

Simulating granular materials by energy minimization

D. Krijgsman¹ · S. Luding¹

Received: 30 April 2015 / Revised: 25 January 2016 / Accepted: 5 February 2016 / Published online: 14 March 2016
© The Author(s) 2016. This article is published with open access at Springerlink.com

Abstract Discrete element methods are extremely helpful in understanding the complex behaviors of granular media, as they give valuable insight into all internal variables of the system. In this paper, a novel discrete element method for performing simulations of granular media is presented, based on the minimization of the potential energy in the system. Contrary to most discrete element methods (i.e., soft-particle method, event-driven method, and non-smooth contact dynamics), the system does not evolve by (approximately) integrating Newton's equations of motion in time, but rather by searching for mechanical equilibrium solutions for the positions of all particles in the system, which is mathematically equivalent to locally minimizing the potential energy. The new method allows for the rapid creation of jammed initial conditions (to be used for further studies) and for the simulation of quasi-static deformation problems. The major advantage of the new method is that it allows for truly static deformations. The system does not evolve with time, but rather with the externally applied strain or load, so that there is no kinetic energy in the system, in contrast to other quasi-static methods. The performance of the algorithm for both types of applications of the method is tested. Therefore we look at the required number of iterations, for the system to converge to a stable solution. For each single iteration, the required computational effort scales linearly with the number of particles. During the process of creating initial conditions, the required number of iterations for two-dimensional systems scales with the square root of the number of particles in

the system. The required number of iterations increases for systems closer to the jamming packing fraction. For a quasi-static pure shear deformation simulation, the results of the new method are validated by regular soft-particle dynamics simulations. The energy minimization algorithm is able to capture the evolution of the isotropic and deviatoric stress and fabric of the system. Both methods converge in the limit of quasi-static deformations, but show interestingly different results otherwise. For a shear amplitude of 4 %, as little as 100 sampling points seems to be a good compromise between accuracy and computational time needed.

Keywords Energy minimization · Quasi-static · Granular materials · Discrete element method · Soft-particle · Initial conditions

1 Introduction

Granular materials are omnipresent in many industrial and natural processes, yet their complex behaviors are far from understood. Discrete element methods are able to capture these complex behaviors and help us understand them. These methods all have in common, that all the particles in the system are individually modeled, and thus give full information on all positions, velocities, and forces. The three most well-known discrete element methods are the soft-particle method, the event-driven method, and the non-smooth contact dynamics.

In the soft-particle method [1], the forces on all particles are calculated from the positions and velocities according to some contact law. Once all these forces are known, they are integrated using Newton's second law. The major advantages of this method are, that it is relatively easy to program and understand and that it allows for a whole range of different

✉ D. Krijgsman
d.krijgsman@utwente.nl

S. Luding
s.luding@utwente.nl

¹ Univeristy of Twente, PO Box 217, 7500 AE Enschede, Netherlands

contact laws. The method has the major disadvantage that it is computational extremely expensive, due to the sharp restriction on the allowed time step. Several attempts have been made to reduce its computational complexity, for example by coarse-graining (i.e., reducing the number of particles and thus degrees of freedom) or by artificially decreasing the particle stiffness (increasing the allowed time step). The second approach can help for dynamic flows, since for simulations usually a certain amount of time needs to be simulated. For quasi-static flows reducing the stiffness of the particles does increase the maximum allowable time step, but it also decreases the wave and information propagation speed. Therefore, in many situations, the required simulation time needs to be increased by the same amount to keep the inertial number constant, i.e., there is no gain.

In the event-driven method [2–4], particle collisions are instantaneous and binary. Once two particles touch each other their contact is handled according to some collision rule and the particles will move apart. After the collision, the algorithm jumps to the time frame where the next two particles interact. The advantage of this method is, that in essence, the particles are modeled as infinitely stiff, without the steep restriction on the time step. However, if a dense state is approached, the allowable time step between two consecutive contacts diminishes, increasing the numerical effort. In the presence of dissipation, this will eventually lead to an artificial phenomenon known as inelastic collapse, i.e., a complete stop of the simulation. Attempts to solve this problem have been made, for example by González [5] and Luding [6].

In the non-smooth contact dynamics [7–10] method, all particles are modeled as perfectly rigid and the interactions are modeled by means of a Signorini condition. Either two particles are not in contact and thus also have no forces associated with them, or they are in contact and the forces associated with the contact are determined from the fact that the two particles cannot overlap. The main advantage of this method is that in the integration process much larger time steps can be made compared to the soft-particle method, due to the way of modeling of particle interactions. Disadvantages are that the uniqueness of the solutions at each time step is not guaranteed and that the computational effort scales super linearly with the number of particles. Krabbenhoft et al. [8] extended the method with an infinite “time step” to perform quasi-static simulations.

As said earlier, soft-particle methods integrate Newton’s equation of motion for each particle, where the forces on each particle consist of damping, stiffness, and external forces. The damping terms mainly depend on the velocities of the particles and the stiffness terms depend mainly on the (history of the) positions of the particles. In quasi-static simulations, however, the applied load changes slowly over time with respect to the inertial forces. Therefore, one can obtain a significant speed up in the simulation if one does not have

to fully resolve the timescale of the inertial forces, but just the one for the external forces. In this paper, this is done by not integrating the movement of all particles over time, but by just searching for equilibrium positions at certain time intervals, where the number of time intervals depends on the rate of change of the external forces and not on the internal forces. Finding these equilibrium positions for the granular system is naturally equivalent to minimizing the potential energy of the system.

One has to note however that soft-particle method simulations with slow deformations will not necessarily yield the same result as energy minimization simulations with a large number of time intervals. This can be understood by imaging the potential energy landscape of the system. At each time step, the system will be in a minimum of this energy landscape. Due to the slowly changing external forces the energy landscape will gradually change. Usually the minimum will just shift slowly with changing external force. In this case, both the soft-particle method and the energy minimization method will move the particles to the exact same positions. However, it is also quite possible that the nature of the equilibrium position changes from a minimum to a saddle point. In this case, the system will slide from the saddle point toward a new minimum position; both methods are then not necessarily ending up in the same energy minimum, no matter how slow the deformation may be.

Luckily in most applications, it is not necessary that both methods converge to exactly the same particle positions. This is, for example, also not the case for soft-particle simulations. Due to the chaotic nature of granular materials it is possible that the positions of particles in two systems with identical initial particle positions diverge over time due to machine inaccuracy. It is, however, important that the trend and characteristics of global properties (like fabric, stress, and potential energy) are the same for all simulation techniques.

Another way of justifying the energy minimization approach is by realizing that soft-particle methods essentially are minimization methods, at least, whenever some damping terms are present. In this paper, however, different, more efficient, techniques are employed to minimize the potential energy. This idea was also used by O’Hern et al. [11], where energy minimization is used to generate packings of particles at different packing fractions, but not to simulate quasi-static deformations.

