

## Stabilized quasi-Newton optimization of noisy potential energy surfaces

Bastian Schaefer, S. Alireza Ghasemi, Shantanu Roy, and Stefan Goedecker

Citation: *The Journal of Chemical Physics* **142**, 034112 (2015); doi: 10.1063/1.4905665

View online: <http://dx.doi.org/10.1063/1.4905665>

View Table of Contents: <http://scitation.aip.org/content/aip/journal/jcp/142/3?ver=pdfcov>

Published by the [AIP Publishing](#)

---

### Articles you may be interested in

[A constrained reduced-dimensionality search algorithm to follow chemical reactions on potential energy surfaces](#)

*J. Chem. Phys.* **138**, 214102 (2013); 10.1063/1.4807743

[Fast numerical design of spatial-selective rf pulses in MRI using Krotov and quasi-Newton based optimal control methods](#)

*J. Chem. Phys.* **137**, 054203 (2012); 10.1063/1.4739755

[Quasi-Newton parallel geometry optimization methods](#)

*J. Chem. Phys.* **133**, 034116 (2010); 10.1063/1.3455719

[Methods for finding transition states on reduced potential energy surfaces](#)

*J. Chem. Phys.* **132**, 234110 (2010); 10.1063/1.3445772

[Interpolating moving least-squares methods for fitting potential energy surfaces: Using classical trajectories to explore configuration space](#)

*J. Chem. Phys.* **130**, 144107 (2009); 10.1063/1.3111261

---



**AIP** | The Journal of  
Chemical Physics

### Meet The New Deputy Editors

	<b>Peter Hamm</b>		<b>David E. Manolopoulos</b>		<b>James L. Skinner</b>
---	-------------------	---	------------------------------	---	-------------------------

# Stabilized quasi-Newton optimization of noisy potential energy surfaces

Bastian Schaefer,<sup>1</sup> S. Alireza Ghasemi,<sup>2</sup> Shantanu Roy,<sup>3</sup> and Stefan Goedecker<sup>1,a)</sup>

<sup>1</sup>*Department of Physics, University of Basel, Klingelbergstrasse 82, CH-4056 Basel, Switzerland*

<sup>2</sup>*Institute for Advanced Studies in Basic Sciences, P.O. Box 45195-1159, IR-Zanjan, Iran*

<sup>3</sup>*Computational and Systems Biology, Biozentrum, University of Basel, CH-4056 Basel, Switzerland*

(Received 2 December 2014; accepted 28 December 2014; published online 16 January 2015)

Optimizations of atomic positions belong to the most commonly performed tasks in electronic structure calculations. Many simulations like global minimum searches or characterizations of chemical reactions require performing hundreds or thousands of minimizations or saddle computations. To automatize these tasks, optimization algorithms must not only be efficient but also very reliable. Unfortunately, computational noise in forces and energies is inherent to electronic structure codes. This computational noise poses a severe problem to the stability of efficient optimization methods like the limited-memory Broyden–Fletcher–Goldfarb–Shanno algorithm. We here present a technique that allows obtaining significant curvature information of noisy potential energy surfaces. We use this technique to construct both, a stabilized quasi-Newton minimization method and a stabilized quasi-Newton saddle finding approach. We demonstrate with the help of benchmarks that both the minimizer and the saddle finding approach are superior to comparable existing methods. © 2015 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4905665>]

## I. INTRODUCTION

Stationary points are the most interesting and most important points of potential energy surfaces. The relative energies of local minima and their associated configuration space volumes determine thermodynamic equilibrium properties.<sup>1</sup> According to transition state theory, dynamical properties can be deduced from the energies and the connectivity of minima and transition states.<sup>2</sup> Therefore, the efficient determination of stationary points of potential energy surfaces is of great interest to the communities of computational chemistry, physics, and biology. Clearly, optimization and, in particular, minimization problems are present in virtually any field. This explains why the development and mathematical characterization of iterative optimization techniques are important and longstanding research topics, which resulted in a number of highly sophisticated methods like, for example, direct inversion of the iterative subspace (DIIS),<sup>3,4</sup> conjugate gradient (CG),<sup>5</sup> or quasi-Newton methods like the Broyden–Fletcher–Goldfarb–Shanno (BFGS) algorithm<sup>6–9</sup> and its limited memory variant (L-BFGS).<sup>10,11</sup> Since for a quadratic function, Newton's method is guaranteed to converge within a single iteration, it is not surprising that the BFGS and L-BFGS algorithms belong to the most efficient methods for minimizations of atomic systems.<sup>1</sup>

If the potential energy surface can be computed with an accuracy on the order of the machine precision, the above mentioned algorithms usually work extremely well. In practice, however, computing the energy surface at this high precision is not possible for physically accurate but computationally demanding levels of theory like, for example, density functional theory (DFT). At DFT level, this is due to the finitely

spaced integration grids and self consistency cycles that have to be stopped at small, but non-vanishing thresholds. Therefore, optimization algorithms that are used at these accurate levels of theory must not only be computationally efficient but also tolerant of noise in forces and energies. Unfortunately, the very efficient L-BFGS algorithm is known to be noise-sensitive and, therefore, frequently fails to converge on noisy potential energy surfaces. For this reason, the fast inertial relaxation engine (FIRE) has been developed.<sup>12</sup> FIRE is a method of the damped molecular dynamics (MD) class of optimizers.<sup>13,14</sup> It accelerates convergence by mixing the velocity at every MD step with a fraction of the current steepest descent direction. A great advantage of FIRE is its simplicity. However, FIRE does not make use of any curvature information and, therefore, usually is significantly less efficient than the Newton or quasi-Newton methods.

Potential energy surfaces are bounded from below, and therefore, descent directions guarantee that a local minimum will finally be found. Furthermore, the curvature at a minimum is positive in all directions. This means, all directions can be treated on the same footing during a minimization. The situation is different for saddle point optimizations. A saddle point is a stationary point at which the potential energy surface is at a maximum with respect to one or more particular directions, and at a minimum with respect to all other directions. Close to a saddle point, it is therefore not possible to treat all directions on the same footing. Instead, one has to single out the directions that have to be maximized. Furthermore, far away from a saddle point, it is usually impossible to tell, which search direction guarantees to finally end up in a saddle point. Therefore, saddle point optimizations typically are more demanding and significantly less reliable than minimizations.

In this contribution, we present a technique that allows to extract curvature information from noisy potential energy

<sup>a)</sup>stefan.goedecker@unibas.ch

surfaces. We explain how to use this technique to construct a stabilized quasi-Newton minimizer (SQNM) and a stabilized quasi-Newton saddle (SQNS) finding method. Using benchmarks, we demonstrate that both optimizers are robust and efficient. The comparison of SQNM to L-BFGS and FIRE and of SQNS to an improved dimer method<sup>15,16</sup> reveals that SQNM and SQNS are superior to their existing alternatives.

## II. METHODS

### A. Newton's and quasi Newton's method

The potential energy surface of an  $N$ -atomic system is a map  $E: \mathbb{R}^{3N} \mapsto \mathbb{R}$  that assigns to each atomic configuration  $\mathbf{R}$  a potential energy. It is assumed that a second order expansion of  $E(\mathbf{R})$  about a point  $\mathbf{R}^i$  is possible

$$E(\mathbf{R}) \approx E(\mathbf{R}^i) + [\mathbf{R} - \mathbf{R}^i]^T \nabla E(\mathbf{R}^i) + \frac{1}{2} [\mathbf{R} - \mathbf{R}^i]^T H_{\mathbf{R}^i} [\mathbf{R} - \mathbf{R}^i], \quad (1)$$

$$\nabla E(\mathbf{R}) \approx \nabla E(\mathbf{R}^i) + H_{\mathbf{R}^i} [\mathbf{R} - \mathbf{R}^i]. \quad (2)$$

Here,  $H_{\mathbf{R}^i}$  is the Hessian of the potential energy surface evaluated at  $\mathbf{R}^i$ . If  $\mathbf{R}$  is a stationary point, the left hand side gradient of Eq. (2) vanishes and Newton's optimization method follows

$$\mathbf{R}^{i+1} = \mathbf{R}^i - H_{\mathbf{R}^i}^{-1} \nabla E(\mathbf{R}^i). \quad (3)$$

In the previous equation,  $\mathbf{R}$  was renamed to  $\mathbf{R}^{i+1}$  in order to emphasize the iterative character of Newton's method for non-quadratic potential energy surfaces.

