

# A preconditioning scheme for minimum energy path finding methods

Cite as: J. Chem. Phys. **150**, 094109 (2019); <https://doi.org/10.1063/1.5064465>

Submitted: 04 October 2018 . Accepted: 13 February 2019 . Published Online: 05 March 2019

Stela Makri , Christoph Ortner , and James R. Kermode 



View Online



Export Citation



CrossMark

# A preconditioning scheme for minimum energy path finding methods

Cite as: *J. Chem. Phys.* **150**, 094109 (2019); doi: [10.1063/1.5064465](https://doi.org/10.1063/1.5064465)

Submitted: 4 October 2018 • Accepted: 13 February 2019 •

Published Online: 5 March 2019



View Online



Export Citation



CrossMark

Stela Makri,<sup>1</sup>  Christoph Ortner,<sup>2</sup>  and James R. Kermode<sup>1</sup> 

## AFFILIATIONS

<sup>1</sup>Warwick Centre for Predictive Modelling, School of Engineering, University of Warwick, CV4 7AL Coventry, United Kingdom

<sup>2</sup>Mathematics Institute, University of Warwick, CV4 7AL Coventry, United Kingdom

## ABSTRACT

Popular methods for identifying transition paths between energy minima, such as the nudged elastic band and string methods, typically do not incorporate potential energy curvature information, leading to slow relaxation to the minimum energy path for typical potential energy surfaces encountered in molecular simulation. We propose a preconditioning scheme which, combined with a new adaptive time step selection algorithm, substantially reduces the computational cost of transition path finding algorithms. We demonstrate the improved performance of our approach in a range of examples including vacancy and dislocation migration modeled with both interatomic potentials and density functional theory.

© 2019 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/1.5064465>

## I. INTRODUCTION

In computational chemistry, structural biology, materials science, and engineering, the time taken for processes is often dominated by transitions between energy minima in a potential energy landscape. The computational evaluation of the Minimum Energy Path (MEP) of the transition is a familiar technique used to find the energy barrier  $\Delta E$  of such a transition.<sup>1</sup> The objective is the evaluation of the transition rate to leading order which is given by  $\nu \sim \nu_0 \exp(-\Delta E/k_B T)$ ,<sup>2,3</sup> where the attempt rate  $\nu_0$  may be estimated using Eyring's heuristic derivation,<sup>2</sup> or approximated with Harmonic Transition State Theory,<sup>4</sup>  $k_B$  is the Boltzmann constant, and  $T$  is the temperature of the system. Knowing the transition rate enables the simulation of the transition on the mesoscale using, for example, the kinetic Monte Carlo method.<sup>5</sup>

We restrict our focus to “double ended” cases where both energy minima are known. The most notable techniques in this case are the string method<sup>6–8</sup> and the Nudged Elastic Band (NEB) method.<sup>9,10</sup> Both methods find the MEP by iteratively relaxing a discretised path, of  $N$  images, until convergence to an approximate MEP is achieved. Typically, the path is evolved in the energy landscape via a steepest descent-like optimisation technique, which

may converge slowly when the potential is ill-conditioned, that is, the Hessian matrix of the potential along the path has a large condition number.<sup>11</sup> Such a situation arises, for example, in large computational domains or if bonds with significant stiffness variations are present. Preconditioning is commonly used in linear algebra and optimisation to effectively reduce the condition number and thus improve the rate of convergence of an iterative scheme.<sup>11</sup>

It has been shown, for example, in Refs. 12–14 how to construct and invert effective preconditioners for the potential energy landscape of materials and molecules at a cost comparable to the evaluation of an interatomic potential and much lower than the cost of evaluating a density functional theory (DFT) model. When used correctly, preconditioning leads to a substantial reduction in the number of force calls and thus is expected to significantly improve computing times.<sup>12,15</sup>

In this paper, we introduce a simple yet effective way to precondition the standard NEB and string methods to obtain efficient and robust algorithms for computing MEPs in ill-conditioned geometries. Our scheme is further enhanced by a novel adaptive step length selection method to improve the robustness of the method. We demonstrate the effectiveness of this combination on a range of material modeling examples.

## II. THE NEB AND STRING METHODS

Let  $\mathbf{x} \in \mathbb{R}^M$ ,  $M \in \mathbb{N}$ , be a state, or configuration, of the dynamical system in question. We denote by  $V(\mathbf{x})$  the potential energy of  $\mathbf{x}$  and assume that  $V$  is twice differentiable and that it has at least two local minima, which we denote by  $\mathbf{x}_A$  and  $\mathbf{x}_B$ , separated by a single saddle point  $\mathbf{x}_S$  of Morse index 1 (to ensure that there is a unique direction of steepest descent at  $\mathbf{x}_S$ <sup>16</sup>). An MEP of the transition from  $\mathbf{x}_A$  to  $\mathbf{x}_B$  is defined as the intrinsically parameterised path  $\mathbf{x}^*(s)$ ,  $s \in [0, 1]$ , satisfying

$$\nabla^\perp V(\mathbf{x}^*) \equiv \mathbf{0}, \quad (1)$$

with end points at the local minima  $\mathbf{x}^*(0) = \mathbf{x}_A$ ,  $\mathbf{x}^*(1) = \mathbf{x}_B$ , where  $\nabla^\perp V(\mathbf{x}) = \left( \mathbf{I} - \frac{\mathbf{x}' \otimes \mathbf{x}'}{\|\mathbf{x}'\|^2} \right) \nabla V(\mathbf{x})$  and where  $\mathbf{x}' = \frac{d\mathbf{x}}{ds}$ . [We note that, strictly speaking,  $\nabla^\perp V$  depends on  $\mathbf{x}'$  as well as  $\mathbf{x}$ , but for the sake of simplicity of notation, we will only write  $\nabla^\perp V(\mathbf{x})$ .] We only present our derivation of preconditioning and numerical tests for the original string method<sup>6</sup> but not the simplified string method,<sup>7</sup> which seems to be used less in practice. However, this is not a fundamental restriction, and we expect no major changes when applying our preconditioning ideas to the simplified string method.

