary Transformations* (University of Stellenbosch, Stellenbosch, South Africa, 2003).
 - [7] Y. Itto and S. Abe, *Found. Phys.* **42**, 377 (2012).
 - [8] C. P. Heidbrink and G. S. Uhrig, *Eur. Phys. J. B* **30**, 443 (2002).
 - [9] S. K. Kehrein and A. Mielke, *J. Phys. A* **27**, 4259 (1994).
 - [10] D. M. Basko, I. L. Aleiner, and B. L. Altshuler, *Annals Phys.* **321**, 1126 (2006).
 - [11] R. Nandkishore and D. A. Huse, *Annu. Rev. Condens. Matter Phys.* **6**, 15 (2015).
 - [12] D. Pekker, B. K. Clark, V. Oganesyan, and G. Refael, *Phys. Rev. Lett.* **119**, 075701 (2017).
 - [13] L. Rademaker and M. Ortuño, *Phys. Rev. Lett.* **116**, 010404 (2016).
 - [14] V. L. Quito, P. Titum, D. Pekker, and G. Refael, *Phys. Rev. B* **94**, 104202 (2016).
 - [15] S. Kehrein, *The Flow Equation Approach to Many-Particle Systems* (Springer, Berlin Heidelberg, 2006).
 - [16] J. R. Dormand and P. J. Prince, *J. Comput. Appl. Math.* **6**, 19 (1980).
 - [17] R. I. McLachlan and G. R. W. Quispel, *J. Phys. A* **39**, 5251 (2006).
 - [18] E. Hairer, *BIT Num. Math.* **41**, 996 (2001).
 - [19] F. E. Cellier, E. Kofman, G. Migoni, and M. Bortolotto, *Proc. GCMS*, 504 (2008).
 - [20] G. Migoni, E. Kofman, and F. Cellier, *Simulation* **88**, 387 (2011).
 - [21] H. F. Trotter, *Proc. Am. Math. Soc.* **10**, 545 (1959).
 - [22] Note that this implies that \mathbf{H} is actually symmetric, \mathbf{U} is orthogonal, and η is antisymmetric. Again, extending this work to complex matrices should not be difficult.
 - [23] Decapitalized matrix equations therefore represent the entry-wise Hadamard analog of their bolded versions.
 - [24] The function atan2 has the minor advantage over the usual inverse tangent in that there is no $\pm\pi$ ambiguity.
 - [25] A. M. Steane, [arXiv:1312.3824](https://arxiv.org/abs/1312.3824).
 - [26] S. R. White, *J. Chem. Phys.* **117**, 7472 (2002).

- [27] T. D. Morris, N. M. Parzuchowski, and S. K. Bogner, *Phys. Rev. C* **92**, 034331 (2015).
- [28] M. Hénon, *Phys. Rev. B* **9**, 1921 (1974).
- [29] The Toda flow was rediscovered and introduced to physics by Mielke [64] and approaches the diagonalized form with the eigenvalues sorted in descending order, where all the θ 's are zero. Remarkably, it preserves the structure of the banded matrices to which it is applied [28]. It can be shown to be a discretization of the well-known continuous Korteweg-de Vries nonlinear partial differential equation.
- [30] F. J. Dyson, *J. Math. Phys.* **3**, 1191 (1962).
- [31] M. L. Mehta, *Random Matrices* (Elsevier Science, Cambridge, Massachusetts, 2004).
- [32] J. Robertson, *Coll. Math. J.* **28**, 271 (1997).
- [33] J. H. Seinfeld, L. Lapidus, and M. Hwang, *Ind. Eng. Chem. Fund.* **9**, 266 (1970).
- [34] G. D. Byrne and A. C. Hindmarsh, *J. Comp. Phys.* **70**, 1 (1987).
- [35] M. Hochbruck and A. Ostermann, *Acta Numerica* **19**, 209 (2010).
- [36] R. I. McLachlan and G. R. W. Quispel, *Acta Numerica* **11**, 341 (2002).