In practice, it is in most cases either impossible to calculate an analytic Hessian or it is too time consuming to compute it numerically by means of finite differences at every iteration. Therefore, quasi-Newton methods use an approximation to the exact Hessian that is computationally less demanding. Using a constant multiple of the identity matrix as an approximation to the Hessian results in the simple steepest descent method. In most cases, such a choice is a very poor approximation to the true Hessian. However, improved approximations can be generated from local curvature information which is obtained from the history of the last  $n_{\text{hist}}$  displacements  $\Delta \mathbf{R}^i := \mathbf{R}^i - \mathbf{R}^{i-1}$  and gradient differences  $\Delta \mathbf{g}^i := \nabla E(\mathbf{R}^i) - \nabla E(\mathbf{R}^{i-1})$ , where  $i = 1, \dots, n_{\text{hist}}$ .

### B. Significant subspace in noisy optimization problems

In noisy optimization problems, the noisy components of the gradients can lead to displacement components that correspond to erratic movements on the potential energy surface. Consequently, curvature information that comes from the subspace spanned by these displacement components must not be used for the construction of an approximate Hessian. In contrast to this, the non-noisy gradient components promote locally systematic net-movements, which do not tend to cancel each other. In this sense, the displacement components that correspond to these well defined net-movement span a significant subspace from which meaningful curvature information can be extracted and used for building an approximate Hessian.

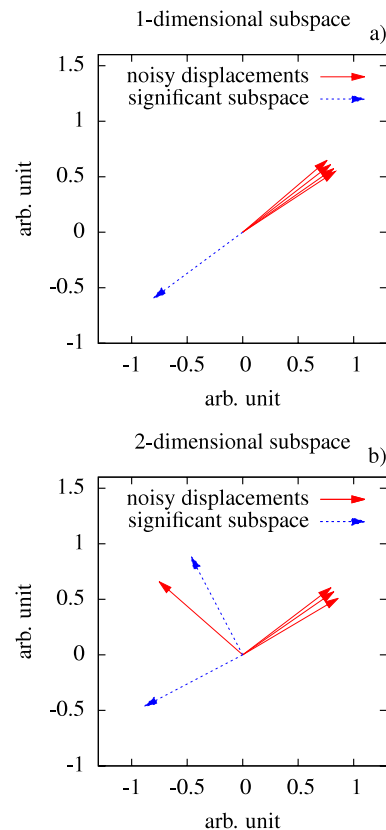


FIG. 1. Illustrated are significant subspaces spanned by the displacements in a model atomic coordinate space. Only from the significant subspace, it is meaningful to extract curvature information. The red solid arrows simulate displacements made under the influence of noisy forces. The blue dashed arrows show significant subspaces from which it is meaningful to extract curvature information. Panel (a) shows a case in which the significant subspace is only one-dimensional. Panel (b) shows an example in which curvature information can be extracted from the full 2-dimensional space. The significant subspaces that are shown here were computed using the method outlined in Sec. II B.

The situation is depicted in Fig. 1 where the red solid vectors represent the history of normalized displacements and the blue dashed vectors constitute a basis of the significant subspace. All the red solid vectors in Fig. 1(a) point into similar directions. Therefore, curvature information should only be extracted from a one-dimensional subspace, as, for example, is given by the blue dashed vector. Displacement components perpendicular to this blue dashed vector come from the noise in the gradients. In contrast to Fig. 1(a), Fig 1(b) shows a displacement that points into a considerably different direction than all the other displacements. For this reason, significant curvature information can be extracted in the full two-dimensional space.

To define the significant subspace more rigorously, we first introduce the set of normalized displacements

$$\widehat{\Delta \mathbf{R}}^i := \frac{\Delta \mathbf{R}^i}{|\Delta \mathbf{R}^i|}, \quad (4)$$

where  $i = 1, \dots, n_{\text{hist}}$ . With  $\sum_k |\omega_k|^2 = 1$ , linear combinations  $\mathbf{w}$  of the normalized displacements are defined as

$$\mathbf{w} := \sum_{k=1}^{n_{\text{hist}}} \omega_k \widehat{\Delta \mathbf{R}}^k. \quad (5)$$

Furthermore, we define a real symmetric overlap matrix  $S$ , as

$$S_{kl} := \widehat{\Delta \mathbf{R}}^k \cdot \widehat{\Delta \mathbf{R}}^l. \quad (6)$$

It can be seen from

$$\mathbf{w} \cdot \mathbf{w} = \boldsymbol{\omega}^T S \boldsymbol{\omega}, \quad (7)$$

that  $|\mathbf{w}|$  is made stationary by coefficient vectors  $\boldsymbol{\omega}^i$  that are eigenvectors of the overlap matrix. In particular, the longest and shortest vectors that can be generated by linear combinations with normalized coefficient vectors  $\boldsymbol{\omega}$  correspond to those eigenvectors of the overlap matrix that have the largest and smallest eigenvalues. As motivated above, the shortest linear combinations of the normalized displacements correspond to noise.

From now on, let the  $\boldsymbol{\omega}^i$  be eigenvectors of  $(S_{kl})$  and let  $\lambda_i$  be the corresponding eigenvalues. With

$$\widetilde{\Delta \mathbf{R}}^i := \frac{1}{\sqrt{\lambda_i}} \sum_{k=1}^{n_{\text{hist}}} \omega_k^i \widehat{\Delta \mathbf{R}}^k, \quad (8)$$

we finally define the *significant subspace*  $\mathfrak{S}$  as

$$\mathfrak{S} := \text{span} \left( \left\{ \widetilde{\Delta \mathbf{R}}^i \mid \lambda_i / \max_j \{\lambda_j\} > \epsilon \right\} \right), \quad (9)$$

where  $0 \leq \epsilon \leq 1$ . In all applications presented in this work,  $\epsilon = 10^{-4}$  has proven to work well. Henceforth, we will refer to the dimension of  $\mathfrak{S}$  as  $n_{\text{dim}}$ . By construction, it is guaranteed that  $n_{\text{dim}} \leq 3N$ . It should be noted that at each iteration of the optimization algorithms that are introduced below, the significant subspace and its dimension  $n_{\text{dim}}$  can change. The history length  $n_{\text{hist}}$  usually lies between 5 and 20.

Our procedure is analogous to Löwdin's canonical orthogonalization,<sup>17-19</sup> which is used in the electronic structure community to remove linear dependencies from chemical basis sets.