One can also look on the minimization approach as a kind of multi body contact problem [12–14] or a local-global solution strategy [15]. In these approaches, just as in the non-smooth contact dynamics method, contacts are generally modeled as infinitely stiff. An approach to solve such problems is using* penalty regularization or Lagrange multipliers. Using a square penalty method results into a comparable algorithm as the energy minimization methods suggested in this paper. The philosophy, however for both methods is quite

different. In our method, the goal is not to model infinitely stiff particles, but to work with deformations of the particles at their contacts to resemble their elastic deformation. In the light of such methods (both rigid contacts and those with finite deformations of contacting particles have been fully considered in literature), the solution strategy suggested in this paper, i.e., a trust-region method in combination with the conjugate gradient algorithm is novel to our knowledge. This solutions strategy allows for an elegant algorithm, which does not require any additional, special treatment of rattlers.

The algorithm is analyzed concerning its computational efficiency and convergence. Furthermore, its accuracy has been tested by comparing it with an alternative simulation method, which has only rarely been performed.

In this paper, only a simple contact model (penalty or minimization criterion) without tangential forces is considered. Even though this might be an oversimplification as compared to more physically sophisticated models, it already leads to quite interesting phenomenology in sheared two-dimensional systems. The advantage is that the simple model allows for a more direct interpretation of the performance results of the algorithm, since no additional non-linearity arise from the contact model. When friction and cohesion are employed even more interesting behavior is observed. [15–18].

In Sect. 2, we will explain the energy minimization process in more detail, while in Sect. 3, we will test the algorithm. The first test will be a simple relaxation to demonstrate the algorithm’s convergence. We will show the scaling of the computational time based on the number of particles and the packing fraction. A second test is performed to compare the energy minimization method with soft-particle simulations. For this, we use a quasi-static pure shear test case, showing that the global behavior of the fabric and stress for both approaches is the same. Furthermore, we report on the scaling of the computational time based on the number of simulation time intervals. Finally, in Sect. 4, we end with some conclusions and recommendations. In this section also a small discussion is started for possible extensions of the algorithm to more complex particle interaction laws.

2 Energy minimization

In the energy minimization method, the granular material is simulated by keeping the potential energy of the system in a local minimum at all “times” (i.e., after each step). In this paper, the potential energy of the granular system is defined as

$$E = \sum_{c \in C} \frac{1}{2} k \delta_c^2, \tag{1}$$

where C is the collection of all contacts in the system, k is the particle stiffness, and δ_c is the overlap between the two particles associated with contact c . This definition of the potential energy corresponds to a soft-particle method where particle–particle interactions are modeled as linear (repulsive) springs. This is one of the simplest contact models available and serves as a good candidate to show the power of the simulation technique. Some more realistic particle system might require more complex particle–particle interactions. A potential approach for modeling such systems will be discussed in the outlook and is a good candidate for further research. Furthermore, in this paper a constant particle stiffness is used, but the model in general allows for a variable stiffness (i.e., depending on particle sizes) without major alterations. Additional simulations have been performed to check the scaling with the particle radius. Results have shown no influence of the particle size on the resulting potential energy and stress (data not shown).

Finding a (local) minimum in the potential energy is difficult because it is highly non-linear in the positions of all the particles. This non-linearity mainly comes from the opening and closing of contacts, but also, in a lesser extent due to particles rotating around and sliding past each other. Furthermore, realistic granular systems consist of countless numbers of particles, leading to a huge number of degrees of freedom, making the minimization problem computational expensive.

One method which is quite capable of solving those kind of problems is the trust-region approach coupled to Steihaugs algorithm for solving the trust-region problem [19,20]. Both are described in Sects. 2.2 and 2.3 respectively, however, these methods require the Jacobian (first derivative or force vector) and Hessian (second derivative or stiffness matrix, in more dynamic situations related to the dynamic matrix) of the system to be known and we shall derive them first.

2.1 Jacobian and Hessian

For a single contact c , the potential energy is defined as

$$E_c = \frac{1}{2} k \delta_c^2, \tag{2}$$

with δ_c the overlap, which is defined as

$$\delta_c = |\mathbf{x}_1 - \mathbf{x}_2| - r_1 - r_2, \tag{3}$$

where \mathbf{x}_1 and \mathbf{x}_2 are the positions of particle 1 and 2 and r_1 and r_2 their respective radii. The first derivatives of the single contact potential energy are thus

$$\begin{aligned}\frac{\partial E_c}{\partial \mathbf{x}_1} &= \mathbf{f}_c = k\delta_c \mathbf{n}_c \\ \frac{\partial E_c}{\partial \mathbf{x}_2} &= -\mathbf{f}_c = -k\delta_c \mathbf{n}_c,\end{aligned}\quad (4)$$

with

$$\mathbf{n}_c = \frac{\mathbf{x}_1 - \mathbf{x}_2}{L_c}.\quad (5)$$

and

$$L_c = |\mathbf{x}_1 - \mathbf{x}_2|.\quad (6)$$

The second derivatives

$$\begin{aligned}\frac{\partial E_c}{\partial \mathbf{x}_1 \partial \mathbf{x}_1} &= \frac{\partial E_c}{\partial \mathbf{x}_2 \partial \mathbf{x}_2} \\ &= \bar{\bar{K}}_c = k \left(\mathbf{n}_c \otimes \mathbf{n}_c \left(1 - \frac{\delta_c}{L_c} \right) + \bar{\bar{I}} \frac{\delta_c}{L_c} \right) \\ \frac{\partial E_c}{\partial \mathbf{x}_1 \partial \mathbf{x}_2} &= -\bar{\bar{K}}_c = -k \left(\mathbf{n}_c \otimes \mathbf{n}_c \left(1 - \frac{\delta_c}{L_c} \right) + \bar{\bar{I}} \frac{\delta_c}{L_c} \right),\end{aligned}\quad (7)$$

with \otimes the dyadic product of two vectors and $\bar{\bar{I}}$ the identity matrix. From these single contact force vector \mathbf{f}_c and stiffness matrix $\bar{\bar{K}}_c$ we can assemble the full force vector \mathbf{f} and stiffness matrix $\bar{\bar{K}}$ of the system and apply the minimization algorithm.

2.2 Trust-region

Trust-region methods are a general type of methods used to minimize or maximize an objective function. They assume that in a subset of the domain (i.e., the trust-region), the objective function can be modeled by a much simpler model function. If an adequate model of the objective function is found within the trust-region, the region can be expanded; conversely if the approximation is insufficient the region should be contracted.

