Solution of boundary-element problems using the fast-inertial-relaxation-engine method

Yunong Zhou,[1] Michael Moseler,[2] and Martin H. Müser[1]

[1] *Department of Materials Science and Engineering, Universität des Saarlandes, 66123 Saarbrücken, Germany*

[2] *Faculty of Physics, University of Freiburg, Hermann-Herder-Strasse 3, 79104, Freiburg, Germany*

(Dated: Received: date / Accepted: date)

The fast inertial-relaxation engine (FIRE) has proven to efficiently find local minima of potential energies or related penalty functions although its implementation requires only few, additional lines of code in a molecular-dynamics or steepest-descent program. So far, FIRE has been predominantly applied to particle-based or low-dimensional problems. In this work, we demonstrate that it can also benefit the solution of boundary-value problems. Towards this end, we study the mechanical contact between an elastic body and a rigid indenter of varying complexity by augmenting Green's function molecular dynamics (GFMD) with FIRE. We find a rather remarkable speed-up, which can be further enhanced when choosing the masses associated with the eigenmodes of the free elastic solid appropriately. For the investigated adhesive and randomly rough indenter with typical system size, 100 mass-weighted FIRE-GFMD iterations suffice to relax the excess energy to 10^{-3} of its original value. The standard GFMD method needs 25 times more iterations. For the investigated problems, FIRE-GFMD even appears to slightly outperform conjugate-gradient based optimization.

I. INTRODUCTION

Finding minima of some penalty function in a highly-dimensional variable space is a standard computational problem. For example, identifying mechanically stable structures requires minima of potential-energy surfaces to be located. Various traditional minimization methods, such as steepest descent, take steps in the variable or configuration space, which are antiparallel to the gradient direction, i.e., parallel to the force[1–4]. Such algorithms can lead to undesired zig-zag motion, in which case the minimum is only reached slowly. The conjugate-gradient method (CGM) avoids the zig-zag motion and converges quickly to a minimum if the penalty function is quadratic. However, when the penalty function is far from being quadratic, CGM-based minimization may have to be restarted many times before satisfying convergence is reached. Moroever, it requires the use of libraries or writing relatively elaborate code.

The fast internal-relaxation engine (FIRE) is a minimization method, which can suppress undesired zig-zag motion[5]. It only necessitates a few lines of code to be added to any existing molecular-dynamics (MD) or steepest-descent program. The implementation and basic ideas of FIRE work as follows: Inertia are assigned to the variables leading to an implicit averageing of the gradient direction over past iterations and turning a steepest-descent program into a MD code. At the same time, the instantaneous velocity is slightly biased toward the steepest-descent direction. Moreover, the time step size can be increased with each iteration, since true dynamics do not matter. Once the vector product of velocity and forces (negative gradients) is negative, all velocities are set to zero, the time step is set back to a small value, and the procedure is restarted with the original, small time step.

So far, FIRE has been applied to traditional particle-based problems. The underlying idea should, however, be more generally applicable, in particular to the solution of partial-differential equations (PDEs). The reason is that the solution of PDEs can be mapped onto an MD problem after discretization of the variable space. Assume $f(\mathbf{r})$ is the field to be found. Then any function that is a local minimum if and only if $f(\mathbf{r})$ satisfies the PDE plays the role of energy, while $f(\mathbf{r})$ at each discretization point corresponds to a (generalized) coordinate. In a similar way, FIRE is applicable to classical boundary-value problems (BVP).

As one such example, we consider the mechanical contact between an adhesive indenter and a flat, linearly elastic solid, whose static, elastic energy is fully defined by the shape of the surface and the elastic tensor. Green's function molecular dynamics (GFMD) is a technique that allows such BVPs to be addressed within the framework of MD[6–8]. The central idea is to use the Fourier transforms of the surface modes as coordinates, which are propagated according to Newton's equations of motion. The latter are usually augmented with a damping term in order to pull energy out of the system until it has settled in an energy minimum. In this study, we replace the damping term in GFMD with a FIRE-based algorithm and investigate how this modification effects the rate of convergence.

In this work, we also explore further optimizations, which we expect to benefit not only the solution of BVPs ressembling the studied contact problem but also the solution of PDEs. In many systems – most notably particle-based systems and second-order differential equations – long-wavelength modes relax more slowly than short-wavelength modes. Assigning wavelength-dependent inertia can then help to match frequencies, in particular when the problem is dominated by the diagonal in a Fourier representation.

In the given context, the restoring force on a free-surface mode is essentially proportional to $-q\tilde{u}(q)$ at large q, where q is the wave vector and $\tilde{u}(q)$ the Fourier transform of the displacement field to be determined. Thus, assigning inertia or masses that increase linearly with q at large q, must be expected to reduce characteristic long-wavelength time scales relative to those of short wavelengths.

In addition to the methods proposed here, we also explore how quickly a conjugate-gradient (CG) based minimization finds the proper solution. Toward this end, CG – as implemented by Bugnicourt, Sainsot, *et al.*[9] – was added as an option to our code. The CG method by Bugnicourt and coworkers had not only outrun regular GFMD in the contact-mechanics challenge[10]. In our understanding, the CG implementation of that group had led to the overall most quickly convergent solution, although other CG-based contact-mechanics methods[11–14] may well be on par. The contact-mechanics challenge was a publicly announced large-scale contact problem for three-dimensional solids having the added complexity of short-range adhesion and hard-wall repulsion. More than one dozen groups participated in the exercise using a similarly large number of solution strategies.

In the remaining part of this paper, the problems are defined in Sect. II, while the numerical methods are described in Sect. III. Numerical results are presented in Sect. IV and conclusions are drawn in the final Sect. V.