### C. Obtaining curvature information on the significant subspace

We define the projection  $\widetilde{H}$  of the Hessian  $H$  onto  $\mathfrak{S}$  as

$$\begin{aligned} \widetilde{H} &:= PHP \\ &= \sum_{ij} H_{ij} \widetilde{\Delta \mathbf{R}}^i \left( \widetilde{\Delta \mathbf{R}}^j \right)^T, \end{aligned} \quad (10)$$

where for all  $\widetilde{\Delta \mathbf{R}}^i \in \mathfrak{S}$ ,

$$P := \sum_{i=1}^{n_{\text{dim}}} \widetilde{\Delta \mathbf{R}}^i \left( \widetilde{\Delta \mathbf{R}}^i \right)^T$$

and

$$H_{ij} := \left( \widetilde{\Delta \mathbf{R}}^i \right)^T H \widetilde{\Delta \mathbf{R}}^j.$$

Using Eq. (2) and defining

$$\widetilde{\Delta \mathbf{g}}^i := \frac{1}{\sqrt{\lambda_i}} \sum_{k=1}^{n_{\text{hist}}} \frac{\omega_k^i}{|\Delta \mathbf{R}^k|} \Delta \mathbf{g}^k, \quad (11)$$

where  $i = 1, \dots, n_{\text{dim}}$ , one obtains an approximation for each matrix element  $H_{ij}$ ,

$$H_{ij} \approx \widetilde{\Delta \mathbf{g}}^i \cdot \widetilde{\Delta \mathbf{R}}^j. \quad (12)$$

In practice, we explicitly symmetrize  $H_{ij}$  in order to avoid asymmetries introduced by anharmonic effects

$$H_{ij} \approx \frac{1}{2} \left( \widetilde{\Delta \mathbf{g}}^i \cdot \widetilde{\Delta \mathbf{R}}^j + \widetilde{\Delta \mathbf{g}}^j \cdot \widetilde{\Delta \mathbf{R}}^i \right). \quad (13)$$

Because the projection  $P$  is the identity operator on  $\mathfrak{S}$ , the curvature  $c(\widehat{\mathbf{d}})$  on the potential energy surface along a normalized  $\widehat{\mathbf{d}} \in \mathfrak{S}$  is given by

$$c(\widehat{\mathbf{d}}) = \widehat{\mathbf{d}}^T \widetilde{H} \widehat{\mathbf{d}}. \quad (14)$$

Given the normalized eigenvectors  $\mathbf{v}^i$  and corresponding eigenvalues  $\kappa_i$  of the  $n_{\text{dim}} \times n_{\text{dim}}$  matrix  $(H_{ij})$ , one can write normalized eigenvectors  $\widetilde{\mathbf{v}}^i \in \mathfrak{S}$  of  $\widetilde{H}$  with eigenvalues  $\kappa_i$  as

$$\widetilde{\mathbf{v}}^i = \sum_{k=1}^{n_{\text{dim}}} \mathbf{v}_k^i \widetilde{\Delta \mathbf{R}}^k, \quad (15)$$

where  $\mathbf{v}_k^i$  is the  $k$ -th element of  $\mathbf{v}^i$ . As can be seen from Eq. (14), the  $\kappa_i$  give the curvatures of the potential energy surface along the directions  $\widetilde{\mathbf{v}}^i$ .

### D. Using curvature information on the significant subspace for preconditioning $\nabla E$

The gradient  $\nabla E$  can be decomposed into a component lying in  $\mathfrak{S}$  and a component lying in its orthogonal complement

$$\nabla E = \nabla E_{\mathfrak{S}} + \nabla E_{\perp}, \quad (16)$$

where  $\nabla E_{\mathfrak{S}} := P' \nabla E$ ,  $\nabla E_{\perp} := (I - P') \nabla E$ , and  $P' := \sum_i \widetilde{\mathbf{v}}^i \left( \widetilde{\mathbf{v}}^i \right)^T$ . In this section, we motivate how the  $\kappa_i$  can be used to precondition  $\nabla E_{\mathfrak{S}}$ . Furthermore, we explain how  $\nabla E_{\perp}$  can be scaled appropriately with the help of a feedback that is based on the angle between two consecutive gradients.

Let us assume that the Hessian  $H$  at the current point of the potential energy surface is non-singular and let  $\nu_i$  and  $\mathbf{V}^i$  be its eigenvalues and normalized eigenvectors. In Newton's method (Eq. (3)), the gradients are conditioned by the inverse Hessian. For the significant subspace component  $\nabla E_{\mathfrak{S}}$ , it follows

$$H^{-1} \nabla E_{\mathfrak{S}} = \sum_{i=1}^{3N} \sum_{j=1}^{n_{\text{dim}}} \left[ \left( \frac{\nabla E \cdot \widetilde{\mathbf{v}}^j}{\nu_i} \right) \left( \widetilde{\mathbf{v}}^j \cdot \mathbf{V}^i \right) \mathbf{V}^i \right]. \quad (17)$$

As outlined in Sec. II C, we know the curvature  $\kappa_j$  along  $\widetilde{\mathbf{v}}^j$ . Therefore, at a first thought, Eq. (17) suggests to simply replace  $\nu_i$  by  $\kappa_j$ , where  $i = 1, \dots, 3N$  and  $j = 1, \dots, n_{\text{dim}}$ . Indeed, if the optimization was restricted to the subspace  $\mathfrak{S}$ , this choice would be appropriate. However, with respect to the complete domain of the potential energy surface, one is at risk to underestimate the curvature  $\nu_i$  if the overlap  $O_{ij} := \widetilde{\mathbf{v}}^j \cdot \mathbf{V}^i$  is non-vanishing.

In particular, if  $O_{ij}$  is far from being negligible, underestimating the curvature  $\nu_i$  can be particularly problematic because coordinate changes in the direction of  $\mathbf{V}^i$  might be too large. This can render convergence difficult to obtain in practice.

We therefore replace  $v_i$  in Eq. (17) by

$$\kappa'_j := \sqrt{\kappa_j^2 + r_j^2}, \quad (18)$$

where  $r_j$  is chosen in analogy to the residue of Weinstein's criterion<sup>20,21</sup> as

$$r_j := \left| H\tilde{\mathbf{v}}^j - \left( (\tilde{\mathbf{v}}^j)^T H\tilde{\mathbf{v}}^j \right) \tilde{\mathbf{v}}^j \right|. \quad (19)$$

Using Eqs. (11), (14), and (15), this residue can be approximated by

$$r_j \approx \left| \sum_{k=1}^{n_{\text{dim}}} \left[ \mathbf{v}_k^i \Delta \tilde{\mathbf{g}}^k \right] - \kappa_j \tilde{\mathbf{v}}^j \right|. \quad (20)$$

With this choice for  $\kappa'_j$ , the preconditioned gradient  $\nabla E_{\mathfrak{E}}^{\text{P}}$  is finally given by

$$\nabla E_{\mathfrak{E}}^{\text{P}} := \sum_{j=1}^{n_{\text{dim}}} \left( \frac{\nabla E \cdot \tilde{\mathbf{v}}^j}{\kappa'_j} \right) \tilde{\mathbf{v}}^j. \quad (21)$$

Clearly, the residue  $r_j$  can only alleviate the problem of curvature underestimation, but it does not rigorously guarantee that every single  $v_i$  is estimated appropriately. However, in practice, this choice works very well. The reason for this can be seen from Fig. 2. In Fig. 2(a), a histogram of the quality and safety measure  $q_{ij} := \sqrt{\kappa_j^2 + r_j^2} - v_i$  is shown. If  $q_{ij} < 0$ , the curvature  $v_i$  is underestimated; if  $q_{ij} \approx 0$ , the curvature  $v_i$  is well estimated; and finally, if  $q_{ij} > 0$ , the curvature is overestimated. Overestimation leads to too small step sizes and, therefore, to a more stable algorithm, albeit at the cost of a performance loss. Critical underestimation of the curvature ( $q_{ij} \ll 0$ ) is rare. Fig. 2(b) shows the averages of the overlap  $O_{ij}$  in the corresponding bins. If  $\tilde{\mathbf{v}}^j$  has on average a large overlap with  $\mathbf{V}^i$ , the curvature along  $\mathbf{V}^i$  is estimated accurately (histogram in Fig. 2(a) peaks at  $q_{ij} \approx 0$ ).