The NEB and string methods discretise a path  $\mathbf{x}(s)$  by interpolating  $N$  discrete points  $\{\mathbf{x}_n\}_{n=1}^N$ . In the present work, we will employ cubic spline interpolation,<sup>17</sup> imposing the “not-a-knot” boundary condition, but the methods we discuss can be readily extended to other interpolation schemes as well.

To evolve the discrete path to equilibrium, we introduce a pseudo-temporal coordinate  $\tau$  and write  $\dot{\mathbf{x}} = \frac{d\mathbf{x}}{d\tau}$ . The evolution of  $\mathbf{x}_n(\tau)$  is then described by the system of ODEs,

$$\dot{\mathbf{x}}_n = -\nabla^\perp V(\mathbf{x}_n) + \boldsymbol{\eta}, \quad (2)$$

where  $\boldsymbol{\eta} = \mathbf{0}$  leads to the string method, while the NEB method introduces elastic interactions between adjacent images along the path by adding the term

$$\boldsymbol{\eta} = \boldsymbol{\eta}_{\text{neb}} = \kappa \left( \mathbf{x}'' \cdot \frac{\mathbf{x}'}{\|\mathbf{x}'\|} \right) \frac{\mathbf{x}'}{\|\mathbf{x}'\|}.$$

System (2) can be solved with any ordinary differential equation (ODE) numerical integrator. Most commonly, Euler’s method<sup>7</sup> is used, which yields an update step of the form

$$\mathbf{x}_n^{k+1} = \mathbf{x}_n^k + \alpha^k \left[ -\nabla^\perp V(\mathbf{x}_n^k) + \boldsymbol{\eta}_n^k \right], \quad (3)$$

where  $\boldsymbol{\eta}_n^k = \boldsymbol{\eta}((\mathbf{x}_n^k)', (\mathbf{x}_n^k)'')$  and  $\alpha^k$  is the time step at iteration  $k$ .

While for NEB, the presence of the elastic interaction  $\boldsymbol{\eta}$  enforces an approximate equidistribution of the nodes along the path, the string method reparameterises the path after each iteration to ensure that the images remain equidistant with respect to a suitable metric. In the continuous limit, as  $N \rightarrow \infty$ , a converged discretised path tends to the correct MEP, independently of the choice of the reparameterisation metric.<sup>8</sup> We initially use the standard  $\ell^2$ -norm defined by  $\|\mathbf{x}\|^2 = \mathbf{x} \cdot \mathbf{x}$ , but we will introduce a different notion of distance later on.

To summarise, the updating relations are given by (3) where, for the string method only, there is an additional redistribution of the images after the update step. We follow precisely the approach described in Eq. (12) in Ref. 7 but for simplicity of presentation do not make this step explicit.

The updating steps Eq. (3) for the string and NEB methods as well as the subsequent analysis were defined in terms of total derivatives of the path variable  $\mathbf{x}$  (i.e., in terms of  $\mathbf{x}'$  and  $\mathbf{x}''$ ), as they are motivated from the respective laws of classical dynamics. This information is available at each iteration at no extra cost as we use cubic spline interpolation to find an expression for  $\mathbf{x}(s)$ .<sup>6,9</sup>

## III. PRECONDITIONING

The NEB and string methods have slow convergence rates when they are subjected to ill-conditioned energy landscapes  $V$ . However, a suitable preconditioner  $\mathbf{P} \in \mathbb{R}^{M \times M}$  that is cheap to compute can be used to reduce the condition number of the Hessian  $\nabla \nabla V$  along the path. In the steepest descent optimisation, preconditioning has related but distinct interpretations: (a) as an approximation of the Hessian,  $\mathbf{P} \approx \nabla \nabla V$ , in analogy to Newton’s scheme or (b) as a coordinate transformation in the state space,  $\mathbf{x} \mapsto \mathbf{P}^{1/2} \mathbf{x}$ , that captures information of the local curvature of the potential landscape (mapping hyperellipsoids to balls).<sup>11</sup>

We will now describe a preconditioning technique for NEB and string methods. The same preconditioners used in geometry optimisation of interatomic potentials<sup>12,13</sup> are expected to be valid for the purposes of preconditioning each image separately. We first present our construction of the preconditioned string method which has a simpler updating step.