II. MODEL AND PROBLEM DEFINITION

A. Treatment of elasticity

The unit for pressure is defined by the contact modulus $E^* = E/(1-\nu^2)$, where E is the Young's modulus. The Poisson number is taken to be $\nu = 1/4$. This choice is a matter of convenience. It allows an isotropic elastic body to be represented as a simple bead-spring system, since the use of pair potential restricts the elastic tensor to satisfy the Cauchy relation $C_{1122} = C_{1212}$, or, in Voigt notation to $C_{12} = C_{44}$ in the case of isotropic solids.

The height h of the deformable body is set to half its width L. We see this as the smallest natural width at which the solid can be rightfully claimed to be almost infinite in the direction normal to the interface. This is because (static) stress undulations decay roughly as $\exp(-qz)$ into the substrate. ~~Therefore, the damping~~ This factor turns out to be $e^{-\pi} \approx 0.04$ when evaluated on the surface opposite to the side touching the indenter for the system's smallest, non-zero wavevector, $q = 2\pi/L$. A constant normal pressure is applied to that opposite side squeezing the elastic body against the rigid indenter, which is fixed in space.

In order to be in a position to compare the efficiency of our solution to a full, all-atom simulation, we restrict our attention primarily to (1+1)-dimensional solids, which means that the rigid indenters correspond to cylinders, whose symmetry axes are oriented parallel to the z-axis. This means that all our energies are, in fact, line energy densities.

Two approaches are pursued to compute the elastic energy of the system. In one approach, the complete (cross-section of the) solid is discretized into a square lattice. Nearest neighbors and next-nearest neighbors interact with springs of "stiffness" $k_1 = 0.75\ E^*$ and $k_2 = 0.375\ E^*$, respectively. (True spring stiffnesses have to multiplied with the length of the cylinder in z-direction.) These values do not change with the discretization of our effectively two-dimensional elastic body. The equilibrium length of the two springs are set to the equilibrium nearest and next-nearest neighbor distance, r_1 and r_2, respectively. The spring energy (line density) can thus be written as

$$V_{\mathrm{el}} = \frac{1}{2}\sum_{i,j>i} k_{ij}\left\{r_{ij} - r_{ij}^{\mathrm{eq}}\right\}^2. \tag{1}$$

In the second approach, we exploit that all information on the (static) elastic energy is encoded in the displacement field $\mathbf{u}(x)$ of the bottom layer (indenter comes from below), as it is frequently done for the solution of BVPs. The displacements may occur in both directions that are orthogonal to the cylinder's axis, i.e., parallel to x and y. The elastic energy is best computed in "reciprocal space", i.e., in terms of the Fourier transforms of the displacements, for which the following convention is used

$$\tilde{u}_\alpha(q) = \frac{1}{N}_x \sum_{n=1}^{N_x} u_{n\alpha}\exp\{iqx\} \tag{2}$$

$$u_{n\alpha}(x) = \sum_q \tilde{u}_\alpha(q)\exp\{-iqx\}. \tag{3}$$

The wave numbers satisfy $-\pi N_x/L \le q < \pi N_x/L$, where qL/π takes all integer numbers from $-N_x/2$ to $N_x/2-1$. In fact, not only the energy is evaluated in reciprocal space when using GFMD, but the entire dynamics are conducted in the Fourier representation. Moreover, Greek indices numerate Cartesian coordinates, $\alpha = 1$ corresponding to the x-coordinate and $\alpha = 2$ to y, while the Latin letter n enumerates grid points.

With these definitions, the elastic energy (line density) of the deformed solid (finite height h, constant pressure acting on the surface opposite to the indenter) can be expressed as

$$V_{\mathrm{ela}} = \frac{L}{4}\sum_q \sum_{\alpha\beta} qM_{\alpha\beta}(q)\tilde{u}^*_\alpha(q)\tilde{u}_\beta(q), \tag{4}$$

where the matrix coefficients $M_{\alpha\beta}$ contain all needed information on the elastic coupling between different

modes[15]. They read

$$M_{11}(qh) = (1-r)\frac{\cosh(qh)\sinh(qh) - rqh}{\|f(qh)\|}C_{1111}$$
$$M_{12}(qh) = \frac{1-r}{1+r}\frac{(1-r)\sinh^2(qh) - 2(rqh)^2}{\|f(qh)\|}C_{1111}$$
$$M_{22}(qh) = (1-r)\frac{\cosh(qh)\sinh(qh) + rqh}{\|f(qh)\|}C_{1111},$$

where

$$r = \frac{1-s}{1+s}$$
$$s = \frac{C_{1212}}{C_{1111}}$$

and

$$\|f(qh)\| = \cosh^2(qh) - (rqh)^2 - 1.$$

B. Geometries and interaction

As a first example, we consider a generic, single-asperity contact, which is sketched in Fig. 1. The elastic layer defined in Sect. II A is in contact with a parabolic rigid indenter. The profile of the indenter is given by

$$h(x) = -x^2/2R_c, \tag{5}$$

where R_c is the radius of curvature. The elastic solid is placed such that all y_n are set to zero as long as no external forces are acting on the body. The lateral equilibrium positions are set to $x_n^{eq} = nL/N_x$. The normal distance g_n of a grid point n from the indenter g_n then reads

$$g_n = u_{n,y} - h_s(x_n), \tag{6}$$

where $x_n = x_n^{eq} + u_{n,x}$ is the lateral position of the discretization point n.

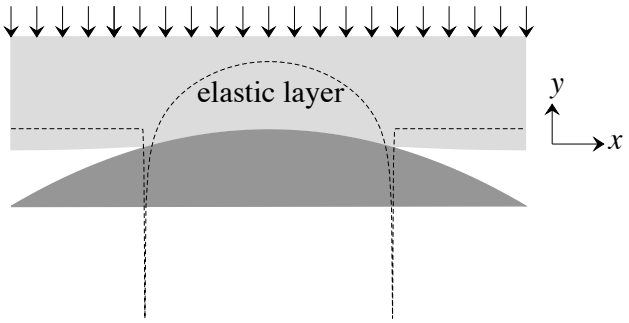

FIG. 1: A weakly adhesive, elastic solid of finite thickness compressed by a parabolic indenter. The dotted line shows the associated stress profile.