In this paper, the trust-region is defined as a higher dimensional sphere with radius Δ around the current positions of all particles, \mathbf{x} . From these current positions an improvement step \mathbf{p} is made to reduce the potential energy. To calculate this improvement, we use the model function, $m(\mathbf{p})$, which is simply a quadratic Taylor expansion of the objective function (i.e., the potential energy):

$$E(\mathbf{x} + \mathbf{p}) \approx m(\mathbf{p}) \equiv E(\mathbf{x}) + \mathbf{f}^T \mathbf{p} + \frac{1}{2} \mathbf{p}^T \bar{\bar{K}} \mathbf{p}\quad (8)$$

with \mathbf{f} the force vector and $\bar{\bar{K}}$ the stiffness matrix as defined in the previous section. The problem of finding a minimum in the potential energy is now reduced to finding a minimum of the model function, within the trust-region:

$$\min_{\mathbf{p}} m(\mathbf{p}) \text{ with } |\mathbf{p}| \leq \Delta.\quad (9)$$

where $||$ is the l^2 or Euclidean norm. To solve this trust-region problem we use the Steihaug algorithm (see Sect. 2.3). Once we have calculated an (approximate) solution, \mathbf{p} , we check the quality, ρ , of our trusted region by comparing the reduction in energy of our approximation with the real reduction in energy

$$\rho = \frac{E(\mathbf{x}) - E(\mathbf{x} + \mathbf{p})}{m(\mathbf{0}) - m(\mathbf{p})}.\quad (10)$$

If the quality is low (i.e., $\rho < \frac{1}{4}$) the model function is not an adequate model of the objective function and we need to reduce the size of our trust-region. On the other hand, when the quality is good (i.e., $\rho > \frac{3}{4}$) and the solution lies on the border of the trust-region ($|\mathbf{p}| = \Delta$), we expand it. Furthermore, we accept updates only if they reduce the potential energy (or equivalently when $\rho > 0$). This trust-region algorithm is described in more detail in Algorithm 1.

Data: Given \mathbf{x} and Δ

Result: A local minimum in the potential energy (Eq. 1)

while *Not converged* **do**

 Calculate \mathbf{f} and $\bar{\bar{K}}$;

 Obtain \mathbf{p} by (approximately) solving (Eq. 9);

 Evaluate ρ from (Eq. 10);

if $\rho > 0$ **then** // Accept improvement

 | set $\mathbf{x} = \mathbf{x} + \mathbf{p}$;

end

if $\rho < \frac{1}{4}$ **then** // Low quality

 | set $\Delta = \frac{1}{4} |\mathbf{p}|$;

end

if $\rho > \frac{3}{4}$ **and** $|\mathbf{p}| = \Delta$ **then** // High quality

 | set $\Delta = 2\Delta$;

end

end

return \mathbf{x}

Algorithm 1: Trust-Region

2.3 Steihaug algorithm

The Steihaug algorithm tries to find a solution to the trust-region problem (Eq. 9). It is based on the conjugate gradient algorithm, which is able to solve quadratic unbounded minimization problems for positive definite Hessian. However since our domain is bounded (i.e., it has to be within the trust-region) and the Hessian is not necessarily positive definite, two additional exit constraints are added. The first kicks in if the iteration would give a solution outside the trust-region. The second kicks in if the current search direction is one of negative curvature in the Hessian. In both cases, the algorithm returns with a final solution on the boundary of the

trust-region domain. This Steihaug algorithm is described in more detail in Algorithm 2. In this algorithm, \mathbf{p}_i is the solution, \mathbf{r}_i the residual and \mathbf{d}_i the search direction.

```

Data: Given  $\mathbf{f}$ ,  $\bar{K}$ ,  $\Delta$  and  $\epsilon$ 
Result: An approximate solution to Eq. 9
set  $\mathbf{p}_0 = 0$ ; // Solution
set  $\mathbf{r}_0 = \mathbf{f}$ ; // Residual
set  $\mathbf{d}_0 = -\mathbf{r}_0$ ; // Search direction
for  $j = 1, 2, 3, \dots$  do
  if  $\mathbf{d}_j^T \bar{K} \mathbf{d}_j \leq 0$  then // Negative curvature
    Find  $\tau$  such that  $\mathbf{p} = \mathbf{p}_j + \tau \mathbf{d}_j$  minimizes  $m(\mathbf{p})$ , within
     $|\mathbf{p}| \leq \Delta$ ;
    return  $\mathbf{p}_j + \tau \mathbf{d}_j$ ; // Return solution on edge
    of trust-region
  end
  set  $\alpha_j = \frac{\mathbf{r}_j^T \mathbf{r}_j}{\mathbf{d}_j^T \bar{K} \mathbf{d}_j}$ ;
  set  $\mathbf{p}_{j+1} = \mathbf{p}_j + \alpha_j \mathbf{d}_j$ ; // Update solution
  set  $\mathbf{r}_{j+1} = \mathbf{r}_j + \alpha_j \bar{K} \mathbf{d}_j$ ; // Update residual
  if  $|\mathbf{p}_{j+1}| \geq \Delta$  then // Solution outside
  trust-region
    Find  $\tau$  such that  $\mathbf{p} = \mathbf{p}_j + \tau \mathbf{d}_j$  minimizes  $m(\mathbf{p})$ , within
     $|\mathbf{p}| \leq \Delta$ ;
    return  $\mathbf{p}_j + \tau \mathbf{d}_j$ ; // Return solution on edge
    of trust-region
  end
  if  $|\mathbf{r}_{j+1}| < \epsilon |\mathbf{r}_0|$  then // Accuracy reached
    return  $\mathbf{p}_{j+1}$ ;
  end
  set  $\beta_{j+1} = \frac{\mathbf{r}_{j+1}^T \mathbf{r}_{j+1}}{\mathbf{r}_j^T \mathbf{r}_j}$ ;
  set  $\mathbf{d}_{j+1} = \mathbf{r}_{j+1} + \beta_{j+1} \mathbf{d}_j$ ; // Update search
  direction
end

```

Algorithm 2: Steihaug algorithm

2.4 Rattlers

The trust-region and Steihaug algorithms require no special action for rattlers. This can easily be understood by distinguishing the case of real rattlers (i.e., particles without any contacts) and pseudo rattlers (i.e., particles with too few contacts to be stable). The real rattlers have no contacts at all and thus do not give any contribution to the force vector or stiffness matrix. The Steihaug algorithm will not give any displacement to these particles in the new improvement step. If new strong contacts arise due to the movement of other particles, the quality of the improvement step in the trust-region method will be negative, the improvement will be discarded and a smaller trust-region will be chosen to calculate a better improvement. Due to this smaller trust-region no new strong contact will arise. Pseudo rattlers do have at least one contact and thus will show up in the force vector and stiffness matrix. Naturally, these particles give an easy way to reduce

the potential energy. This is already picked up by the Steihaug algorithm in the first iteration. Due to the force inequality the pseudo rattler will be pushed away from its contact partners. Just as with other particles, the pseudo rattler may pick up new contacts during this displacement, but the trust-region method will make sure that these new contacts are weak and do not give a substantial contribution to the potential energy.

3 Test cases

To demonstrate the method two test cases are studied. The first is a simple relaxation test of a static packing. The goal of this test is to check the performance of the method. Also the scaling of its computational time with the number of particles and the packing fraction are reported. The second test is a pure shear deformation test case, to compare the results of the energy minimization scheme with regular soft-particle simulations. The soft-particle simulations are performed with the open access particle mechanics simulator MercuryDPM [21–24]. In all test cases, we use a uniform distribution of particle radii between r_{\min} and $2r_{\min}$, to prevent crystallization [25], i.e., long-range order and an associated increase in packing efficiency and thus reduction of potential energy [26].

3.1 Relaxation

In the relaxation test, we let a particle system with random initial conditions evolve, until a minimum in potential energy is obtained (see Fig. 1 for a schematic). More specifically, we start with a system of N_p particles, randomly distributed over a unit square domain (i.e., particles can have arbitrarily large overlaps), where the minimum radius, r_{\min} , is chosen such that the packing fraction of the system equals a selected ϕ .