### A. Preconditioned string method

Let us first consider the simple case where  $\mathbf{P}$  is constant in  $\mathbf{x}$ . Starting from the coordinate transformation

$$\mathbf{x} \mapsto \mathbf{P}^{-1/2} \mathbf{x} := \tilde{\mathbf{x}}, \quad (4)$$

with corresponding  $\tilde{V}(\tilde{\mathbf{x}}) = V(\mathbf{P}^{1/2} \tilde{\mathbf{x}})$ , it is trivial to deduce that  $\frac{\partial \tilde{x}_i}{\partial x_j} = P_{ij}^{1/2}$ . The string method in the transformed space has the updating step  $\tilde{\mathbf{x}}_n^{k+1} = \tilde{\mathbf{x}}_n^k - \alpha^k \nabla^\perp \tilde{V}(\tilde{\mathbf{x}}_n^k)$  which for convenience we rewrite as

$$\begin{aligned} \tilde{\mathbf{x}}_n^{k+1} &= \tilde{\mathbf{x}}_n^k - \alpha^k \left( \mathbf{I} - \tilde{\mathbf{t}}_n^k \otimes \tilde{\mathbf{t}}_n^k \right) \nabla_{\tilde{\mathbf{x}}} \tilde{V}(\tilde{\mathbf{x}}_n^k), \\ \tilde{\mathbf{t}}_n^k &= \frac{(\tilde{\mathbf{x}}_n^k)'}{\|(\tilde{\mathbf{x}}_n^k)'\|}. \end{aligned} \quad (5)$$

Reversing the coordinate transformation, we obtain an equivalent formulation in the original coordinates with the updating step

$$\begin{aligned} \mathbf{x}_n^{k+1} &= \mathbf{x}_n^k - \alpha^k \left( \mathbf{P}^{-1} - \mathbf{t}_{p,n}^k \otimes \mathbf{t}_{p,n}^k \right) \nabla_{\mathbf{x}} V(\mathbf{x}_n^k), \\ \mathbf{t}_{p,n}^k &= \frac{(\mathbf{x}_n^k)'}{\|(\mathbf{x}_n^k)'\|_{\mathbf{P}}}, \end{aligned} \quad (6)$$

where care needs to be taken to normalise the tangents  $\mathbf{x}'$  with respect to the  $\mathbf{P}$ -norm,  $\|\mathbf{y}\|_{\mathbf{P}} = (\mathbf{y} \cdot \mathbf{P} \mathbf{y})^{1/2}$ , instead of the usual  $\ell^2$ -norm,  $\|\mathbf{y}\| = (\mathbf{y} \cdot \mathbf{y})^{1/2}$ .

Expressing the reparameterisation step in terms of coordinates in the configuration space is trivial, as it suffices to replace the usual  $\ell^2$ -norm with the  $\mathbf{P}$ -norm, due to linearity of the  $\frac{d}{ds}$  operator.

The systems of interest, however, are described by preconditioners that are not constant in the configuration space,<sup>12</sup> which leads to a Riemannian metric framework, and, in particular, the analogue of Eq. (5) involves the evaluation of  $\nabla \mathbf{P}^{1/2}(\mathbf{x}_n^k)$  which is computationally expensive. We circumvent these issues entirely by dropping these terms. Preliminary tests (which we do not discuss here) showed that this does not lead to any loss of performance. Thus, we obtain the preconditioned string method

$$\mathbf{x}_n^{k+1} = \mathbf{x}_n^k - \alpha^k \nabla^\perp V_P(\mathbf{x}_n^k), \quad (7)$$

where we defined the quantity

$$\begin{aligned} \nabla^\perp V_P(\mathbf{x}_n^k) &= \left( [\mathbf{P}_n^k]^{-1} - \mathbf{t}_{p,n}^k \otimes \mathbf{t}_{p,n}^k \right) \nabla_x V(\mathbf{x}_n^k), \\ \mathbf{t}_{p,n}^k &= \frac{(\mathbf{x}_n^k)'}{\|(\mathbf{x}_n^k)'\|_{\mathbf{P}_n^k}}, \end{aligned}$$

in terms of  $\mathbf{P}_n^k = \mathbf{P}(\mathbf{x}_n^k)$ . We are left to specify how to re-parameterise the path. Recall that in the continuous limit, we are free to use any parameterisation for the path. In our setting, the premise is that  $\|\cdot\|_{\mathbf{P}}$  is a more natural notion of distance than the standard  $\ell^2$ -norm  $\|\cdot\|$ ; hence, we will use the following notion of distance along the path:

$$d_P(\mathbf{x}, \mathbf{y}) := \left( (\mathbf{x} - \mathbf{y}) \cdot \left( \frac{\mathbf{P}(\mathbf{x}) + \mathbf{P}(\mathbf{y})}{2} \right) (\mathbf{x} - \mathbf{y}) \right)^{1/2}. \quad (8)$$

We note that  $d_P$  is *not* a metric in the technical sense as it does not satisfy the triangle inequality. However, it is an approximation (discretisation) of the geodesic distance on the Riemannian manifold induced by the preconditioner  $\mathbf{P}$ ; hence, it is reasonable to expect that it can be used for the reparameterisation of the path. In practice, we have not encountered any difficulties related to this issue. The details of the preconditioned reparameterisation algorithm are given in [Appendix](#).