FIG. 2. Panel (a) is a histogram of  $q_{ij} := \sqrt{\kappa_j^2 + r_j^2} - v_i$  for  $i = 1, \dots, 3N$  and  $j = 1, \dots, n_{\text{dim}}$ .  $q_{ij}$  is a measure for the quality of the estimation of the eigenvalue  $v_i$  of the exact Hessian. Panel (b) shows the bin-averaged overlap  $O_{ij}$ . The frequency of severe curvature underestimation drops quickly in the region  $q_{ij} < 0$ . The histogram in panel (a) peaks in the region of good estimation ( $q_{ij} \approx 0$ ) which coincidences with the region of large overlap  $O_{ij}$ , shown in panel (b). The data for this figure come from 100 minimizations of a Si<sub>20</sub> system described by the Lenosky-Silicon<sup>22,23</sup> force field.

What remains to be discussed is how the gradient component  $\nabla E_{\perp}$  should be scaled. By construction,  $\nabla E_{\perp}$  lies in the subspace for which no curvature information is available. We therefore treat this gradient component by a simple steepest descent approach that adjusts the step size  $\alpha > 0$  at each iteration. For the minimizer that is outlined in Sec. II F, the adjustment is based on the angle between the complete gradient  $\nabla E$  and the preconditioned gradient  $\nabla E^{\text{P}}$ . If the cosine of this intermediate angle is larger than 0.2,  $\alpha$  is increased by a factor of 1.1, otherwise  $\alpha$  is decreased by a factor of 0.85. For the saddle search algorithm, the feedback is slightly different and will be explained in Sec. II G. The exact numbers for the scaling factors were determined heuristically. The only constraints are that the scaling factors must increase the step size if the complete gradient and the preconditioned gradient point into similar directions and decrease the step size otherwise. According to our experience, the above choices offer a good efficiency.

In conclusion, the total preconditioned gradient  $\nabla E^{\text{P}}$  is given by

$$\nabla E^{\text{P}} := \nabla E_{\mathfrak{E}}^{\text{P}} + \alpha \nabla E_{\perp}. \quad (22)$$

In Sec. II E, we explain how this preconditioned gradient can be further improved for biomolecules.

The preconditioned subspace gradient  $\nabla E_{\mathfrak{E}}^{\text{P}}$  was obtained under the assumption of a quadratic potential energy surface. However, if the gradients at the current iteration are large, this assumption is probably not satisfied. Displacing along  $\nabla E_{\mathfrak{E}}^{\text{P}}$  in these cases can reduce the stability of the optimization. Hence, if the  $|\nabla E|$  exceeds a certain threshold, it can be useful to set the dimension of  $\mathfrak{E}$  to zero for a certain number of iterations. This means that  $\nabla E_{\perp} = \nabla E$  and, therefore,  $\nabla E^{\text{P}} = \alpha \nabla E$ . In that case,  $\alpha$  is also adjusted according to the above described gradient feedback. As this fallback to steepest descent is intended as a last final fallback, it should have the ability to deal with arbitrarily large forces. Therefore, we also check that  $\alpha \nabla E$  does not displace any atom by more than a user-defined trust radius. However, to our experience, this fallback is not necessary in most cases. Indeed, all the benchmarks presented in Secs. III were performed without this fallback.

## E. Additional efficiency for biomolecules

Many large molecules like biomolecules or polymers are floppy systems in which the largest and smallest curvatures can be very different from each other. Steepest descent optimizers are very inefficient for these ill-conditioned systems, because the high curvature directions force to use step sizes that are far too small for an efficient optimization in the directions of small curvatures. Put more formally, the optimization is inefficient for those systems, because the condition number, which is the fraction of largest and smallest curvature, is large.<sup>24</sup> For biomolecules, the high-curvature directions usually correspond to bond stretchings, that is, movements along inter-atomic displacement vectors of bonded atoms. For the current purpose, we regard two atoms to be bonded if their inter-atomic distance is smaller than or equal to 1.2 times the sum of their covalent radii. For  $i = 1, \dots, N$ , let  $\mathbf{r}^i \in \mathbb{R}^3$  be the coordinate vector of the  $i$ -th atom. For a system with  $n_{\text{bond}}$  bonds, we define

for each bond a bond vector  $\mathbf{b}^m \in \mathbb{R}^{3N}$ ,  $m = 1, \dots, n_{\text{bond}}$ ,

$$\mathbf{b}^m := \begin{pmatrix} \tilde{\mathbf{b}}^{\sim m,1} \\ \mathbf{b} \\ \tilde{\mathbf{b}}^{\sim m,2} \\ \vdots \\ \tilde{\mathbf{b}}^{\sim m,N} \\ \mathbf{b} \end{pmatrix}, \quad (23)$$

where the  $\tilde{\mathbf{b}}^{\sim m,k} \in \mathbb{R}^3$ ,  $k = 1, \dots, N$ , are defined as

$$\tilde{\mathbf{b}}^{\sim m,i} := -\tilde{\mathbf{b}}^{\sim m,j} := \begin{cases} \mathbf{r}^j - \mathbf{r}^i, & \text{if atoms } i \text{ and } j \text{ are} \\ & \text{bonded by the } m\text{-th bond,} \\ \mathbf{0}, & \text{otherwise.} \end{cases} \quad (24)$$

The  $\mathbf{b}^m$  are sparse vectors with six non-zero elements.

We separate the total gradient  $\nabla E$  into its bond-stretching components  $\nabla E_{\text{str}}$  and all the remaining components  $\nabla E_r$ ,

$$\nabla E = \nabla E_{\text{str}} + \nabla E_r. \quad (25)$$

Let  $c_m \in \mathbb{R}$  be coefficients that allow the bond-stretching components to be expanded in terms of the bond vectors

$$\nabla E_{\text{str}} := \sum_{m=1}^{n_{\text{bond}}} c_m \mathbf{b}^m. \quad (26)$$

Using definition Eq. (26), left-multiplying Eq. (25) with a bond vector  $\mathbf{b}^n$  and requiring the  $\nabla E_r$  to be orthogonal to all the bond vectors, one obtains the following linear system of equations, which determines the coefficients  $c_m$  and, with it, the bond stretching gradient defined in Eq. (26):

$$\mathbf{b}^n \cdot \nabla E = \sum_m c_m \mathbf{b}^n \cdot \mathbf{b}^m. \quad (27)$$

For the optimization of a biomolecule, the bond-stretching components are minimized in a simple steepest descent fashion. The atoms are displaced by  $-\alpha_s \nabla E_{\text{str}}$ . The bond-stretching step size  $\alpha_s$  is a positive number, which is adjusted in each iteration of the optimization by simply counting the number of projections  $\mathbf{b}^m \cdot \nabla E$  that have not changed signs since the last iteration. If more than two thirds of the signs of the projections have remained unchanged, the bond-stretching step size  $\alpha_s$  is increased by 10%. Otherwise,  $\alpha_s$  is decreased by a factor of 1/1.1. The non-bond-stretching gradients  $\nabla E_r$  are preconditioned using the stabilized quasi-Newton approach presented in Secs. II B—II D. It is important to note that in Secs. II B—II D, all  $\nabla E$  have to be replaced by  $\nabla E_r$  when using this biomolecule preconditioner. In particular, this is also true for the gradient feedbacks that are described in Secs. II D and II G.