Results from a single relaxation simulation with $N_p = 10^4$ particles at a packing fraction of $\phi = 0.86$ are shown in Fig. 2. In Fig. 2a, the potential energy is plotted as a function of the iteration number on a double logarithmic scale, while in Fig. 2b the reduction in potential energy per iter-

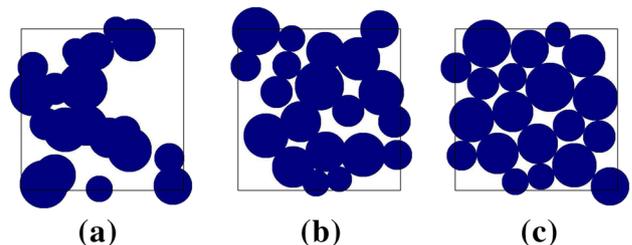


Fig. 1 Schematic evolution of relaxation test. **a** Initial. **b** Middle. **c** Final

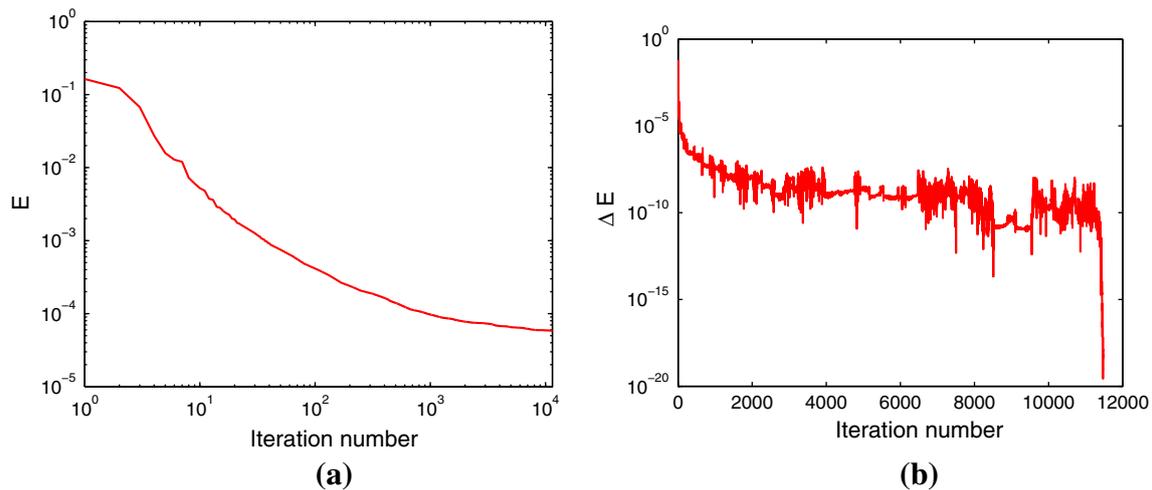


Fig. 2 Potential energy and reduction in potential energy per iteration as a function of the iteration number for random initial conditions with $N_p = 10^4$ and $\phi = 0.86$. **a** Potential energy. **b** Reduction in potential energy

ation is plotted as a function of the iteration number. From these graphs we can clearly distinguish three regimes. In the first few iterations, the potential energy reduces quite quickly, then for most of the simulation the reduction in energy per iteration fluctuates around a constant value, while in the end the energy reduction per iteration drops sharply until computer accuracy is reached.

These three regimes can be explained by looking at what happens in the system. In the beginning, particles can have large overlaps due to the initialization procedure. Local, small rearrangements can easily reduce these overlaps and thus the potential energy, until all particles are roughly homogeneously spread throughout the domain and no large overlaps remain. In a potential energy landscape perspective, the simulation starts somewhere on the slope of a large hill and quickly runs into one of the valleys below.

In the second phase, medium- to large-scale rearrangements are required to further reduce the potential energy, until the system reaches a local minimum. These large-scale rearrangements lead to particles losing or gaining contacts, such that the approximate potential energy is only accurate in a small neighborhood. In the potential energy landscape perspective, the system walks through different valleys looking for a deeper local minimum.

In the final phase, the system is close to a local minimum and quickly ends up there. In this stage, no new contacts are formed or old contacts are broken, therefore the approximate potential energy is quite accurate, such that the conjugate gradient algorithm almost becomes an exact solver. In the potential energy landscape perspective, the system is able to see the local minimum and walks straight into it without making any detours.

3.1.1 Scaling with packing fraction

The influence of the packing fraction on the required number of iterations to reach the potential energy minimum is shown in Fig. 3a. For this study, simulations with 10^4 particles at different packing fractions are performed. Each simulation is repeated 50 times with different initial conditions. From the graph, it becomes clear that the required number of iterations increases when the system is closer to the average jamming packing fraction ($\phi_j \approx 0.8425$). This behavior has three possible explanations. Firstly, the number of local potential energy minima could decrease as the packing fraction comes closer to the jamming packing fraction. If there are fewer minima, the system has to travel further to reach one of those minima and thus the second phase of the relaxation will take longer. Secondly, there could be the same number of local minima, but the valleys between them can become shallower and curved, resulting into a slower approach. Thirdly, the scale of the eigenmodes of the system changes from local to global once the packing fraction approaches the jamming volume fraction. These larger scale eigenmodes need less energy to activate and will result into a more smooth energy landscape, where the system has to travel further to reach a local minimum. One has to note however that there is no certain jamming packing fraction, but rather a range of packing fractions where a system can be in both jammed and unjammed states. The green line in the figure is a least-squares fit to the data of the form

$$N_{it} = \frac{C_\phi}{|\phi - \phi_j|^{\alpha_\phi}}, \quad (11)$$

with $\alpha_\phi = 1.5$ and $C_\phi = 29$.

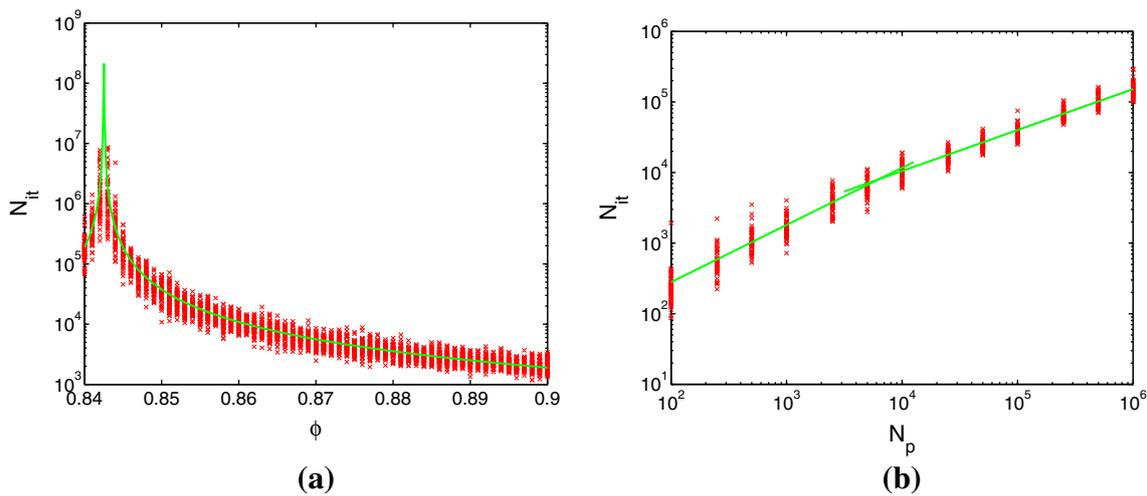


Fig. 3 Required number of iterations to reach a potential energy minimum from random initial conditions, for different packing fractions with $N_p = 10^4$ particles (left) and different number of particles at packing fraction $\phi = 0.86$. Red crosses are single simulations, green lines are

least-squares fits (Eqs. 11 and 12). **a** Influence of packing fraction at $N_p = 10^4$. **b** Influence of number of particles at $\phi = 0.86$. (Color figure online)