## B. Preconditioned NEB method

An entirely analogous argument yields the preconditioned NEB method,

$$\mathbf{x}_n^{k+1} = \mathbf{x}_n^k + \alpha^k [-\nabla^\perp V_P(\mathbf{x}_n^k) + (\boldsymbol{\eta}_{\text{neb},P})_n^k], \quad (9)$$

where

$$(\boldsymbol{\eta}_{\text{neb},P})_n^k = \kappa \left( (\mathbf{x}_n^k)'' \cdot \mathbf{P}_n^k \frac{(\mathbf{x}_n^k)'}{\|(\mathbf{x}_n^k)'\|_{\mathbf{P}_n^k}} \right) \frac{(\mathbf{x}_n^k)'}{\|(\mathbf{x}_n^k)'\|_{\mathbf{P}_n^k}}.$$

Notice that this class of preconditioning schemes disregards the interactions between images and, therefore, the preconditioner aids the convergence of the path only in the transverse direction. This is justified when the main source of ill-conditioning is due to the potential energy landscape, which is the case when only few images are used as is often done in practice. To summarise, the preconditioned updating relations are given by

$$\mathbf{x}_n^{k+1} = \mathbf{x}_n^k + \alpha^k [-\nabla^\perp V_P(\mathbf{x}_n^k) + (\boldsymbol{\eta}_P)_n^k], \quad (10)$$

where, in analogy to our earlier notation,  $(\boldsymbol{\eta}_P)_n^k = 0$  for the string method and  $(\boldsymbol{\eta}_P)_n^k = (\boldsymbol{\eta}_{\text{neb},P})_n^k$  for NEB.

## C. ODE solvers and steepest descent

The optimisation step Eq. (3) was derived by applying Euler's method to the first order differential equation (2), but any ODE solver can be used instead. Here, we use an adaptive ODE solver based on Ref. 18 to allow for some adaptivity in the step selection mechanism.

The user supplies an absolute and a relative tolerance *atol* and *rtol*, which control the accuracy of the solution. We will demonstrate that choosing these two parameters is more intuitive and more robust than choosing the step length of the static method.

We modify an adaptive ODE solver, *ode12*.<sup>18</sup> To begin, we compute a trial step  $\tilde{\mathbf{x}}_n^{k+1}$  using Eq. (10) with a given step-length  $\alpha^k$ . Next, we use  $\tilde{\mathbf{x}}_n^{k+1}$  to compute a second-order solution to the underlying ODE system via

$$\tilde{\mathbf{x}}_n^{k+1} = \mathbf{x}_n^k + \frac{1}{2} \alpha^k [\mathbf{f}_n^k + \mathbf{f}_n^{k+1}],$$

where  $\mathbf{f}_n^k = -\nabla^\perp V_P(\mathbf{x}_n^k) + (\boldsymbol{\eta}_P)_n^k$  is the driving force on image  $n$  at time step  $k$ . We can then use the difference  $\tilde{\mathbf{x}}_n^{k+1} - \mathbf{x}_n^{k+1}$  or equivalently the difference  $\mathbf{f}_n^k - \mathbf{f}_n^{k+1}$  as an error indicator.

Taking this as a starting point and following, for example, Ref. 19 to implement an adaptive time-stepping algorithm, we obtain an algorithm that underestimates the local error in the neighbourhood of equilibria and, in particular, will not converge as  $k \rightarrow \infty$ . To overcome this, we add a second step-length selection mechanism based on minimising the residual. In essence, the adaptive ODE step selection should be used in the pre-asymptotic regime while minimising the residual is a suitable mechanism in the asymptotic regime.

This leads to the following step-length selection algorithm, which we label *ode12r*: we define the re-scaled residual error

$$R^{k+1} = \max_n \left\| \mathbf{P}_n^k \nabla^\perp V_P(\mathbf{x}_n^k) \right\|_\infty \quad (11)$$

and local error

$$E^{k+1} = \max_{nj} \left\{ \frac{\frac{1}{2} |(\mathbf{f}_n^k - \mathbf{f}_n^{k+1})_j|}{\max\{|(\mathbf{x}_n^k)_j|, |(\mathbf{x}_n^{k+1})_j|, \frac{\text{atol}}{\text{rtol}}\}} \right\},$$


where the index  $j$  denotes vector components. We then accept the proposed  $\mathbf{x}_n^{k+1}$  if the scaled residual error satisfies either one of the two following conditions:

- (1)  $R^{k+1} \leq R^k (1 - c_1 \alpha^k)$ ,
- (2)  $R^{k+1} \leq R^k c_2$  AND  $E^{k+1} \leq \text{rtol}$ ,

for contraction and growth parameters  $c_1$  and  $c_2 \in \mathbb{R}$ .

Whether the step is accepted or rejected, we now compute two step-length candidates using (1) the adaptive solver and (2) a simple line-search procedure.

The step-length candidate given by the *ode12* solver is  $\alpha_{\text{ode12}}^{k+1} = \frac{1}{2} \alpha^k \sqrt{\text{rtol}/E^{k+1}}$ . For the second candidate, we approximate the driving force along the previous search direction by its linear



**FIG. 1.** Convergence rate of the string method applied to vacancy migration in a 249-atom bcc W supercell modeled with the EAM4 potential.<sup>20</sup> Optimal static time stepping, time stepping with *ode12*, and time stepping with *ode12r* were used with a path consisting of 5 images.

interpolant  $(1-\theta)\mathbf{f}_n^k + \theta\mathbf{f}_n^{k+1}$ . We then minimise  $\|(1-\theta)\mathbf{f}_n^k + \theta\mathbf{f}_n^{k+1}\|_{\mathbf{p}_n}^2$  with respect to  $\theta$  to obtain  $\alpha_{\text{ls}}^{k+1} = \theta\alpha^k$ .