## F. Finding minima—The SQNM method

The pseudocode below demonstrates how the above presented techniques can be assembled into an efficient and stabilized quasi-newton minimizer (SQNM). The pseudo code contains 4 parameters explicitly.  $\alpha_{\text{start}}$  and  $\alpha_{s,\text{start}}$  are initial step sizes that scale  $\nabla E_{\perp}$  and  $\nabla E_{\text{str}}$ , respectively.  $m$  is the maximum length of the history list from which the significant subspace  $\mathcal{S}$  is constructed.  $E_{\text{thresh}}$  is an energy-threshold that is used to determine whether a minimization step is accepted or not. It

should be adapted to the noise level of the energies and forces. The history list is discarded if the energy increases, because an increase in energy is an indication for inaccurate curvature information. In this case, the dimension of the significant subspace is considered to be zero. Furthermore, line 17 implicitly contains the parameter  $\epsilon$ , which is described in Sec. II B. The optimization is considered to be converged if the norm of the gradient is smaller than a certain threshold value. Of course, other force criteria like, for example, using the maximum force component instead of the force norm, are possible.

```

1.  $\alpha \leftarrow \alpha_{\text{start}}$ ;  $\alpha_s \leftarrow \alpha_{s,\text{start}}$ ;
2. accepted  $\leftarrow$  true;
3.  $k \leftarrow 1$ ;
4. Initialize  $\mathbf{R}_k$  with coordinates;
5.  $E_k \leftarrow E(\mathbf{R}_k)$ ;
6. repeat
7.   if optimizing biomolecule then
8.     if accepted then
9.       Compute  $\nabla E_{\text{str}}$  for  $\mathbf{R}_k$ , as outlined in Sec. II E;
10.      Adjust  $\alpha_s$  based on the feedback described in Sec. II E;
11.       $\mathbf{g}_k \leftarrow \nabla E(\mathbf{R}_k) - \nabla E_{\text{str}}$ ;
12.       $\mathbf{R}_k \leftarrow \mathbf{R}_k - \alpha_s \nabla E_{\text{str}}$ ;
13.    end if
14.  else
15.     $\mathbf{g}_k \leftarrow \nabla E(\mathbf{R}_k)$ ;
16.  end if
17.  Based on the  $\{\mathbf{g}_j, \mathbf{R}_j\}_{j \leq k}$  in the history list, compute the
    preconditioned gradient  $\nabla E^{\text{P}}$  as outlined in Secs. II B—II E;
18.   $\mathbf{R}_{k+1} \leftarrow \mathbf{R}_k - \nabla E^{\text{P}}$ ;
19.  if  $E(\mathbf{R}_{k+1}) > E_k + E_{\text{thresh}}$  and  $\alpha > \alpha_{\text{start}}/10$  then
20.    accepted  $\leftarrow$  false;
21.    Remove  $\{\mathbf{g}_j, \mathbf{R}_j\}_{j < k}$  from the history list;
22.     $\alpha \leftarrow \alpha/2$ ;
23.  else
24.    accepted  $\leftarrow$  true;
25.     $E_{k+1} \leftarrow E(\mathbf{R}_{k+1})$ ;
26.    Adjust  $\alpha$  based on the gradient feedback described in
    Sec. II D;
27.    if  $k > m$  then
28.      Remove  $\mathbf{R}_{k-m}$  and  $\mathbf{g}_{k-m}$  from storage;
29.    end if
30.     $k \leftarrow k+1$ ;
31.  end if
32. until convergence.

```

## G. Finding saddle points—The SQNS method

In this section, we describe a stabilized quasi-Newton saddle finding method (SQNS) that is based on the same principles as the minimizer in Sec. II F. SQNS belongs to the class of the minimum mode following methods.<sup>15,25,26</sup>

For simplicity, we will denote the Hessian eigenvector corresponding to the smallest eigenvalue as minimum mode. Broadly speaking, a minimum mode following method maximizes along the direction of the minimum mode and it minimizes in all other directions. The optimization is considered to be converged if the curvature along the minimum mode is negative and if the norm of the gradient is smaller than a certain threshold. As for the minimization, other force criteria are possible.

The minimum mode of the Hessian can be found by minimizing the curvature function  $c: \mathbb{R}^{3N} \mapsto \mathbb{R}$ ,

$$c(\mathbf{d}) = \frac{\mathbf{d}^T \mathbf{H} \mathbf{d}}{\mathbf{d}^T \mathbf{d}} \approx \frac{\Delta \mathbf{g} \cdot \Delta \mathbf{R}}{h^2}, \quad (28)$$

where along with  $h \ll 1$ , the following definitions were used:  $\Delta \mathbf{R} := h \frac{\mathbf{d}}{|\mathbf{d}|}$  and  $\Delta \mathbf{g} := \nabla E(\mathbf{R} + \Delta \mathbf{R}) - \nabla E(\mathbf{R})$ . The vector  $\mathbf{R}$  is the position at which the Hessian  $H$  is evaluated at. For the minimization of  $c(\mathbf{d})$ , we use the algorithm described in Sec. II F where the energy as objective function is replaced by  $c(\mathbf{d})$ . In the pseudocode below, the here discussed minimization is done at line 6. Under the constraint of normalization, the gradient  $\nabla c(\mathbf{d})|_{|\mathbf{d}|=1}$  is given by

$$\nabla c(\mathbf{d})|_{|\mathbf{d}|=1} = 2(\mathbf{H} \mathbf{d} - c(\mathbf{d}) \mathbf{d}) \approx 2 \left( \frac{\Delta \mathbf{g}}{h} - \left( \frac{\Delta \mathbf{g} \cdot \Delta \mathbf{R}}{h^3} \right) \Delta \mathbf{R} \right). \quad (29)$$

Blindly using the biomolecule preconditioner of Sec. II E for the minimization of  $c(\mathbf{d})$  would mean that the gradient of Eq. (29) was projected on the bond vectors of  $\mathbf{d}$ . Obviously, the bond vector as defined in Sec. II E has no meaning for  $\mathbf{d}$ . Therefore, Eq. (29) instead is projected onto the bond vectors of  $\mathbf{R} + \Delta \mathbf{R}$ .

At a stationary point, systems with free boundary conditions have six vanishing eigenvalues. The respective eigenvectors correspond to overall translations and rotations.<sup>1</sup> Instead of directly using Eq. (29) for the minimization of the curvature of those systems, it is advantageous to remove the translations and rotations from  $\Delta \mathbf{R}$  and  $\nabla c(\mathbf{d})|_{|\mathbf{d}|=1}$  in Eq. (29).<sup>1,27,28</sup>

The convergence criterion for the minimization of  $c(\mathbf{d})$  has a large influence on the total number of energy and force evaluations needed to obtain convergence. It therefore must be chosen carefully.

The minimum mode is usually not computed at every iteration, but only if one of the following conditions is fulfilled:

1. at the first iteration of the optimization,
2. if the integrated length of the optimization path connecting the current point in coordinate space and the point at which the minimum mode has been calculated previously exceeds a given threshold value  $r_{\text{recomp}}$ ,
3. if the curvature along the minimum mode is positive, and the curvature has not been recomputed for at least  $n_{\text{recomp}}$  iterations,
4. if the curvature along the minimum mode is positive and the norm of the gradient falls below the convergence criterion,
5. at convergence (optional).

In the pseudocode, these conditions are checked in line 5. Among these conditions, condition no. 2 is, with respect to the performance, the most important one. The number of energy and gradient evaluations needed for converging to a saddle point can be strongly reduced if a good value for  $r_{\text{recomp}}$  is chosen. Conditions 3 and 4 can be omitted for most cases. However, for some cases, they can offer a slight reduction in the number of energy and gradient evaluations. For example, for the alanine dipeptide system used in Sec. III, these two

conditions offered a performance gain of almost 10%. Although possible, we usually do not tune  $n_{\text{recomp}}$ , but typically use  $n_{\text{recomp}} = 10$ . In our implementation, condition 5 is optional. It can be used if very accurate directions of the minimum mode at the saddle point are needed. In this case, this last minimum mode computation can also be done at a tighter convergence criterion. Additional energy and gradient computations are saved in our implementation by using the previously computed minimum mode as the starting mode for a new curvature minimization.