The interaction between the indenter and the elastic body consists of a short-range adhesion and an even shorter-ranged repulsion. Specifically, we define the interaction energy (line density) as follows:

$$V_{int} = \frac{L}{N_x}\sum_n \gamma_1 \exp(-2g_n/\rho) - \gamma_2 \exp(-g_n/\rho) \tag{7}$$

where γ_i has the unit energy per surface area and ρ of length. We chose them to be $\rho \approx 2.56 \times 10^{-4}\ R_c$, $\gamma_1 \approx 2.10 \times 10^3\ E^* R_c$, and $\gamma_2 \approx 2.05\ E^* R_c$. With these choices, the equilibrium gap between the two solids would be $\rho_{eq} \approx 1.95 \times 10^{-3}\ R_c$ and a surface energy of $\gamma_{eq} = 5.0 \times 10^{-4}\ E^* R_c$ would we gained at a gap of ρ_{eq}. This leads to a Tabor coefficient of $\mu_T \gtrsim 3$, i.e., the adhesion is (mildly) short ranged, as evidenced by a small adhesive neck, which is visible at the contact edge shown in Fig. 1.

Assuming that ρ_{eq} is of atomic dimension, $\rho_{eq} \approx 3$ Å and the interfactial interactions are dispersive, i.e., $\gamma_{eq} \approx 50$ mJ/m^2, the studied system would correspond to $R_c \approx 150$ nm and $E^* \approx 650$ MPa. This value is representative of a thermoplastic polymer, such as polypropylene. By choosing alternative numbers for ρ_{eq} and γ_{eq}, additional examples of potential practical interest can be related to the dimensionless numbers used in the presented case study.

In the second example, the properties of the elastic body remain unchanged but an indenter with a randomly rough, self-affine shape is considered, see Fig. 2. Its power spectrum $C(q)$ for a $D = 1+1$ dimensional solid is defined as follows[16]

$$C(q) \propto q^{-2H-1}\Theta\left(q_{max} - q\right),$$

where $H = 0.8$ is called the Hurst exponent, $\Theta(\bullet)$ is the Heavyside step function, and $q_{max} = 1024\, q_0$ with $q_0 = 2\pi/L$.

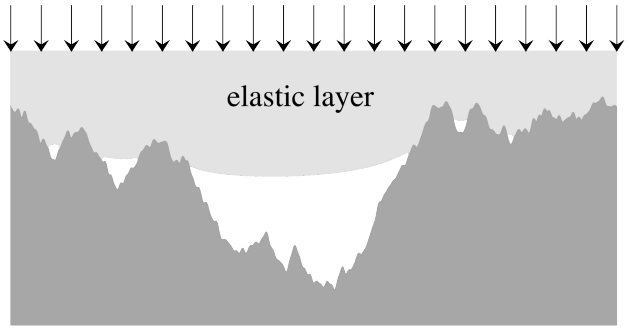

FIG. 2: An elastic layer of thickness h in contact with a random roughness substrate. The figure is not to scale, i.e., the resolution in y direction is enhanced. The height of the elastic body, which is not fully shown, is set to half its width.

The third and final example, which is only touched upon peripherally, is the problem defined in the contact-mechanics challenge[10]. It is similar to the second example, however, the interface is two-dimensional

and the short- but finite-range repulsion consists of a non-overlap constraint. For more details, the reader is referred to the original literature[10]. Configurations and exact problem definitions can currently be downloaded at https://www.lmp.uni-saarland.de/index.php/research-topics/contact-mechanics-challenge-announcement/. The surface profiles are also stored at http://musam.imtlucca.it/wikisurf.html. So far, we have not yet succeeded in implementing a functioning mass-weighted GFMD method with a non-overlap constraint so that we only test the performance of FIRE-GFMD for this final example. Mass-weigted and FIRE-GFMD are introduced next in Sect. III.

III. NUMERICAL METHODS

A. Mass-weighted GFMD

As mentioned in the introduction, GFMD solves Newton's equations of motion in Fourier space rather than in real space. The equation of motion for each surface mode thus reads

$$m(q)\,\ddot{\tilde{\mathbf{u}}}(q) = \tilde{\mathbf{f}}(q), \tag{8}$$

where the Fourier transform of the generalized total force, $\tilde{\mathbf{f}}(q)$, can be divided into an elastic, an interfacial, and an external contribution. m is an inertia or mass, which in our case, would be of dimension kg/m. ~~The force coefficients~~ The coefficients of the force vector are computed, respectively, as

$$\tilde{f}_{\mathrm{ela},\alpha}(q) = -\frac{qE^*}{2}\sum_\beta M_{\alpha\beta}(q)\tilde{u}_\beta(q) \tag{9}$$

$$\tilde{f}_{\mathrm{int},\alpha}(q) = \frac{1}{N_x}\sum_n \frac{\partial V_{\mathrm{int}}}{\partial r_\alpha}\exp\left(iqx_n^{\mathrm{eq}}\right) \tag{10}$$

$$\tilde{f}_{\mathrm{ext},\alpha}(q) = p_0\,\delta_{\alpha 2}\,\delta_{q0}, \tag{11}$$

where p_0 is the external force divided by the linear length of the system in x-direction. Imposing non-isotropic normal or shear stresses on the top surface of the elastic body rather than a constant normal pressure would necessitate additional force terms, which can be deduced from the literature[15]. The such computed force (per cylinder length) is the negative gradient of the total potential-energy line density, V_{tot},

$$V_{\mathrm{tot}} = V_{\mathrm{ela}} + V_{\mathrm{int}} - p_0\,N_x\tilde{u}_y(0), \tag{12}$$

When solving static problems, V_{tot} needs to be minimized.