If the current step  $\mathbf{x}^{k+1}$  is accepted, then the next step-length candidate is chosen to be

$$\alpha^{k+1} = \max\left(\frac{1}{4}\alpha^k, \min\left(4\alpha^k, \alpha_{\text{ls}}^{k+1}, \alpha_{\text{ode12}}^{k+1}\right)\right).$$

If the step  $\mathbf{x}^{k+1}$  is rejected, then the new step-length candidate starting from  $\mathbf{x}^k$  is

$$\alpha^k = \max\left(\frac{1}{10}\alpha^k, \min\left(\frac{1}{4}\alpha^k, \alpha_{\text{ls}}^{k+1}, \alpha_{\text{ode12}}^{k+1}\right)\right).$$

Figure 1 demonstrates how *ode12* effectively selects appropriate step lengths in the pre-asymptotic regime but stagnates in the asymptotic regime for the case of vacancy migration in tungsten modeled with the EAM4 class of the Embedded Atom Model (EAM) interatomic potential proposed by Marinica *et al.*<sup>20</sup> The convergence rate of the modified *ode12r* agrees with the results of *ode12* in the pre-asymptotic regime but successfully converges upon reaching the asymptotic regime.

## IV. RESULTS

We tested our preconditioning scheme for a variety of examples. First, we looked at examples using interatomic potentials which are not the main target as these are typically fast models and constructing a preconditioner may not be computationally efficient relative to force evaluations. These examples, however, demonstrate how the number of force evaluations can be reduced with the use of the preconditioner. Further fine-tuning the preconditioner implementation and application (e.g., our current implementation updates the preconditioner after each iteration, which could be avoided), one would still obtain significant practical speed-ups for severely ill-conditioned cases.

We then compare with a density functional theory (DFT) model to confirm our earlier results. In the following tables, we compare the number of force evaluations per image needed to converge to “coarse” and “fine” target accuracies (maximum force less than  $10^{-1}$  eV/Å and  $10^{-3}$  eV/Å, respectively) using unpreconditioned and preconditioned schemes with either static or adaptive *ode12r* step selection. The criterion for convergence is the magnitude of the residual error  $R^{k+1}$  as defined in Eq. (11). For the use of the *ode12r*

step selection, fitting the *rtol* and *atol* parameter was simple as it was observed that *rtol* = 0.1 was sufficient in most cases for convergence but other values *rtol* = 1 and *rtol* = 0.01 were occasionally more appropriate. The value of *atol* was chosen so that *atol/rtol* = 1 in all cases except the 2D vacancy of Sec. IV A, where *atol/rtol* = 0.01 had to be used instead.

### A. Vacancy migration

First, we consider the diffusion of a vacancy in a two dimensional 60-atom triangular lattice governed by a Lennard-Jones potential  $V(r) = 4\epsilon[(\sigma/r)^{12} - 2(\sigma/r)^6]$  with parameters  $\epsilon = 1.0$ ,  $\sigma = 2^{-\frac{1}{6}}$ . The vacancy is located at the centre of the cell initially and migrates in the *y* direction by one lattice spacing. Periodic boundary conditions are imposed in the *x* and *y* directions. Table I shows the number of force calls per image required for convergence. The exponential preconditioner (Exp) introduced in Packwood *et al.*<sup>12</sup> with parameters  $A = 3.0$  and  $r_{\text{cut}} = 2.5$ , which utilises bond-connectivity information to treat the ill-conditioning of the system allowed convergence beyond the  $10^{-3}$  tolerance, which the unpreconditioned case could not achieve within a reasonable number of iterations. The latter came as a surprise to us as on the contrary to the real vacancy migration systems that we study next, this artificial setup exhibits more severe ill-conditioning. We note that for the unpreconditioned case when using the *ode12r* time stepping for the string method, we had to use *atol/rtol* = 0.01. The absolute differences  $\|\mathbf{x}_1 - \mathbf{x}_2\|_{\infty}$  between the positions  $\mathbf{x}_1$  and  $\mathbf{x}_2$  for the image nearest the saddle in converged paths with and without preconditioning, were of the order of  $8 \times 10^{-3}$ .

Next, we considered a three dimensional system containing a vacancy, specifically a 107-atom Cu fcc supercell in a fixed cell with periodic boundary conditions. Interactions were modeled with a Morse potential with parameters  $A = 4.0$ ,  $\epsilon = 1.0$  and the nearest neighbour distance  $r_0 = 2.55$  Å with interactions between atoms expressed by  $V(r) = \epsilon(e^{-2A(r/r_0-1)} - 2e^{-A(r/r_0-1)})$ . The exponential preconditioner introduced in Packwood *et al.*<sup>12</sup> was used with parameters  $A = 3.0$  and  $r_{\text{cut}} = 2.2r_0 = 5.62$  Å. Table II shows the number of force evaluations per image needed for convergence to two preset tolerance limits. This example demonstrates how the *ode12r* solver can aid the performance of the string and NEB methods if

**TABLE I.** Number of force evaluations per image required by the string and NEB methods to converge the vacancy migration MEP in a 9 image path of a 60-atom 2D cell modeled with a Lennard-Jones potential, with either the static or *ode12r* step length selection methods. In the cases marked \*, the algorithm did not converge within a reasonable number of iterations.

Step selection	2D Vacancy			
	Static		<i>ode12r</i> solver	
	$10^{-1}$	$10^{-3}$	$10^{-1}$	$10^{-3}$
String	197	*	52	*
String (p)	16	38	12	33
NEB	200	*	53	*
NEB (p)	19	60	14	67

**TABLE II.** Force evaluations per image needed for the string and NEB methods for the migration of a vacancy in a 107-atom Cu fcc supercell modeled by a Morse potential. The MEP was discretised with 5 images.