3.1.2 Scaling with number of particles

The influence of the number of particles on the required number of iterations to reach the potential energy minimum is shown in Fig. 3b. For this study, simulations at packing fraction $\phi = 0.86$ with different number of particles are performed. Each simulation is repeated 50 times with different initial conditions. From the graph, it becomes clear that the required number of iterations increases with the number of particles. This can also be explained by the fact that in systems with more particles also more large scale eigenmodes exists. Due to these large-scale eigenmodes the energy landscape flattens, therefore the system has to travel further in the energy landscape resulting in a larger required number of iterations.

Furthermore, from the results we can distinguish two regimes, one for a small system size ($N_p < 7000$) and one for larger system sizes ($N_p > 7000$); both are fitted by a power law function.

$$N_{it} = C_p N_p^{\alpha_p} \tag{12}$$

For the small systems sizes, the exponent $\alpha_p = 0.81$ and for large systems sizes $\alpha_p = 0.56$ (values for C_p are 6.68 and 50.12 for small and large systems sizes, respectively). The fact that for large system sizes the numbers of iterations approximately scales with the square root of the number of particles, confirms the idea that the required number of iterations depends on the scale of the largest eigenmode. Since in two-dimensional systems, the system size scales with the square root of the number of particles. This also leads us to believe that for three-dimensional systems, the required

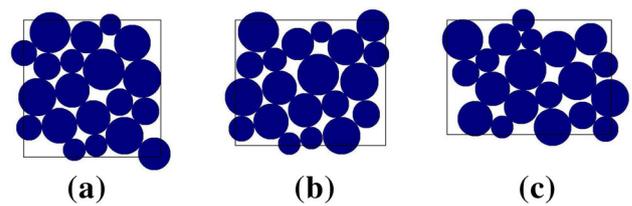


Fig. 4 Schematic evolution of shear test (shear amplitude is exaggerated). **a**Initial. **b** Middle. **c** Final

number of iterations should scale with the cubic root of the number of particles, however, this has not yet been confirmed.

3.2 Shear

In this test, we compare the energy minimization technique with regular soft-particle dynamics simulations. Therefore, it is chosen to run a two-dimensional pure shear test experiment with periodic boundaries on one of the packings generated with the energy minimization scheme. A shear test case is chosen since it usually gives a more severe restriction on the required inertial number for quasi-static flows than compression or expansion tests. More specifically, we use a system with $N_p = 10^4$ particles at packing fraction $\phi = 0.86$ and shear it, keeping the volume conserved, until its true shear strain equals $\gamma_{max} = 0.04$, where γ is defined as

$$\gamma = \log \frac{x}{x_0} - \log \frac{y}{y_0} \tag{13}$$

where x_0 and y_0 are the initial domain width and length and x and y are the current domain width and length (see Fig. 4).

To compare the two methods, we look at the evolution of the non-dimensional fabric \bar{F} and the stress $\bar{\sigma}$ in the system, defined as (other definitions are possible [27])

$$\bar{F} = \frac{1}{A} \sum_{c \in C} \pi (r_1^2 + r_2^2) \mathbf{n}_c \otimes \mathbf{n}_c, \quad (14)$$

$$\bar{\sigma} = \frac{k}{A} \sum_{c \in C} \delta_c L_c \mathbf{n}_c \otimes \mathbf{n}_c. \quad (15)$$

with A the surface area between the periodic walls. More specifically, we look at the isotropic and deviatoric parts (denoted by subscripts *iso* and *dev*, respectively) of these tensorial quantities, where the isotropic part is defined as the mean of the two eigenvalues and the deviatoric part is half the difference between the two eigenvalues.

3.2.1 Soft-particle dynamics simulation

In the soft-particle dynamics simulation, the stiffness, mass, and damping are chosen such that the collision time for the smallest particles is $t_c = 10^{-3}$ [s] and that they have a restitution coefficient of $r = 0.8$.¹ Furthermore, a small background damping has been added equal to a tenth of the particle collision damping to provide some additional damping for global particle movements. The periodic walls are moved sinusoidal, such that the system undergoes a smooth transition from static to moving and no shock waves are triggered. For more details see, for example, Krijgsman [28] and Luding [29].

The deformation speed is determined by the number, N_{sim} , of time steps ($\Delta t = t_c/50 = 2 \times 10^{-5}$ [s]) simulated before the maximum shear strain amplitude is reached. This number of time steps has been varied to test how slow we need to go to be in the quasi-static regime with maximal strain rate $\dot{\gamma}_{\text{max}} \approx 0.06/(N_{\text{sim}} \Delta t)$. In Fig. 5, we plot the kinetic energy of the system versus the shear strain for simulations with different shear rates. The kinetic energy of the system for the fastest simulations is determined by the global motion of the particles. For the slowest simulation, however, the curve has a clear lower bound that represents the global (affine) motion of the particles, but there are also a lot of smaller individual peaks. These peaks correspond to the small and large scale, non-affine, reorganizations that happen in the system during shear. For these slow deformation speeds, these fluctuations clearly dominate the kinetic energy and we can say that the system is in the quasi-static regime.