Standard GFMD assumes $m(q)$ to be constant, in which case identical dynamics are obtained as if the equations of motion were integrated in real space. More natural dynamics would be produced – at least for semi-infinte solids – if $m(q)$ were chosen proportional to $1/q$, because the elastic deformation of an undulation with wavevector λ penetrates $O(\lambda)$ deep into the elastic body. Efficient dynamics, which are useful to quickly identify a minimum of the total energy, are obtained if the effective masses are chosen proportional to the stiffness at wavevector q. For a free surface, this would be $m(q) \propto qE^*$, in the limit of large thicknesses h. In the presence of an external force, an additional contribution arises due to the contact stiffness k_{cont}, which couples, in particular to the center-of-mass or $q = 0$ mode. In this case, a sensible choice for a quick convergence to the minimum is

$$m(q) \propto \sqrt{(qE^*)^2 + \theta\,(k_{\mathrm{cont}}/A)^2}, \tag{13}$$

where A is the apparent contact area and θ a number of order unity. If k_{cont} is known reasonably well, the natural time scales of small variations around a minimum are of similar order of magnitude, which implies that long-wavelength modes relax to their minima with a similar characteristic time scale as short-wavelength modes. Sometimes, k_{cont} may not be known to high precision. However, a systematic slowing down with system size – or with increased small-scale resolution – is prevented from happening even if the estimate of the optimim choice for k_{cont} is off by a factor of 10 or 100.

In the case of randomly rough surfaces, k_{cont} can often be roughly estimated to be a small but finite fraction of the external pressure divided by the root-mean-square height $\bar{h}$, say $k_{\mathrm{cont}} \approx p_0/(10\bar{h})$.

B. Mass-weighted FIRE algorithm

The basic idea of the FIRE algorithm is sketched in the introduction. Technical details pertaining to our study are reported here below. Our system is propagated without damping, as long as the power

$$P = \mathbf{F}\cdot\mathbf{v} \tag{14}$$

is positive, where $\mathbf{F}$ and $\mathbf{v}$ are vectors containing the (generalized) forces and velocities of the considered degrees of freedom. The time step was increased in each iteration by 2%. Moreover, we redirected the instantaneous velocity by changing amounts towards the current direction of steepest descent while keeping the magnitude of the velocity constant. This is done such that $\mathbf{v} \to (1-\xi)\mathbf{v} + \xi\mathbf{f}v/f$, where $\xi = 0.1$ initially and after each FIRE restart. Otherwise, $\xi(t+1) = 0.99\xi(t)$, where t is the time step.

We call our method "mass-weighted", because the dynamics are formulated in Fourier space and the inertia chosen proportional to the expected curvature of a given Fourier mode. We also attempted to do conventional dynamics using q-independent masses and evaluate an effective power, in which increased weights ~~w_q~~ were assigned to the slow modes, ~~i.e.,~~ e.g., by replacing the expression in Eq. (14) with $\sum_q m(q)\tilde{\mathbf{F}}^*(q)\cdot\mathbf{v}(q)$. The effect of this and related modifications to FIRE was meant to make the slow modes relax for as long as possible before restarting

the engine. However, we did not find a parameterization that would lead to an improved convergence, which is why we abonded the idea of a mass-weighting on FIRE without doing the same mass-weighting in the dynamics.

IV. NUMERICAL RESULTS

A. Hertzian indenter

In this section, the efficiency of various minimization techniques is evaluated on a contact-mechanics BVP with one simple and one complex geometry. We start with the simple Hertzian contact geometry, which has become a benchmark for numerical solution techniques in contact mechanics[17]. Due to the short-range adhesion that was added to the regular Hertz problem, the surface topography has features at small scales in addition to the long-range elastic deformation so that different parts of the problem relate to rather distinct scales.

Fig. 3 compares how quickly various solution strategies minimize the energy at a fixed system size. Towards this end, we first compute the excess energy, which is defined as the total potential energy minus the total potential of the fully relaxed structure. The excess energy is then divided by the value obtained for the initial structure, which is set up such that the elastic manifold is flat and located ρ_{eq} above the highest indenter coordinate.

The all-atom simulation turns out to be the most inefficient of all investigated methods, even when measured in (global) time steps, which does not account for the many layers that need to be simulated. The reason is the large disparity of frequencies in the system. The fastest mode, which limits the time step, has an intrinsic frequency ω_{max} which is $O(N_x)$ higher than that of the slowest mode, ω_{min}, for which the damping is chosen to be roughly critical. Adding FIRE to the all-atom simulations leads to an increase of the convergence rate of six for the investigated system size. While this is an improvement, another factor of 2.5 is gained when using conventional GFMD. The ratio $\omega_{\mathrm{max}}/\omega_{\mathrm{min}}$ reduces from $O(N_x)$ to $O(\sqrt{N_x})$, which is at the root of the speed-up compared to natural dynamics.