Step selection	Vacancy in Cu supercell			
	Static		ode12r solver	
	$10^{-1}$	$10^{-3}$	$10^{-1}$	$10^{-3}$
String	8	74	8	41
String (p)	7	38	8	21
NEB	8	57	8	27
NEB (p)	7	37	8	19

a static step is not suitable. Preconditioning gave almost a 2-fold speedup for the higher accuracy results but no improvement for the lower accuracy. The absolute differences of the positions of the converged paths at the saddle, as performed before, were well below  $3 \times 10^{-14}$  Å.

A 53-atom W bcc supercell modeled with the EAM4 potential described in Ref. 20 was examined as well. Periodic boundary conditions were imposed. A *force field* (FF) *preconditioner* was constructed by suitably modifying the EAM Hessian to enforce positivity; see the work of Mones *et al.*<sup>13</sup> (p. 9) for full details. This yields up to 6 times faster convergence for higher accuracies as shown in Table III. The absolute differences of the positions of the converged paths at the saddle were well below  $5 \times 10^{-9}$  Å.

We studied the same 53-atom W vacancy system with density functional theory (DFT) as implemented in the Castep<sup>21</sup> software. The exchange correlation functional was approximated by the Perdew, Burke, and Ernzerhof (PBE) generalised gradient approximation (GGA),<sup>22</sup> with a plane wave energy cutoff of 500 eV and a  $2 \times 2 \times 2$  Monkhorst-Pack grid to sample the Brillouin zone (a comparison of convergence behavior obtained with a  $3 \times 3 \times 3$  k-point grid was carried out which showed that the use of the  $2 \times 2 \times 2$  k-point grid is sufficient). Step selection with *ode12r* step and static step selection schemes was studied. A regularised FF preconditioner based on the EAM Hessian was used,  $\mathbf{P} = (1 - \lambda)\mathbf{P}_{\text{FF}} + \lambda\mathbf{P}_{\text{Exp}} + c\mathbf{I}$ ,

**TABLE III.** Force evaluations per image needed for the string and NEB methods to converge the MEP for vacancy migration in a 53-atom W bcc supercell modeled by the EAM4 potential.<sup>20</sup> The path was discretised by 5 images, and the preconditioner was constructed from the force field.<sup>13</sup>

Step selection	Vacancy in W supercell			
	Static		ode12r solver	
	$10^{-1}$	$10^{-3}$	$10^{-1}$	$10^{-3}$
String	7	77	7	49
String (p)	5	12	5	9
NEB	8	58	7	35
NEB (p)	5	10	8	17


where  $c = 0.05$ ,  $\lambda = 0.4$ ,  $\mathbf{P}_{\text{FF}}$  is described in Ref. 13, p. 9, and the  $\mathbf{P}_{\text{Exp}}$  parameters were fitted to  $\mathbf{P}_{\text{FF}}$ .


The path is made up of 5 images, and traversing the path in subsequent iterations of the NEB and string methods was performed in an alternating order, allowing efficient reuse of previous electronic structure data to start the next optimisation step.

Unlike the EAM case above, the preconditioner we used for the DFT model does not describe the potential energy surface of the DFT model exactly but nevertheless gives a speed-up of a factor of two for an accuracy of  $\sim 10^{-2}$  eV/Å and furthermore allows accuracies of the order of  $\sim 10^{-3}$  eV/Å to be achieved, unlike the unpreconditioned case, as shown in Figs. 2 and 3. The results of Table III suggest that constructing a better preconditioner would improve these results further. Notice further that the number of force evaluations needed for convergence and the time needed for convergence are in agreement (by comparison of the upper and lower panes of Figs. 2 and 3), confirming that the computational cost of constructing the preconditioner model is negligible compared to the cost of computing DFT forces, justifying our earlier assumptions. We note that the gain of preconditioning would be expected to further increase with the system size.<sup>12</sup> The absolute differences of the positions of the converged paths at the saddle were of the order of  $1 \times 10^{-4}$  Å.

## B. Screw dislocation

In the final example, we study a  $\frac{1}{2}\langle 111 \rangle$  screw dislocation in a 562-atom W bcc structure confined in a cylinder of radius equal to 20 Å and surrounded by an 11 Å cylindrical shell of clamped atoms, with periodic boundary conditions along the dislocation line ( $z$ ) direction. The system is simulated with the same EAM4 potential. The dislocation advances by one glide step. Table IV shows the

**FIG. 2.** Convergence of the string and NEB methods with and without preconditioner for a 53-atom bcc W supercell containing a vacancy and modeled with DFT. The upper panel (a) shows the error as a function of the number of force evaluations per image and the lower (b) as a function of the time required to converge. Time stepping with *ode12r* was used with a path of 5 images. Comparison shows that constructing and evaluating the preconditioner is negligible compared to the cost of force computation.



**FIG. 3.** Convergence of the string and NEB methods with and without preconditioning for a 53-atom W bcc supercell containing a vacancy and modeled with DFT. The upper panel (a) shows the error as a function of the number of force evaluations per image and the lower panel (b) shows the error as a function of the time required to converge. The static time step was chosen by extrapolating the *ode12r* data. The path was discretised by 5 images.


computational costs for converging the MEP with the NEB and string methods, using either static or *ode12r* step length selection. A force field preconditioner built from the same EAM potential was used for geometry optimisation.