This observation is also confirmed by looking at the isotropic and deviatoric fabric and stress of the system (see

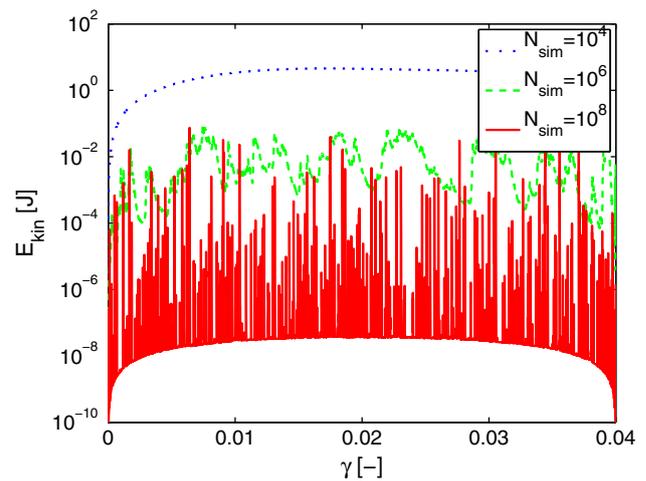


Fig. 5 Kinetic energy as a function of the amount of shear for different shear velocities, during pure shear deformation

Fig. 6). For all deformation rates, we see an increase in the deviatoric values for fabric and stress up to a maximum value. For the isotropic values the global behavior already depends on the deformation speed. At small deformation rates this results in an approximately constant value for the isotropic fabric and a small decreasing value for the isotropic stress (this behavior highly depends on the initial conditions and will be studied in detail elsewhere). Higher deformation rates show a decrease in the isotropic fabric and an increase in the isotropic stress, even though only the static contributions are shown. This shows that for the highest shear velocities the flow is definitely not in the quasi-static regime. When we look closer at the graphs we also see a transition from smooth curves for fast shear to more fluctuating behavior for small deformation rates for both fabric and stress. Especially for the stress at low deformation rates a saw-tooth-like behavior can be seen. The stress generally increases in magnitude with increasing shear strain, but after an amount of stress has been build up, the system is not able to withstand the stress anymore and the particles rearrange, leading to an abrupt reduction of stress. For the faster deformations, the particles almost have no time to settle into their new position before new global rearrangements happen.

3.2.2 Energy minimization

For the energy minimization, there is no such thing as deformation speed or kinetic energy of the system. However, we can choose the number of sampling points (N_{min}) we take between the initial state and the final shear deformation.

Just as for the soft-particle simulations, we plot the isotropic and deviatoric fabric and stress evolution of the system as a function of the shear strain (see Fig. 7). With increasing number of steps, the nature of the curves changes from smooth, for $N_{\text{min}} = 10^1$, to highly fluctuating, for

¹ The results in the paper do not depend qualitatively on r , however, neither a too large, nor a too small restitution coefficient should be used, since both lead to inefficient energy dissipation.

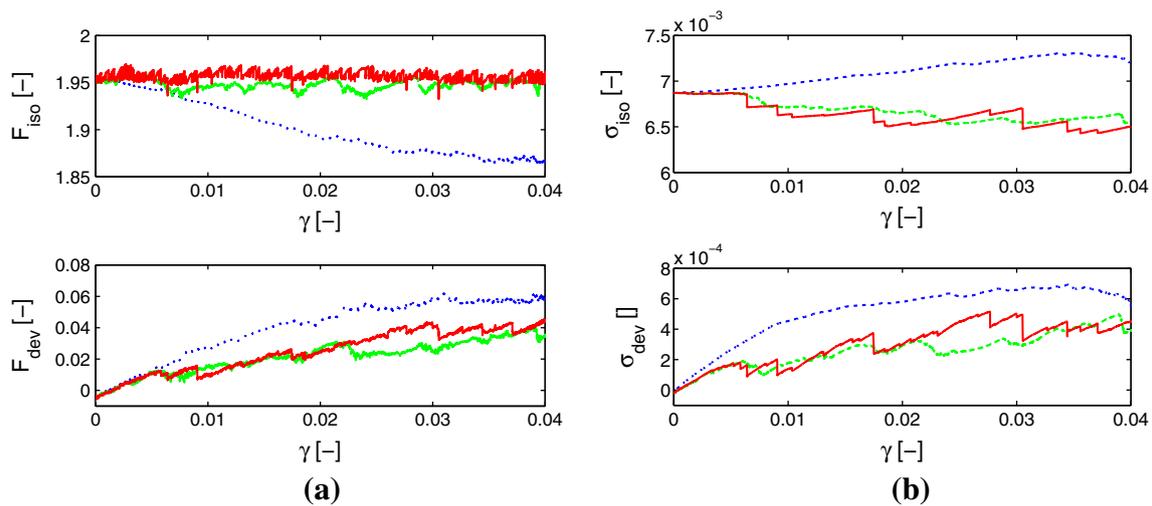


Fig. 6 Comparison of the isotropic and deviatoric fabric and stress between soft-particle dynamics simulations with different numbers of time steps (and thus shear strain rates). Colors and line styles are equal to those in Fig. 5. **a** Fabric. **b** Stress. (Color figure online)

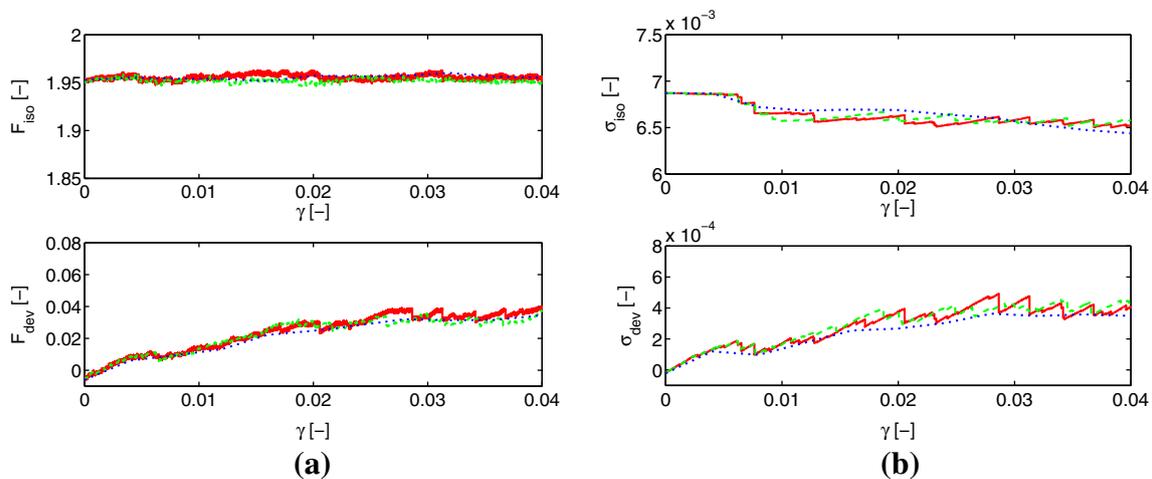


Fig. 7 Comparison of the isotropic and deviatoric fabric and stress between energy minimization simulations with different number of sampling points. Blue dotted lines have $N_{min} = 10$ sampling points, green dashed lines $N_{min} = 10^3$ and red solid lines $N_{min} = 10^5$. **a** Fabric. **b** Stress. (Color figure online)

$N_{min} = 10^5$. The nature of this change, however, differs quite significantly. While for the soft-particle dynamics simulations it was due to the fact that particles have no time to relax after a rearrangement, for the minimization there are simply no additional sampling points. For the minimization in 10 steps, the fabric and stress in the system are only known at 10 different strains and thus automatically result in smooth curves, avoiding the saw-tooth-like behavior. In the minimization algorithm, simulations with a small number of steps already give results quite comparable to the simulations with much more steps, which indicates that the differences between Figs. 6 and 7 are due to the kinetic energy.