When adding FIRE to regular GFMD, the convergence increases yet again by almost a factor of three, which is also slightly faster than our CG-based method, which only leads to a speed-up $\gtrsim 2$ at this system size compared to regular GFMD. In contrast, mass-weighted GFMD leads to a speed-up of a factor of ten compared to conventional GFMD. Adding FIRE to the mass-weighted GFMD increases the convergence rate by another 20%. The overall speed-up is remarkable, ~~even if the fastest method (mass-weighted GFMD) remains a linearly convergent method, i.e., the excess energy is reduced by a decade each 15 iterations.~~ MW-FIRE-GFMD needs 75 iterations to reduce the excess energy to 10^{-5} of its original value compared to $\gtrsim$ 15,000 iterations for natural, that is, all-atom dynamics, while

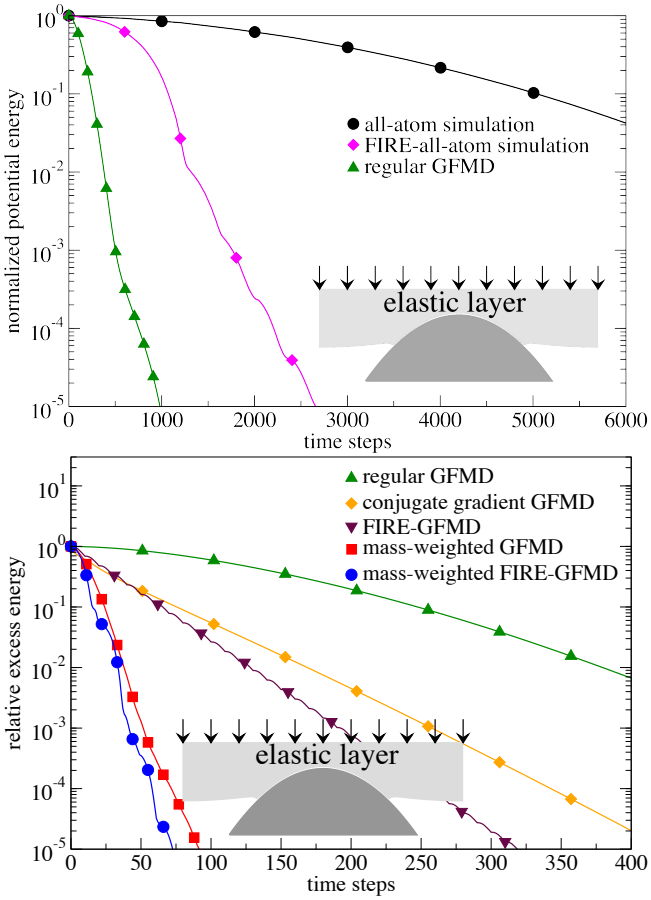

FIG. 3: Relaxation of the relative excess energy $\Delta V(t)/\Delta V(0)$ as a function of time steps for a system with a linear discretization of $N_{\mathrm{x}} = 512$, $\Delta V(t)$ being the potential energy at time t minus that of the fully relaxed structure. Each degree of freedom (DOF) is stepped forward in time once per time step. The GFMD-based methods have N_{x} DOFs, while the all-atom simulation is based on $N_{\mathrm{x}} \times N_{\mathrm{y}}$ DOFs with $N_y = N_x/2$. The top and the bottom graph focus on different methods, regular GFMD being the only one reported in both. Note that the scales on the x-axis differ between the two graphs.

regular GFMD needs 1,000 iterations. The value of 10^{-5} is chosen somewhat arbitrarily. However, the spatially resolved stress profiles start to look very similar to the naked eye as those of the fully relaxed structures when $\Delta V(t)/V(0) < 10^{-5}$. Specifically, errors in the stress are clearly less than 0.5% of the maximum (compressive) contact stress for the investigated and related examples. This number drops to roughly 10^{-3}% when $\Delta V(t)/V(0) < 10^{-8}$.

While the choice of system size – or number of discretization points N_x – is quite common for contact-mechanics problem (give or take a factor of one half ~~or~~ to eight), it is often useful to know how the convergence rate scales with system size. This question is

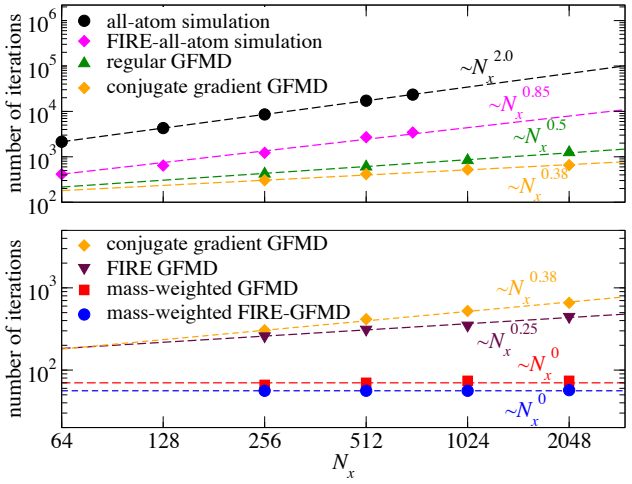

FIG. 4: Number of iterations needed to relax the relative excess energy $\Delta V(t)/\Delta V(0)$ to 10^{-5} for the Hertzian case as a function of the number of grid points.

addressed in Fig. 4. It reveals that the number of iterations needed to decrease the energy to 10^{-5} of its past value scales linearly with N_x (or number of discretization points in the top-most layer) in an all-atom simulation. To achieve this scaling, the damping has to be reduced with increasing N_x. If the damping were kept constant, the long-range modes would be automatically overdamped at large N_x and the scaling would go as N_x^2. Adding FIRE to an all-atom treatment somewhat alleviates the situation, however, the exponent is only reduced to 0.85. Regular GFMD improves that to a square-root dependence. CG-enhanced and FIRE-enhanced improves the scaling to approximately $N^{0.38}$ and $N^{0.25}$, respectively. It thus appears that FIRE slightly outperforms CG-based optimization for this class of problems when systems are very large. For the considered problem, mass-weighting appears to eliminate the size dependence altogether. ~~Similar scaling with system size is found for 2+1 dimensional systems, which we tested explicitly for the various GFMD methods, however, not in the all-atom cases.~~

The scaling shown in Fig. 4 was also found to hold in 2+1-dimensional contact problems, whenever tested, e.g., GFMD, FIRE-GFMD, and MW-GFMD. It also persisted when short-range repulsion was replaced with a hard-wall constraints. However, in this latter case, we did not succeed in getting mass-weighting to work so that regular FIRE-GFMD is the most efficient of the tested methods for hard-wall repulsion.