Upon preconditioning, we observed a 5-fold speed up for the static case for low accuracies but only a 2-fold speed up for the *ode12r* case. For a higher accuracy, a speed up of a factor of 6 was observed and there was a speed up of a factor of at least 2 from using the *ode12r* step selection over the static step selection for both the unpreconditioned and preconditioned cases. This indicates that the fitted static step is only suitable in the pre-asymptotic regime and a larger step size is suitable in the asymptotic regime, showcasing the advantages of using the adaptive *ode12r* scheme over the hand-tuned static step. The absolute differences of the positions of the converged paths at the saddle were below  $2 \times 10^{-3}$  Å.

**TABLE IV.** Computational cost for the NEB and string methods for a screw dislocation in a 562-atom W bcc cylinder simulated with the EAM4 Marinica potential.<sup>20</sup> The circular boundary is fixed at a radius of  $R = 20$  Å. Periodic boundary conditions were imposed in the  $z$  direction. The path was discretised by 9 points.

Step selection	Screw dislocation			
	Static		<i>ode12r</i> solver	
	$10^{-1}$	$10^{-3}$	$10^{-1}$	$10^{-3}$
String	40	272	14	124
String (p)	7	48	9	21
NEB	40	312	14	162
NEB (p)	7	47	7	21

We investigated this system further, focusing on the NEB implementation to allow comparison with the widely used Limited memory Broyden—Fletcher—Goldfarb—Shanno (LBFGS)<sup>23</sup> optimisation algorithm, which can be used with the NEB implementation<sup>10</sup> in the Atomic Simulation Environment (ASE).<sup>24</sup> This required fixing the end points of the path at the minima as is performed in the ASE code. The comparison was carried out on systems of two sizes. A force field preconditioner was used as before for the preconditioned cases. Figure 4 shows the convergence rate of the various NEB schemes for a radius of 20 Å in the upper panel (a) and for a radius of 40 Å in the lower panel (b). Note that although LBFGS gave good convergence in the unpreconditioned case, it lacks robustness. This is because the force field of the NEB algorithm is not conservative, violating one of LBFGS's assumptions. LBFGS constructs a Hessian matrix corresponding to a scalar field, failing to capture the effects of the transport terms of the NEB force field. Moreover, the lack of the energy function prevents the use of line search, required to ensure the method's stability; in the ASE LBFGS implementation, a heuristic is instead used to impose a maximum step length of 0.04 Å. Furthermore, it should be noted that because our preconditioning scheme does not treat the longitudinal force components, it is inappropriate for



**FIG. 4.** Convergence of NEB variants for a screw dislocation in a 562-atom W bcc cylindrical structure (a) and a 1489-atom W bcc cylindrical structure (b) modeled with the EAM4 Marinica potential.<sup>20</sup> Atoms outside outer radii of  $R = 20$  Å and  $R = 40$  Å, respectively, were clamped with periodic boundary conditions along the dislocation line. The path was discretised with 7 images (excluding the minima at each end, which were held fixed). The horizontal axis of the plots was cut after 160 force evaluations per image to focus on the performance of the preconditioned schemes. The static unpreconditioned NEB method converged after 312 force evaluations per image for the  $R = 20$  Å case and after 343 force evaluations per image for the  $R = 40$  Å case.

us to use it together with the LBFGS method for MEP finding methods.

## V. CONCLUSIONS

We have demonstrated that MEP finding techniques such as the NEB and the string method can exhibit slow convergence rates due to poor search direction and step-length selection during the optimisation procedure. We have introduced a new optimisation technique combining an adaptive time-stepping scheme with preconditioning to address ill-conditioning of the energy landscape in directions transverse to the path and to allow faster convergence to the minimum energy path.

We observed that our new scheme gives a significant speed up and improved robustness over currently used approaches for a range of systems using both force fields and DFT. Moreover, it allows higher accuracies to be reached than existing methods.

However, our preconditioning scheme targets transverse ill-conditioning only. The longitudinal terms (e.g., the NEB spring interactions) are unaffected by the preconditioner, suggesting that our scheme provides a baseline for further improvements.

An open source prototype implementation of our technique is available at <https://github.com/cortner/SaddleSearch.jl>.

## ACKNOWLEDGMENTS

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) under Grant Nos. EP/P002188/1, EP/R012474/1, EP/J021377/1, and EP/R043612/1, by ERC Starting Grant No. 335120, and by the Royal Society under Grant No. RG160691. Computing facilities were provided by the Scientific Computing Research Technology Platform of the University of Warwick with support from the Science Research Investment Fund. We thank Petr Grigorev for providing the screw dislocation configurations.

## APPENDIX: REPARAMETERISING IN PRECONDITIONED STRING

The path reparameterisation described in Eq. (12) in Ref. 7 assumes that the  $\ell^2$ -metric is used to measure distance. Here, we briefly describe the modifications required when it is replaced with the metric  $d_p$  defined in (8), used in the preconditioned string method introduced in Sec. III A.