- [37] In fact, White originally proposed using the generator $\eta = j/\delta_0$, using the level spacing from the initial Hamiltonian [26]. While sometimes practical, we find this dependence on the starting point of the flow to be conceptually unsatisfying, particularly because δ_0 could well equal zero, and so will not consider it further.
- [38] T. J. Stuchi, *Braz. J. Phys.* **32**, 958 (2002).
- [39] L. Dieci, R. D. Russell, and E. S. V. Vleck, *SIAM J. Num. Anal.* **31**, 261 (1994).
- [40] B. A. Shadwick and W. F. Buell, *Phys. Rev. Lett.* **79**, 5189 (1997).
- [41] N. D. Buono and L. Lopez, *J. Comput. Appl. Math.* **111**, 133 (1999).
- [42] D. J. Higham, *Appl. Num. Math.* **22**, 217 (1996).
- [43] C. Moler and C. V. Loan, *SIAM Rev.* **45**, 3 (2003).
- [44] F. Diele, L. Lopez, and R. Peluso, *Adv. Comput. Math.* **8**, 317 (1998).
- [45] B. A. Shadwick and W. F. Buell, *J. Phys. A* **34**, 4771 (2001).
- [46] Note that we are implicitly assuming that the Magnus correction does not affect the validity of the decay ansatz in Eq. (32).
- [47] W. Magnus, *Commun. Pure Appl. Math.* **7**, 649 (1954).
- [48] C. Kelley, *Iterative Methods for Linear and Nonlinear Equations* (Soc. Indus. & Appl. Math., Philadelphia, Pennsylvania, 1995).
- [49] A. Margaris, S. Souravlas, and M. Roumeliotis, [arXiv:1403.5805](https://arxiv.org/abs/1403.5805).
- [50] M. Serbyn and J. E. Moore, *Phys. Rev. B* **93**, 041424 (2016).
- [51] A. C. Potter, R. Vasseur, and S. A. Parameswaran, *Phys. Rev. X* **5**, 031033 (2015).
- [52] S. K. Ma, C. Dasgupta, and C. K. Hu, *Phys. Rev. Lett.* **43**, 1434 (1979).
- [53] C. Dasgupta and S. K. Ma, *Phys. Rev. B* **22**, 1305 (1980).
- [54] G. Refael and E. Altman, *C. R. Phys.* **14**, 725 (2013).
- [55] B. Fauseweh and G. S. Uhrig, *Phys. Rev. B* **87**, 184406 (2013).
- [56] J. Krones and G. S. Uhrig, *Phys. Rev. B* **91**, 125102 (2015).
- [57] G. Biroli, A. C. Ribeiro-Teixeira, and M. Tarzia, [arXiv:1211.7334](https://arxiv.org/abs/1211.7334).
- [58] E. Tarquini, G. Biroli, and M. Tarzia, *Phys. Rev. B* **95**, 094204 (2017).
- [59] B. Altshuler, H. Krovi, and J. Roland, *Proc. Nat. Acad. Sci. USA* **107**, 12446 (2010).
- [60] R. Abou-Chakra, D. J. Thouless, and P. W. Anderson, *J. Phys. C* **6**, 1734 (1973).
- [61] B. L. Altshuler, L. B. Ioffe, and V. E. Kravtsov, [arXiv:1610.00758](https://arxiv.org/abs/1610.00758).
- [62] C. Sanderson and R. Curtin, *J. Open Source Software* **1**, (2016).
- [63] P. Yalamanchili, U. Arshad, Z. Mohammed, P. Garigipati, P. Entschew, B. Kloppenborg, J. Malcolm, and J. Melonakos, *ArrayFire - A high performance software library for parallel computing with an easy-to-use API* (AccelerEyes, Atlanta, 2015). Retrieved from: <https://github.com/arrayfire/arrayfire>.
- [64] A. Mielke, *Eur. Phys. J. B* **5**, 605 (1998).