So far, we have only evaluated the performance in terms of the number of iterations. From a practical point of view, it certainly also matters how long each iteration takes for a given system size, which is analyzed in Fig. 5. In all-atom simulations, the CPU time per iteration grows with the square of the linear system size (with the cube for three-dimensional systems), while it only grows linearly (with the square for two-dimensional systems) plus logarithmic corrections in the GFMD-based approaches.

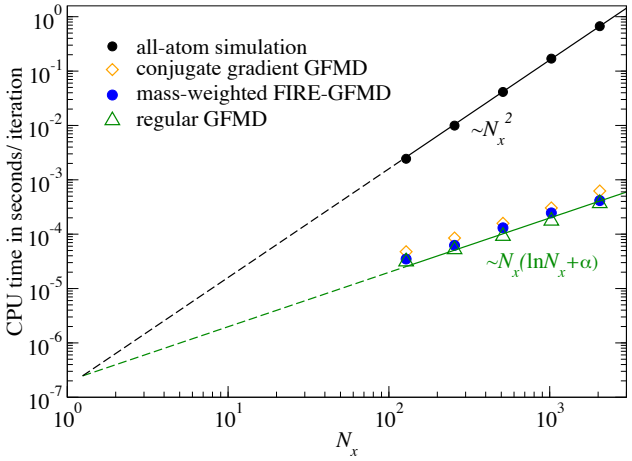

FIG. 5: CPU time in seconds per iteration as a function of the linear system size N_x. The solid lines reflect fits, while the dashed lines are reverse extension of the fits. Adding mass-weighting or FIRE does not significantly affect the time per iteration for typically used numbers. All computations were performed on a personal computer with a 1.6 GHz Intel Core i5 central processor unit (CPU). The FFTW version 3.3.5 is used in our code.

As a final comment in this section, we note that each method can be combined with the regular multi-scale tricks of the trade: A crude result can be obtained at a small discretization and adjustable parameters be gauged (e.g., those needed for FIRE, damping, or mass-weighting). The resolution of the simulation can then be successively increased. This way, the number of iterations is much reduced, while allowing one at the same time to make a Richardson extrapolation of central observables to the continuum limit.

B. Randomly rough indenter

Replacing the simple Hertzian geometry with a more complex self-affine, rigid substrate changes prefactors, but trends remain similar, as can be seen in Fig. 6. Both mass-weighting and FIRE are substantially faster than regular GFMD. Adding FIRE to mass-weighting brings about a speed up by almost a factor of two versus a 25% gain in speed in case of the simpler Hertz geometry.

The exponents of the tested methods remain unchanged within tight margins of ± 0.01 compared to those reported in Fig. 4 for the simple Hertz case. We therefore speculate that they are almost constant for the considered class of contact problems.

C. Application to the contact-mechanics challenge

Since this study is the first time that any of its authors

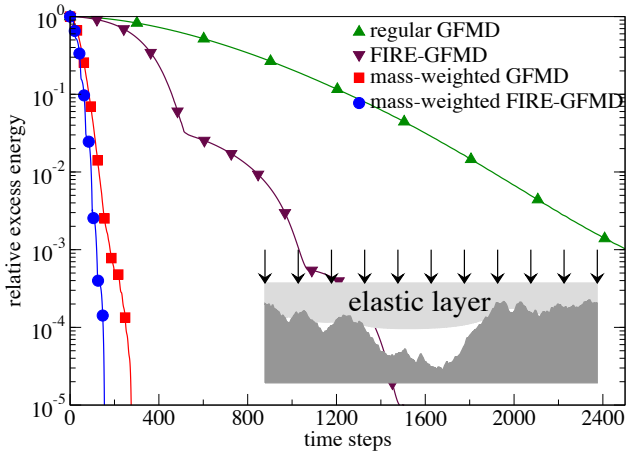

FIG. 6: As in Fig. 3, however, for a randomly rough indenter. The default substrate roughness is set up as follows: We use a Hurst exponent of $H = 0.8$. There is no roll-off so that the system size corresponds to the long wavelength cutoff λ_l. The short wavelength cutoff is chosen to be $\lambda_s = 0.01 L_x$. By default, the system is discretized into $N_x = 1024$ grid points.

actively used a CG method, it was felt that concluding to have outperformed CG bares a non-negligible risk to be erroneous. CG requires some fine tuning of parameters and despite much effort, a finite probability remains of a suboptimal parameter selection. We therefore applied FIRE-GFMD to the problem defined in the contact-mechanics challenge. The team by Bugnicourt reported convergence with their CG method after 3,000 iterations for a discretization of the surface into 32,768 × 32,768 elements, which clearly outrun the regular GFMD-based reference solution in terms of efficiency, for which 30,000 iterations were needed to achieve a similar accuracy at that size. The data submitted by Bugnicourt to the contact-mechanics challenge revealed convergence to a few 10−9 times the maximum compressive stresses. An even greater accuracy would certainly require higher data precision than those obtained when requesting "double" in C++ or "double precision" in Fortran.

The convergence of FIRE-GFMD for the contact-mechanics problem is similar to that identified in this study for related contact problems. This times, a little more than 500 iterations are needed to reduce the excess energy to 10^{-3} of its original value. This value is slightly increased compared to the 300 iterations reported in the previous section. This is due to the contact-mechanics challenge surface being larger and the adhesion being more short ranged than in the problem investigated in the last section. In addition, FIRE-GFMD needs 2,000 iterations to converge the stress to the same accuracy as Bugnicourt. This means that FIRE-GFMD appears to slightly outpeforms the CG-based boundary method by that group. This gain may not be large enough to motivate the replacement of a working CG-based minimization in a code with a new FIRE method. However, when designing new code, the truly easy-to-implement FIRE method appears to be the better choice for this class of problems.