After accepting an optimisation step  $k$  of Eq. (7), the following steps are performed:

1. Compute the relative distances  $d_p(\mathbf{x}_n^k, \mathbf{x}_{n-1}^k)$  between the images  $\{\mathbf{x}_n^k\}_n$ , for all  $n = 2, \dots, N$ .
2. Define

$$s_1 = 0, \\ s_n = \frac{\sum_{m=2}^n d_p(\mathbf{x}_m^k, \mathbf{x}_{m-1}^k)}{\sum_{m=2}^N d_p(\mathbf{x}_m^k, \mathbf{x}_{m-1}^k)}, \text{ for } n = 2, \dots, M. \quad (\text{A1})$$

3. Use cubic spline interpolation<sup>17</sup> of  $\{s_n, \mathbf{x}_n^k\}_{n=1}^N$  to obtain  $\mathbf{x}^k(s) : [0, 1] \rightarrow \mathbb{R}^N$ .
4. The new images are then given by

$$\mathbf{x}_n^k = \mathbf{x}^k\left(\frac{n-1}{N-1}\right), \quad n = 1, \dots, N. \quad (\text{A2})$$

This algorithm does *not* ensure that images will be equidistributed according to  $d_p$ . However it does ensure that images remain bounded away from one another, which is the key property required for the string method.

## REFERENCES

- <sup>1</sup>A. F. Voter, F. Montalenti, and T. C. Germann, *Annu. Rev. Mater. Res.* **32**, 321 (2002).
- <sup>2</sup>H. Eyring, *J. Comput. Phys.* **3**, 107 (1935).
- <sup>3</sup>E. Pollak and P. Talkner, *Chaos* **15**, 026116 (2005).
- <sup>4</sup>G. H. Vineyard, *J. Phys. Chem. Solids* **3**, 121 (1957).
- <sup>5</sup>A. F. Voter, *Radiation Effects in Solids*, NATO Science Series Vol. 235 (Springer Netherlands, 2007), p. 1.
- <sup>6</sup>W. E, W. Ren, and E. Vanden-Eijnden, *Phys. Rev. B* **66**, 052301 (2002).
- <sup>7</sup>W. E, W. Ren, and E. Vanden-Eijnden, *J. Chem. Phys.* **126**, 164103 (2007).
- <sup>8</sup>M. Cameron, R. V. Kohn, and E. Vanden-Eijnden, *J. Nonlinear Sci.* **21**, 193 (2011).
- <sup>9</sup>H. Jónsson, G. Mills, and K. W. Jacobsen, *Classical and Quantum Dynamics in Condensed Phase Simulations* (World Scientific Publishing Co. Pte. Ltd., 1998), p. 385.
- <sup>10</sup>G. Henkelman and H. Jónsson, *J. Chem. Phys.* **113**, 9978 (2000).
- <sup>11</sup>J. Nocedal and S. J. Wright, *Springer Series in Operations Research and Financial Engineering* (Springer, Berlin, 2006).
- <sup>12</sup>D. Packwood, J. Kermodé, L. Mones, N. Bernstein, J. Woolley, N. Gould, C. Ortner, and G. Csányi, *J. Chem. Phys.* **144**, 164109 (2016).
- <sup>13</sup>L. Mones, C. Ortner, and G. Csányi, *Sci. Rep.* **8**, 13991 (2018).
- <sup>14</sup>R. Lindh, A. Bernhardtsson, G. Karlström, and P.-Å. Malmqvist, *Chem. Phys. Lett.* **241**, 423 (1995).
- <sup>15</sup>M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, *Rev. Mod. Phys.* **64**, 1045 (1992).
- <sup>16</sup>W. E and X. Zhou, *Nonlinearity* **24**, 1831 (2011).
- <sup>17</sup>P. Dierckx, *Curve and Surface Fitting with Splines* (Oxford University Press, Inc., New York, NY, USA, 1993).
- <sup>18</sup>E. Hairer, S. P. Nørsett, and G. Wanner, *Solving Ordinary Differential Equations I: Nonstiff Problems, 2nd rev. ed.* (Springer-Verlag New York, Inc., New York, NY, USA, 1993).
- <sup>19</sup>H. Lamba, *BIT Numer. Math.* **40**, 314 (2000).
- <sup>20</sup>M.-C. Marinica, L. Ventelon, M. R. Gilbert, L. Proville, S. L. S. L. Dudarev, J. Marian, G. Bencteux, and F. Willaime, *J. Phys.: Condens. Matter* **25**, 395502 (2013).
- <sup>21</sup>S. Clark, M. Segall, C. Pickard, P. Hasnip, M. Probert, K. Refson, and M. Payne, *Z. Kristallogr.-Cryst. Mater.* **220**, 567 (2005).
- <sup>22</sup>J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).
- <sup>23</sup>D. C. Liu and J. Nocedal, *Math. Program.* **45**, 503 (1989).
- <sup>24</sup>A. H. Larsen, J. J. Mortensen, J. Blomqvist, I. E. Castelli, R. Christensen, M. Dulak, J. Friis, M. N. Groves, B. Hammer, C. Hargus, E. D. Hermes, P. C. Jennings, P. B. Jensen, J. Kermodé, J. R. Kitchin, E. L. Kolsbjerg, J. Kubal, K. Kaasbjerg, S. Lysgaard, J. B. Maronsson, T. Maxson, T. Olsen, L. Pastewka, A. Peterson, C. Rostgaard, J. Schiøtz, O. Schütt, M. Strange, K. S. Thygesen, T. Vegge, L. Vilhelmsen, M. Walter, Z. Zeng, and K. W. Jacobsen, *J. Phys.: Condens. Matter* **29**, 273002 (2017).