When one is not interested in the global trend, but more in specific details at certain points during the simulation, the energy minimization algorithm can provide these quite well. One has to increase the number of sampling points at these

interesting locations, while for the less interesting parts a low number can be used. In this respect, it is also interesting to note that while for the soft-particle mechanics simulation the computational time scales linearly with the number of time steps, the computational time for the energy minimization does not scale linearly, as shown in Fig. 8. In this figure, the total sum of the number of iterations has been plotted versus the number of sampling points. The green line is a power law fit with a power of 0.31, so the required iterations scale approximately with the cubic root of the number of deformation steps.

3.2.3 Comparison

As a last step, the results from both the soft-particle dynamics simulation in the quasi-static case and the energy minimiza-

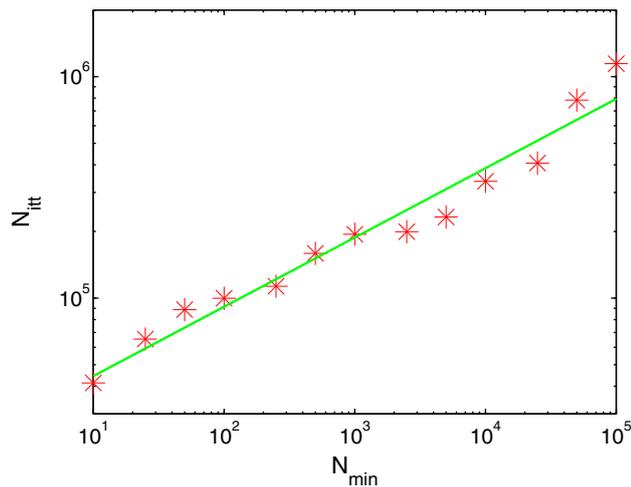


Fig. 8 The total number of iterations used in the minimization scheme versus the number of sampling points. The *red points* are the results of the minimization simulations and the *green line* is a power law fit with a power of 0.31. (Color figure online)

tion method are compared in Fig. 9. Here, the stresses and fabric of the slowest deformation speed in the soft-particle dynamics case and the highest number of sampling points in the minimization algorithm are plotted. As can be seen no perfect agreement between the two has been obtained, but that also was not expected due to the chaotic nature of the system. For sufficiently large systems sizes, all four variables for both simulations show the same trends and structures, indicating that the energy minimization method is able to simulate quasi-static problems. The largest difference here is that the soft-particle dynamics simulations give higher fluctuations in the isotropic fabric. These differences are caused by very weak contacts in the simulation (otherwise it would also have resulted in fluctuations in the stress). Weak contacts

give extremely small forces in the soft-particle dynamics simulations and incredibly slow deformation rates would have to be applied to let the particles move due to this weak forces. In the minimization algorithm, however, all contacts have comparable influence on the next iteration step and the iteration algorithm only quits once the whole system is in perfect equilibrium (up to numerical accuracy), resulting into a much less fluctuating number of contacts and thus a more constant (isotropic) fabric.

The observation that both simulation methods give comparable results is also confirmed by looking at the average isotropic and deviatoric fabric and shear, when continuing the simulation for another 4 % shear. These averages are plotted in Fig. 10 as a function of the simulation time. All curves converge with increasing computational time, which indicates that in the quasi-static regime both approaches give the same results. From the graphs, it also becomes clear that the isotropic fabric gives the most stringent condition on the allowed deformation rate for the soft-particle simulations. For the minimization algorithm, however, both isotropic values already give good results at the first point on the curve (at $t_{\text{comp}} = 1.6 \cdot 10^3$ [s] using just 10 sampling points), while the deviatoric values required more sampling points. There a good compromise between accuracy and computational time seem to be the third point on the curve (at $t_{\text{comp}} = 1.0 \cdot 10^4$ [s] using 100 sampling points).

4 Conclusion

Discrete element or, more general, particle simulation methods are extremely helpful in understanding the complex behaviors of granular media. Different implementations exist, which all excel at a certain range of the possible states in

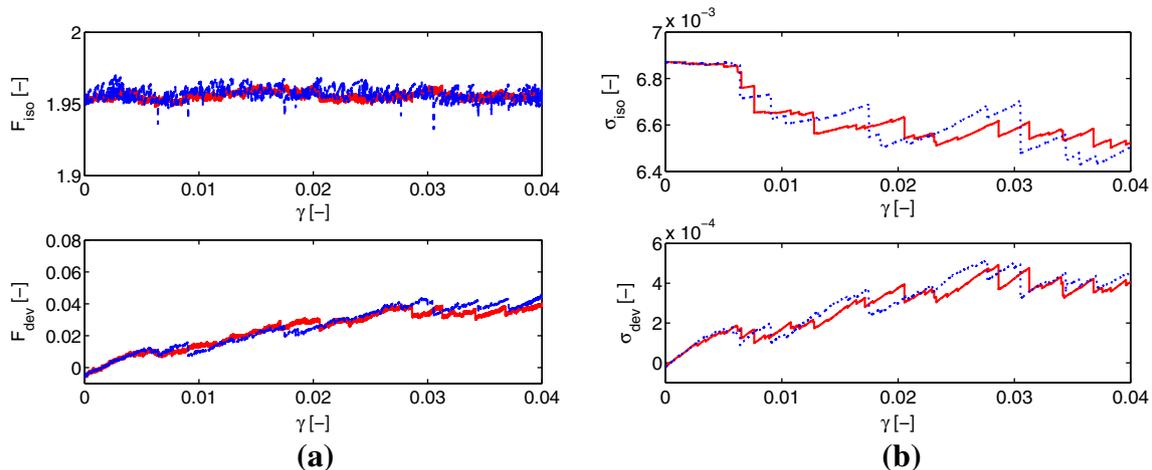


Fig. 9 Comparison of the isotropic and deviatoric fabric and stress between the soft-particle simulations (*blue dotted line*) and the energy minimization (*red solid line*). **a** Fabric. **b** Stress. (Color figure online)

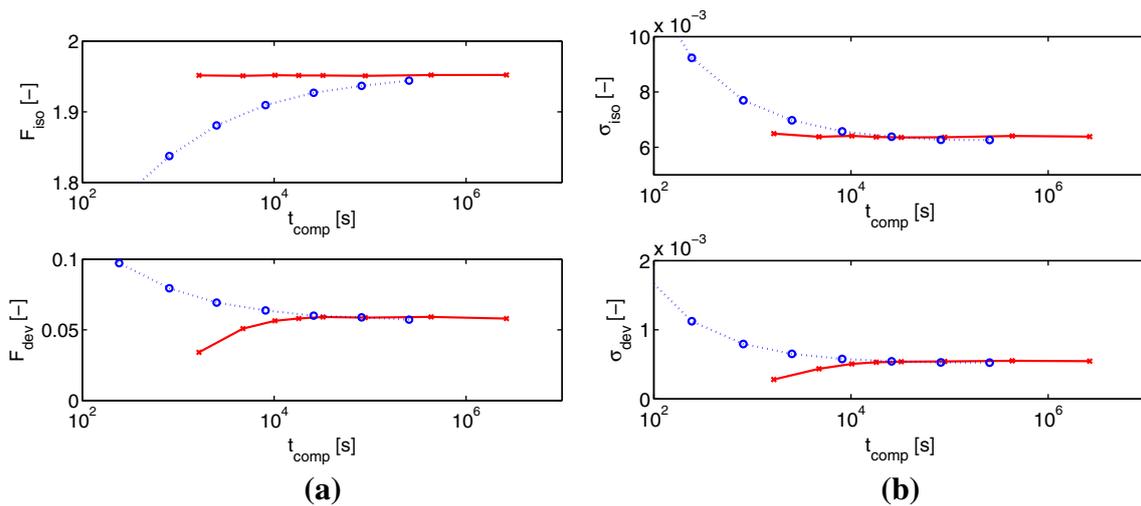


Fig. 10 Comparison of the average isotropic and deviatoric fabric and stress between the soft-particle simulations (*blue dotted line*) and the energy minimization (*red solid line*) as a function of the computational time. For the soft-particle simulations, the number of time steps is var-