V. CONCLUSIONS

The fast inertial-relaxation engine (FIRE) is based on modified Newtonian dynamics of particles[5]. It was developed as an easy-to-code but highly efficient method for structural optimization of molecular-scale models. As such it has proven useful in numerous atomistic studies. A small, selected list of applications include density-functional-theory based simulations of chemically complex systems[18–21], classical interatomic potential based molecular-statics calculations[22–24], numerical localization of transition pathways[25], and even the simulation of hard-sphere based models of condensed-matter[26,27] systems.

FIRE is implemented in many software suites for molecular statics and dynamics calculations, for instance, in LAMMPS[28], ASE[29], and ATOMISTICA[30,31], which has certainly helped to foster its popularity. However, despite its rootage in inertial dynamics, it is not limited to molecular statics. In principle, any numerical problem that can be formulated as a minimization of a penalty function can benefit from FIRE. One such example could be the optimization of an electrostatic charge distribution subjected to a given set of constraints and boundary conditions.

In this article, two contact-mechanics problems were considered to illustrate the extension of FIRE from the atomistic realm into the continuum-mechanics world. Towards this end, we have employed the Green's function molecular dynamics (GFMD) method, which we see as a prototypical example for how a classical boundary-value problem can be tackled using the molecular-dynamics (MD) toolbox. In its original formulation, GFMD sets up a (critically damped) Newtonian dynamics of surface Fourier modes to minimize the overall energy. It allows for the possibility of adhesion, short-range repulsion, and non-overlap constraints. Since GFMD is an MD method, introducing FIRE into GFMD requires the addition of only a few lines to the GFMD code.

It turns out that FIRE can successfully accelerate a regular GFMD calculation resulting in a remarkable speed up of one order of magnitude for typical system sizes and even larger speed ups for larger systems. FIRE can also be combined in a straightforward fashion with other accellerators of the GFMD method, such as an effective choice for the inertia of the modes. The latter induces a narrow distribution of intrinsic frequencies whereby the number of required sweeps to relax the system no longer increases substantially with system size. Even if the relative speed-up due to FIRE in such a mass-weighted GFMD approach is not overwhelming, a factor of two in efficiency can still be useful for pushing the boundaries of large-scale problems on massive parallel supercomputers.

These first promising results suggest that FIRE could

also be a versatile addition to finite-element, solid-mechanics codes especially for engineering problems with many contacts that usually result in highly non-linear equations. Experience from atomic-scale applications indicates that FIRE is always competitive with much more complex mathematical optimization algorithms[5,32,33] (such as quasi-Newton methods) and sometimes FIRE can even be superior[34]. The contact-mechanics problems considered in this study adds another example, where FIRE appears to slighly outperform conjugate-gradient minimization. Therefore, it might be interesting future research to explore the specific conditions and the underlying reasons for a superior performance by FIRE. Its speed depends on a small set of internal adjustable parameters. The original version of FIRE comes along with a robust choice of these parameters that requires no alteration for molecular-scale calculations. For continuum applications, the previous choices might not be the optimum and systematic studies could provide an additional boost of FIRE.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the German research foundation (DFG) for financial support through grant Mu 1694/5-1.

[1] J. Nocedal, S. Wright, *Numerical Optimization* (Springer, New York, 1996)

[2] W. Press, B. Flannery, S. Teukolsky, W. Vetterling, *Numerical Recipes in C* (Cambridge University, UK, 1997)

[3] A. Leach, *Molecular Modelling: Principles and Applications* (Prentice Hall, UK, 2001)

[4] T. Schlick, *Molecular Modelling and simulation: An Interdisciplinary Guide, Interdisciplinary Applied Mathematics*, vol. 21 (Springer, New York, 2002)

[5] E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, P. Gumbsch, Physical Review Letters **97**(17), 170201 (2006)

[6] E. Karpov, G. Wagner, W.K. Liu, International Journal for Numerical Methods in Engineering **62**(9), 1250 (2005)

[7] C. Campaná, M.H. Müser, Physical Review B **74**(7), 075420 (2006)

[8] L.T. Kong, G. Bartels, C. Campañá, C. Denniston, M.H. Müser, Computer Physics Communications **180**(6), 1004 (2009)

[9] R. Bugnicourt, P. Sainsot, D. Dureisseix, C. Gauthier, A.A. Lubrecht, Tribology Letters **66**(1), 29 (2018). doi:10.1007/s11249-017-0980-z. URL `https://doi.org/10.1007/s11249-017-0980-z`

[10] M.H. Müser, W.B. Dapp, R. Bugnicourt, P. Sainsot, N. Lesaffre, T.A. Lubrecht, B.N.J. Persson, K. Harris, A. Bennett, K. Schulze, S. Rohde, P. Ifju, W.G. Sawyer, T. Angelini, H.A. Esfahani, M. Kadkhodaei, S. Akbarzadeh, J.J. Wu, G. Vorlaufer, A. Vernes, S. Solhjoo, A.I. Vakis, R.L. Jackson, Y. Xu, J. Streator, A. Rostami, D. Dini, S. Medina, G. Carbone, F. Bottiglione, L. Afferrante, J. Monti, L. Pastewka, M.O. Robbins, J.A. Greenwood, Tribology Letters **65**(4) (2017). doi:10.1007/s11249-017-0900-2. URL `https://doi.org/10.1007/s11249-017-0900-2`