As stated above, a saddle point is found by maximizing along the minimum mode and minimizing in all other directions. This is done by inverting the preconditioned gradient component that is parallel to the minimum mode. This is shown at line 19 of the pseudocode below. For the case of biomolecules, the component of the bond-stretching gradient that is parallel to the minimum mode is also inverted (line 13). As already mentioned in Sec. II D, the feedback that adjusts the stepsize of  $\nabla E_{\perp}$  is slightly different in case of the saddle finding method. Let  $\hat{\mathbf{d}}_{\text{min}}$  be the normalized direction of the minimum mode. Then, in contrast to minimizations, the stepsize that is used to scale  $\nabla E_{\perp}$  is not based on the angle between the complete  $\nabla E$  and  $\nabla E^{\text{P}}$  but only on the angle between  $\nabla E - (\nabla E \cdot \hat{\mathbf{d}}_{\text{min}}) \hat{\mathbf{d}}_{\text{min}}$  and  $\nabla E^{\text{P}} - (\nabla E^{\text{P}} \cdot \hat{\mathbf{d}}_{\text{min}}) \hat{\mathbf{d}}_{\text{min}}$ . These are the components that are responsible for the minimization in directions that are not the minimum mode direction. Otherwise, the gradient feedback is absolutely identical to that described in Sec. II D.

A saddle point can be higher in energy than the configuration at which the optimization is started at. Therefore, in contrast to a minimization, it is not reasonable to discard the history, if the energy increases. As a replacement for this safeguard, we restore to a simple trust radius approach in which any atom must not be moved by more than a predefined trust radius  $r_{\text{trust}}$ . A displacement exceeding this trust radius is simply rescaled. If the curvature is positive and the norm of the gradient is below the convergence criterion, we also rescale displacements that do not come from bond-stretchings. The displacement is rescaled such that the displacement of the atom that moved furthest is finally given by  $r_{\text{trust}}$ . This avoids arbitrarily small steps close to minima.

On very rare occasions, we could observe for some cluster systems that over the course of several iterations, a few atoms sometimes detach from the main cluster. To avoid this problem, we identify the main fragment and move all neighboring fragments towards the nearest atom of the main fragment.

Below, the pseudocode for SQNS is given. It contains 3 parameters explicitly.  $\alpha'_{\text{start}}$  and  $\alpha'_{\text{s,start}}$  are initial step sizes that scale  $\nabla E_{\perp}$  and  $\nabla E_{\text{str}}$ , respectively.  $m'$  is the maximum length of the history list from which the significant subspace is constructed.

The path-length threshold  $r_{\text{thresh}}$  that determines the re-computation frequency of the minimum mode is implicitly contained in line 5. Lines 14 and 21 imply the trust radius  $r_{\text{trust}}$ .

Besides all the parameters that are needed for the minimizer of Sec. II F, line 6 additionally implies the finite difference step size  $h$  that is used to compute the curvature and its gradient.

Line 18 implicitly contains the parameter  $\epsilon$ , which is described in Sec. II B.

1.  $\alpha' \leftarrow \alpha'_{\text{start}}; \alpha'_s \leftarrow \alpha'_{s,\text{start}};$
2.  $l \leftarrow 1;$
3. Initialize  $\mathbf{R}_l$  with coordinates;
4. repeat
5.   if recompute minimum mode then
6.     Use algorithm of Sec. II F and obtain a normalized minimum  $\hat{\mathbf{d}}_{\text{min}}$  of  $c(\mathbf{d})$  at  $\mathbf{R}_l$ , use the previously computed minimum mode as input;
7.   end if
8.   if optimizing biomolecule then
9.     Compute  $\nabla E_{\text{str}}$  for  $\mathbf{R}_l$ , as outlined in Sec. II E;
10.    Adjust  $\alpha'_s$  based on the feedback described in Sec. II E;
11.     $\mathbf{s} \leftarrow \alpha'_s \nabla E_{\text{str}};$
12.     $\mathbf{g}_l \leftarrow \nabla E(\mathbf{R}_l) - \nabla E_{\text{str}};$
13.     $\mathbf{R}_l \leftarrow \mathbf{R}_l - \mathbf{s} + 2(\mathbf{s} \cdot \hat{\mathbf{d}}_{\text{min}}) \hat{\mathbf{d}}_{\text{min}};$
14.    Check for trust radius condition as described in Sec. II G. Rescale, if needed;
15.   else
16.      $\mathbf{g}_l \leftarrow \nabla E(\mathbf{R}_l);$
17.   end if
18.   Based on the  $\{\mathbf{g}_j, \mathbf{R}_j\}_{j \leq l}$  in the history list, compute the preconditioned gradient  $\nabla E^P$  as outlined in Secs. II B—II E;
19.    $\mathbf{R}_{l+1} \leftarrow \mathbf{R}_l - \nabla E^P + 2(\nabla E^P \cdot \hat{\mathbf{d}}_{\text{min}}) \hat{\mathbf{d}}_{\text{min}};$
20.   Check for trust radius condition and for fragmentation as described in Sec. II G. Rescale and fix fragmentation, if needed;
21.   Adjust  $\alpha'$  based on the gradient feedback described in Sec. II G;
22.   if  $l > m'$  then
23.     Remove  $\mathbf{R}_{l-m'}$  and  $\mathbf{g}_{l-m'}$  from the history list;
24.   end if
25.    $l \leftarrow l + 1;$
26. until convergence.

### III. BENCHMARKS AND COMPARISONS

#### A. Minimizers

We compare the performance of the new SQNM method to the FIRE and L-BFGS minimizers. We did not include the CG method in this benchmark, because FIRE has previously been shown to be significantly more efficient than CG.<sup>12</sup> Both FIRE and L-BFGS belong to the best optimizers in their class. With regard to the required number of energy and force evaluation, L-BFGS is one of the best minimizers available for the optimization of atomic systems. With respect to noise tolerance, the same is true for FIRE. Although more efficient than FIRE, L-BFGS tends to fail if there are inconsistent forces and energies due to computational noise.<sup>12</sup> Such inconsistencies are unavoidable in electronic structure calculations like, for example, DFT.

For Si<sub>20</sub> clusters and the alanine dipeptide biomolecule, benchmarks were performed both at DFT and force field level. For L-BFGS, we used the reference implementation of Nocedal<sup>10,11</sup> which is available from his website. We are not aware of any reference implementation of FIRE. However, FIRE is straightforward to implement, and thus, we used our own code. For the benchmarks of the minimizers at DFT level, all codes were coupled to the BigDFT electronic structure code.<sup>29,30</sup> For the benchmarks at force field level, we used the Assisted Model

Building with Energy Refinement (AMBER) force field in the ff99SB variant as implemented in AMBER Tools<sup>31</sup> and the Lenosky Silicon force field.<sup>22,23</sup>

For alanine dipeptide and Si<sub>20</sub>, we generated test sets by running MD simulations at force field level. At force field level, each test set contains 1000 structures that were taken from the MD trajectories. Subsets containing 100 of these force field structures were used as benchmark systems at DFT level. For each method, we tuned the parameters at force field level for a subset of 100 configurations. Identical parameters were used both at force field and DFT level. The Si<sub>20</sub> system was considered to be converged as soon as the norm of the force fell below  $1.0 \times 10^{-4}$  hartree/bohr. Even if far away from a stationary point, relatively small forces can arise in alanine dipeptide. Therefore, a much tighter convergence criterion of  $1.0 \times 10^{-5}$  hartree/bohr had to be chosen for alanine dipeptide.

Table I gives the results of these benchmarks. In addition to the average number of energy and force calls  $\langle n_{\text{ef}} \rangle$ , we also give the average integrated path length of the optimization path  $\langle r \rangle$ .  $\langle r \rangle$  is computed by summing all the distances between structures for which consecutive energy and force evaluations were performed.