[11] I.A. Polonsky, L.M. Keer, Journal of Tribology **122**(1), 36 (2000). doi:10.1115/1.555326. URL `https://doi.org/10.1115/1.555326`

[12] J.J. Wu, Journal of Physics D: Applied Physics **39**(9), 1899 (2006). doi:10.1088/0022-3727/39/9/027. URL `https://doi.org/10.1088/0022-3727/39/9/027`

[13] S. Medina, D. Dini, International Journal of Solids and Structures **51**(14), 2620 (2014). doi:10.1016/j.ijsolstr.2014.03.033. URL `https://doi.org/10.1016/j.ijsolstr.2014.03.033`

[14] V. Rey, G. Anciaux, J.F. Molinari, Computational Mechanics **60**(1), 69 (2017). doi:10.1007/s00466-017-1392-5. URL `https://doi.org/10.1007/s00466-017-1392-5`

[15] S.P. Venugopalan, L. Nicola, M.H. Müser, Modelling and Simulation in Materials Science and Engineering **25**(3), 034001 (2017)

[16] G. Carbone, B. Lorenz, B. Persson, A. Wohlers, The European Physical Journal E **29**(3), 275 (2009)

[17] L. Kogut, L. Etsion, Tribol. Trans. **46**, 383 (2003)

[18] P. Koskinen, S. Malola, H. Häkkinen, Physical Review Letters **101**(11), 115502 (2008). doi:10.1103/physrevlett.101.115502. URL `https://doi.org/10.1103/physrevlett.101.115502`

[19] J.A. Flores-Livas, A. Sanna, E.K. Gross, The European Physical Journal B **89**(3) (2016). doi:10.1140/epjb/e2016-70020-0. URL `https://doi.org/10.1140/epjb/e2016-70020-0`

[20] M. Walter, M. Moseler, The Journal of Physical Chemistry C **113**(36), 15834 (2009). doi:10.1021/jp9023298. URL `https://doi.org/10.1021/jp9023298`

[21] J.A. Flores-Livas, M. Amsler, T.J. Lenosky, L. Lehtovaara, S. Botti, M.A.L. Marques, S. Goedecker, Physical Review Letters **108**(11), 117004 (2012). doi:10.1103/physrevlett.108.117004. URL `https://doi.org/10.1103/physrevlett.108.117004`

[22] L. Pastewka, P. Pou, R. Pérez, P. Gumbsch, M. Moseler, Physical Review B **78**(16), 161402 (2008). doi:10.1103/physrevb.78.161402. URL `https://doi.org/10.1103/physrevb.78.161402`

[23] S. Lee, J. Im, Y. Yoo, E. Bitzek, D. Kiener, G. Richter, B. Kim, S.H. Oh, Nature Communications **5**(1) (2014). doi:10.1038/ncomms4033. URL `https://doi.org/10.1038/ncomms4033`

[24] N. Nouri, S. Ziaei-Rad, Macromolecules **44**(13), 5481 (2011). doi:10.1021/ma2005519. URL `https://doi.org/10.1021/ma2005519`

[25] D. Sheppard, R. Terrell, G. Henkelman, The Journal of Chemical Physics **128**(13), 134106 (2008). doi:10.1063/1.2841941. URL `https://doi.org/10.1063/1.2841941`

[26] D. Asenjo, F. Paillusson, D. Frenkel, Physical Review Letters **112**(9) (2014). doi:10.1103/physrevlett.112.098002. URL `https://doi.org/10.1103/physrevlett.112.098002`

[27] E. DeGiuli, E. Lerner, C. Brito, M. Wyart, Proceedings of

the National Academy of Sciences **111**(48), 17054 (2014). doi:10.1073/pnas.1415298111. URL https://doi.org/10.1073/pnas.1415298111

[28] S. Plimpton, Journal of Computational Physics **117**(1), 1 (1995). doi:10.1006/jcph.1995.1039. URL https://doi.org/10.1006/jcph.1995.1039

[29] A.H. Larsen, J.J. Mortensen, J. Blomqvist, I.E. Castelli, R. Christensen, M. Dułak, J. Friis, M.N. Groves, B. Hammer, C. Hargus, E.D. Hermes, P.C. Jennings, P.B. Jensen, J. Kermode, 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, K.W. Jacobsen, Journal of Physics: Condensed Matter **29**(27), 273002 (2017). doi:10.1088/1361-648x/aa680e. URL https://doi.org/10.1088/1361-648x/aa680e

[30] M.D.B. Bouchet, C. Matta, B. Vacher, T. Le-Mogne, J. Martin, J. von Lautz, T. Ma, L. Pastewka, J. Otschik, P. Gumbsch, M. Moseler, Carbon **87**, 317 (2015). doi:10.1016/j.carbon.2015.02.041. URL https://doi.org/10.1016/j.carbon.2015.02.041

[31] Y. Jiang, J.A. Harrison, J.D. Schall, K.E. Ryan, R.W. Carpick, K.T. Turner, Tribology Letters **65**(3) (2017). doi:10.1007/s11249-017-0857-1. URL https://doi.org/10.1007/s11249-017-0857-1

[32] D. Asenjo, J.D. Stevenson, D.J. Wales, D. Frenkel, The Journal of Physical Chemistry B **117**(42), 12717 (2013). doi:10.1021/jp312457a. URL https://doi.org/10.1021/jp312457a

[33] C.P. Goodrich, S. Dagois-Bohy, B.P. Tighe, M. van Hecke, A.J. Liu, S.R. Nagel, Physical Review E **90**(2), 022138 (2014). doi:10.1103/physreve.90.022138. URL https://doi.org/10.1103/physreve.90.022138

[34] B. Eidel, A. Stukowski, J. Schröder, PAMM **11**(1), 509 (2011). doi:10.1002/pamm.201110246. URL https://doi.org/10.1002/pamm.201110246