There is no guarantee that minimizations that are started at the same configuration will converge to the identical minimum. Therefore, Table I gives averages for both, the subset of runs that all converged to identical minima and averages over all runs, regardless of whether the final minima were identical or not. Identical configurations were identified by using the recently developed s-overlap fingerprints.<sup>32</sup>

In all benchmarks, FIRE is clearly inferior to L-BFGS and SQNM. With respect to the average number of energy and force evaluations, the L-BFGS method is slightly more efficient than the new SQNM minimizer. However,  $\langle r \rangle$  of L-BFGS is 1.6—2.6 times larger than the corresponding values of the SQNM method. On average, this means that L-BFGS displaces the atoms more violently than SQNM. In DFT calculations, the wavefunction of the previous optimization step can be used as input wave function for the current iteration. Roughly speaking, the less the positions of the atoms have changed, the better this input guess usually is. Therefore, less wavefunction optimizations are needed for convergence. To quantify this, the average number of wavefunction optimization iterations  $\langle n_{\text{woi}} \rangle$  needed for a minimization of the potential energy surface is given in Table I. As a consequence of the smaller displacements in the SQNM method, the L-BFGS and the SQNM methods roughly need the same number of wavefunction optimizations for converging to a minimum of the potential energy surface.

The L-BFGS minimizer needed less energy and force evaluations at force field level than at DFT level. We verified that this is not due to the noise at DFT level, but a consequence of the different natures of both potential energy surfaces. The force field potential energy surface is not a noiseless variant of the DFT potential energy surface but a rather inaccurate approximation to it. In particular, this means that the frequencies of the force field are different from the frequencies of the DFT energy surface, and therefore, the same is true for the condition numbers. Therefore, one cannot expect to obtain the same number of energy and force evaluations at force field and DFT level.



TABLE I. Benchmark results for minimizers. DFT test sets contain 100, force field test sets contain 1000 distinct structures. SQNM runs labeled with “(Bio)” indicate the usage of the preconditioner for biomolecules described in Sec. II E.

System	Level of theory	Method	To same minimum					To arbitrary minimum			
			$n_f^a$	$N^b$	$\langle n_{et} \rangle^c$	$\langle r \rangle^d$	$\langle n_{woi} \rangle^e$	$N^b$	$\langle n_{et} \rangle^c$	$\langle r \rangle^d$	$\langle n_{woi} \rangle^e$
Alanine dipeptide	DFT	FIRE	0	93	454	14.01	7602	100	458	14.14	7662
		L-BFGS	2	93	185	23.41	3876	98	188	24.02	3941
		SQNM (Bio)	0	93	198	14.10	3711	100	207	14.29	3858
	Force field	FIRE	0	954	414	12.21	—	1000	418	12.35	—
		L-BFGS	1	954	156	19.69	—	999	158	20.39	—
		SQNM (Bio)	0	954	188	12.38	—	1000	192	12.57	—
Si <sub>20</sub>	DFT	FIRE	0	46	139	18.83	2458	100	143	19.52	2513
		L-BFGS	30	46	73	27.19	1677	70	74	31.26	1714
		SQNM	0	46	83	16.00	1740	100	86	16.50	1784
	Force field	FIRE	0	486	147	13.26	—	1000	163	15.32	—
		L-BFGS	0	486	57	25.49	—	1000	65	30.44	—
		SQNM	0	486	72	10.82	—	1000	81	11.93	—

<sup>a</sup>Number of failed optimizations.<sup>b</sup>Number of runs over which the averages are taken.<sup>c</sup>Average number of energy and force calls (only successful runs).<sup>d</sup>Average integrated path length of the optimization trajectory in units of bohr.<sup>e</sup>Average number of wavefunction optimization iterations.

The L-BFGS minimizer proved to be unreliable at DFT level. For example, 30% of all Si<sub>20</sub> minimizations failed to converge. In contrast to this, all SQNM runs successfully converged to a minimum. The failure to converge in the L-BFGS method is generally due to failures of the line minimizations in the final part of the optimization where a large fraction of the forces consists of noise.

## B. Saddle finding methods

The SQNS method was compared to an improved version of the dimer method<sup>15</sup> as described in Ref. 16 and as implemented in the EON code.<sup>33</sup> In this improved version, the L-BFGS<sup>10,11,16</sup> algorithm is used for the rotations and translation

of the dimer. Furthermore, the rotational force and the dimer energy are evaluated by means of a first order forward finite difference of the gradients.<sup>16,34,35</sup> The same force fields as for the minimization benchmarks were used. For the DFT calculations, SQNS was coupled to the BigDFT code. The EON codes offer an interface to the Vienna *Ab-initio* Simulation Package (VASP),<sup>36–40</sup> which consequently was used.

The same test sets as for the minimizer benchmarks were used. In particular, this means that the starting configurations are not close to a saddle point, and therefore, these test sets are comparatively difficult for saddle finding methods. Again, parameters were only tuned for a subset of 100 configurations at force field level. With exception to the finite difference step size that is needed to calculate the curvature and its gradient, we

TABLE II. Benchmark results for saddle finding methods. DFT test sets contain 100, force field test sets contain 1000 distinct structures. SQNS runs labeled with “(Bio)” indicate the usage of the preconditioner for biomolecules described in Sec. II E.

System	Level of theory	Method	$n_f^a$	To same saddlepoint		To arbitrary saddlepoint	
				$N^b$	$\langle n_{et} \rangle^c$	$N^b$	$\langle n_{et} \rangle^c$
Alanine dipeptide	DFT	SQNS (Bio)	0	—	—	100	510
	Force field	DIMER	0	87	1324	1000	3146
		SQNS (Bio)	0	87	309	1000	415
		SQNS	0	87	632	1000	757
Si <sub>20</sub>	DFT	DIMER	0	8	234	100	444
		SQNS	0	8	140	100	237
	Force field	DIMER	0	20	264	1000	622
		SQNS	0	20	189	1000	368

<sup>a</sup>Number of failed optimizations.<sup>b</sup>Number of runs over which the averages are taken.<sup>c</sup>Average number of energy and force calls (only successful runs).

used the same parameters at force field and DFT level. Because of noise, the finite difference step size must be chosen larger at DFT level. The same force norm convergence criteria as for the minimization benchmarks were used. In all SQNS optimizations, the minimum mode was recalculated at convergence (condition 5 of Sec. II G).

The test results are given in Table II. In contrast to the minimization benchmarks, we do not give averages for the number of wavefunction optimization iterations, because the two saddle finding methods were coupled to two different electronic structure codes. Therefore, the number of wavefunction optimizations is not comparable.

In particular, in case of the Si<sub>20</sub> system, both methods converged only seldom to the same saddle points, and therefore, the statistical significance of the corresponding numbers given in Table II is limited. However, averages over large sets could be made in the case of convergence to an arbitrary saddle point.

In the cases we considered, the dimer method needed between 1.4 and 7.6 times more energy and force evaluations than the new SQNS method. In particular, for alanine dipeptide, the SQNS approach was far superior to the dimer method. Due to its inefficiency, it was impossible to obtain a significant number of saddle points for alanine dipeptide at DFT level when using the dimer method. For this reason, only benchmark results for the SQNS method are given for alanine dipeptide at DFT level.

#### IV. CONCLUSION

Optimizations of atomic structures belong to the most important routine tasks in fields like computational physics, chemistry, or biology. Although the energies and forces given by computationally demanding methods like DFT are physically accurate, they are contaminated by noise. This computational noise comes from underlying integration grids and from self-consistency cycles that are stopped at non-vanishing thresholds. The availability of optimization methods that are not only efficient but also noise-tolerant is therefore of great importance. In this contribution, we have presented a technique to extract significant curvature information from noisy potential energy surfaces. We have used this technique to create a SQNM and a SQNS algorithm. SQNM and SQNS were demonstrated to be superior to existing efficient and well established methods.

Until now, the SQNM and the SQNS optimizers have been used over a period of several months within our group. During this time, they have performed thousands of optimizations without failure at the DFT level. Because of their robustness with respect to computational noise and due to their efficiency, they have replaced the default optimizers that have previously been used in Minima Hopping<sup>41,42</sup> and Minima Hopping Guided Path Search<sup>43</sup> runs.

Implementations of the minimizer and the saddle search method are made available via the BigDFT electronic structure package. The code is distributed under the GNU General Public License and can be downloaded free of charge from the BigDFT website.<sup>44</sup>

#### ACKNOWLEDGMENTS

We thank the Indo-Swiss Research grant for financial support. Computer time was provided by the Swiss National Supercomputing Centre (CSCS) under Project ID s499.

- <sup>1</sup>D. Wales, *Energy Landscapes: Applications to Clusters, Biomolecules and Glasses* (Cambridge University Press, 2003).
- <sup>2</sup>H. Eyring, *J. Chem. Phys.* **3**, 107 (1935).
- <sup>3</sup>P. Pulay, *Chem. Phys. Lett.* **73**, 393 (1980).
- <sup>4</sup>P. Pulay, *J. Comput. Chem.* **3**, 556 (1982).
- <sup>5</sup>M. R. Hestenes and E. Stiefel, *J. Res. Natl. Bur. Stand.* **49**, 409 (1952).
- <sup>6</sup>C. G. Broyden, *J. Appl. Math.* **6**, 76–90 (1970).
- <sup>7</sup>R. Fletcher, *Comput. J.* **13**, 317–322 (1970).
- <sup>8</sup>D. Goldfarb, *Math. Comput.* **24**, 23–26 (1970).
- <sup>9</sup>D. F. Shanno, *Math. Comput.* **24**, 647–656 (1970).
- <sup>10</sup>J. Nocedal, *Math. Comput.* **35**, 773 (1980).
- <sup>11</sup>D. C. Liu and J. Nocedal, *Math. Program.* **45**, 503–528 (1989).
- <sup>12</sup>E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, and P. Gumbsch, *Phys. Rev. Lett.* **97** (2006).
- <sup>13</sup>F. Tassone, F. Mauri, and R. Car, *Phys. Rev. B* **50**, 10561–10573 (1994).
- <sup>14</sup>M. Probert, *J. Comput. Phys.* **191**, 130–146 (2003).
- <sup>15</sup>G. Henkelman and H. Jonsson, *J. Chem. Phys.* **111**, 7010 (1999).
- <sup>16</sup>J. Kästner and P. Sherwood, *J. Chem. Phys.* **128**, 014106 (2008).
- <sup>17</sup>P.-O. Löwdin, *Adv. Phys.* **5**, 1–171 (1956).
- <sup>18</sup>I. Mayer, *Simple Theorems, Proofs, and Derivations in Quantum Chemistry* (Springer, New York, 2003).
- <sup>19</sup>F. Jensen, *Introduction to Computational Chemistry* (John Wiley & Sons, 2007).
- <sup>20</sup>D. Weinstein, *Proc. Natl. Acad. Sci. U. S. A.* **20**, 529 (1934).
- <sup>21</sup>Y. Suzuki and K. Varga, *Stochastic Variational Approach to Quantum-Mechanical Few-Body Problems* (Springer, 1998).
- <sup>22</sup>T. J. Lenosky, B. Sadigh, E. Alonso, V. V. Bulatov, T. D. d. I. Rubia, J. Kim, A. F. Voter, and J. D. Kress, *Modell. Simul. Mater. Sci. Eng.* **8**, 825–841 (2000).
- <sup>23</sup>S. Goedecker, *Comput. Phys. Commun.* **148**, 124–135 (2002).
- <sup>24</sup>S. Goedecker, F. Lancon, and T. Deutsch, *Phys. Rev. B* **64**, 161102 (2001).
- <sup>25</sup>C. J. Cerjan, *J. Chem. Phys.* **75**, 2800 (1981).
- <sup>26</sup>D. J. Wales, *J. Chem. Soc., Faraday Trans.* **89**, 1305 (1993).
- <sup>27</sup>M. Page and J. W. McIver, *J. Chem. Phys.* **88**, 922 (1988).
- <sup>28</sup>G. Mills and K. W. Jacobsen, in *Classical and Quantum Dynamics in Condensed Phase Simulations*, edited by G. C. B. J. Berne and D. F. Coker (World Scientific, 1998), p. 385.
- <sup>29</sup>L. Genovese, A. Neelov, S. Goedecker, T. Deutsch, S. A. Ghasemi, A. Willand, D. Caliste, O. Zilberberg, M. Rayson, A. Bergman, and R. Schneider, *J. Chem. Phys.* **129**, 014109 (2008).
- <sup>30</sup>S. Mohr, L. E. Ratcliff, P. Boulanger, L. Genovese, D. Caliste, T. Deutsch, and S. Goedecker, *J. Chem. Phys.* **140**, 204110 (2014).
- <sup>31</sup>D. Case, V. Babin, J. Berryman, R. Betz, Q. Cai, D. Cerutti, T. Cheatham III, T. Darden, R. Duke, H. Gohlke, A. Goetz, S. Gusarov, N. Homeyer, P. Janowski, J. Kaus, I. Kolossvy, A. Kovalenko, T. Lee, S. LeGrand, T. Luchko, R. Luo, B. Madej, K. Merz, F. Paesani, D. Roe, A. Roitberg, C. Sagui, R. Salomon Ferrer, G. Seabra, C. Simmerling, W. Smith, J. Swails, R. Walker, J. Wang, R. Wolf, X. Wu, and P. Kollman, *AMBER 14* (University of California, San Francisco, 2014).
- <sup>32</sup>A. Sadeghi, S. A. Ghasemi, B. Schaefer, S. Mohr, M. A. Lill, and S. Goedecker, *J. Chem. Phys.* **139**, 184118 (2013).
- <sup>33</sup>S. T. Chill, M. Welborn, R. Terrell, L. Zhang, J.-C. Berthet, A. Pedersen, H. Jansson, and G. Henkelman, *Modell. Simul. Mater. Sci. Eng.* **22**, 055002 (2014).
- <sup>34</sup>A. Heyden, A. T. Bell, and F. J. Keil, *J. Chem. Phys.* **123**, 224101 (2005).
- <sup>35</sup>R. A. Olsen, G. J. Kroes, G. Henkelman, A. Arnaldsson, and H. Jonsson, *J. Chem. Phys.* **121**, 9776 (2004).
- <sup>36</sup>G. Kresse and J. Hafner, *Phys. Rev. B* **47**, 558–561 (1993).
- <sup>37</sup>G. Kresse and J. Hafner, *Phys. Rev. B* **49**, 14251–14269 (1994).
- <sup>38</sup>G. Kresse and J. Furthmüller, *Comput. Mater. Sci.* **6**, 15–50 (1996).
- <sup>39</sup>G. Kresse, *Phys. Rev. B* **54**, 11169–11186 (1996).
- <sup>40</sup>G. Kresse, *Phys. Rev. B* **59**, 1758–1775 (1999).
- <sup>41</sup>S. Goedecker, *J. Chem. Phys.* **120**, 9911 (2004).
- <sup>42</sup>S. Goedecker, W. Hellmann, and T. Lenosky, *Phys. Rev. Lett.* **95**, 055501 (2005).
- <sup>43</sup>B. Schaefer, S. Mohr, M. Amsler, and S. Goedecker, *J. Chem. Phys.* **140**, 214102 (2014).
- <sup>44</sup>See <http://bigdft.org> for source codes for SQNM and SQNS.