ied between 10^3 and $10^{7.5}$. For the energy minimization, the number of sampling points is varied between 10^1 and 10^6 . **a** Fabric. **b** Stress. (Color figure online)

which granular media occur. In this paper, a novel method is developed based on the minimization of the potential energy in the system; the new method is then compared to traditional soft-particle simulations. To minimize the energy we use the trust-region minimization technique in combination with the Steihaug algorithm to solve the trust-region problem. The energy minimization framework is more general in the sense that different minimization techniques could be employed; however, this is not studied here. The method can be applied to either rapidly generate initial conditions or to perform truly static deformation steps of jammed granular materials, which closely resembles very slow quasi-static soft-particle simulations.

One advantage of the method is that it allows for static deformations without fluctuations, since in the method the system does not evolve with inertia during time, but rather with the externally applied strain or load. This makes sure that at no point in the simulation there is any kinetic energy and that the system is always in complete mechanical equilibrium. The stiffness matrix of the system is semi-positive definite and it can be used to calculate the small strain stiffness of the system, without performing any additional simulations. Rattlers that do not contribute to the contact network are handled automatically by the algorithm. Also the method can be easily extended to increase the range of viable simulation types, for example by adding a term to the potential energy to simulate gravity. Compared to the soft-particle method large steps can be made between snapshots, due to the fact that the small scales of particle–particle interactions do not have to be fully resolved. The method differs from the non-smooth contact dynamics in that it allows for overlaps to mimic the deformation of the particles.

The energy minimization method works for both two- and three-dimensional systems, but has been tested for two-dimensional systems only. In the first test, we have shown that, when creating static initial conditions, the required number of iterations in the minimization scheme, for sufficiently large systems sizes, scales with the square root of the number of particles in the system. This has been explained by looking at the spatial-scale of the largest eigenmodes, which in two dimensions also scales with the square root of the number of particles. Furthermore, this leads us to believe that for three-dimensional systems the computational time scales with the cubic root of the number of particles, but this has not been tested. Also the dependence of the required number of iterations on the packing fraction has been tested, where it has been shown that for systems closer to the jamming volume fraction the number of required iterations increases.

A second test has been run to show the performance of the energy minimization method in comparison with quasi-static deformation simulations. Here, we have run a pure shear experiment and compared the new method against soft-particle simulations. It has been shown that in the limit of quasi-static deformations (extremely small deformation rates for the soft-particle simulations and a high number of sampling points in the energy minimization scheme) the values of the isotropic and deviatoric fabric and stress converge for both simulation techniques. For large strains, both methods do not exactly agree, but feature the same mean values and variations in amplitude and frequency. In some cases, we tested that both methods strictly agree for some extremely small strain amplitude (data not shown). Furthermore, while, in the quasi-static regime, for the soft-particle simulation

the isotropic fabric gives the most stringent condition on the allowable deformation speed, in the energy minimization scheme the deviatoric quantities are instrumental for the minimum required number of sampling points. For a large shear amplitude of 4 %, only about 100 sampling points seems to be a good compromise between accuracy and computational time needed in contrast to millions of time steps required in the soft-particle algorithm. Finally, it has been shown that the required computational effort of the minimization algorithm scales with the cubic root of the number of sampling points.

In its current form, the algorithm only is used with one specific definition of the potential energy (or contact model). Extending this to other definitions, which only depend on properties of the current time step (i.e., no history-dependent information is required) is trivial. An example of such a model is the non-constant particle stiffness Hertzian-type contact model. Extensions to models with static friction is expected to be a bit more cumbersome and requires additional research. The problem with these contact models is that friction generally depends on history parameters and that sticking contacts can become sliding even for very small time- or strain-steps and can lead to different phenomenology [17, 18]. A possibility to incorporate such contact models could be to make the modeled potential energy a function of multiple time frames, such that history can be taken into account.

Modeling a Hertz–Mindlin contact for example requires the addition of all tangential spring lengths to the properties of the current time frame. The problem however is, that these lengths are not free variables in the sense that in a new frame any value can be assigned to them. Instead, for each individual contact, the previous tangential springs length has to be used and a contribution has to be added dependent on the movement of the particles and the state the spring is in (sticking or sliding). Essentially, the tangential spring length is integrated in a way similar to what soft-particle mechanics simulations are doing. However, for accurate integration between two time frames, both frames should be similar and a control parameter has to be used to control this accuracy. An idea to apply this is by re-using the trust-region idea, which requires further research that goes beyond the scope of this study.

Acknowledgments This research is supported by the Dutch Technology Foundation STW, which is the applied science division of NWO, and the Technology Program of the Ministry of Economic Affairs, project STW-VICI Grant 10828.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Cundall PA, Strack ODL (1979) A discrete numerical model for granular assemblies. *Geotechnique* 29:47–65
2. Bannerman MN, Sargant R, Lue L (2011) Dynamo: a free $o(n)$ general event-driven molecular dynamics simulator. *J Comput Chem* 32(15):3329–3338
3. Lubachevsky B (1991) How to simulate billiards and similar systems. *J Comput Phys* 94(2):255–283
4. Miller S, Luding S (2004) Event-driven molecular dynamics in parallel. *J Comput Phys* 193(1):306–316
5. González S, Thornton A, Luding S (2011) An event-driven algorithm for fractal cluster formation. *Comput Phys Commun* 182(9):1842–1845
6. Luding S, McNamara S (1998) How to handle the inelastic collapse of a dissipative hard-sphere gas with the tc model. *Granular Matter* 1(3):113–128
7. Jean M (1999) The non-smooth contact dynamics method. *Comput Methods Appl Mech Eng* 177:235–257
8. Krabbenhoft K, Lyamin A, Huang J, da Silva M (2012) Granular contact dynamics using mathematical programming methods. *Comput Geotech* 43:165–176
9. Moreau J (1988) Unilateral contact and dry friction in finite freedom dynamics. In: Moreau J, Panagiotopoulos P (eds) *Nonsmooth mechanics and applications*. International centre for mechanical sciences, vol 302. Springer, Vienna, pp 1–82
10. Radjai F, Richefeu V (2009) Contact dynamics as a nonsmooth discrete element method. *Mech Mater* 41(6):715–728
11. O’Hern C, Langer S, Liu A, Nagel S (2002) Random packings of frictionless particles. *Phys Rev Lett* 88:075,507
12. Alart P, Curnier A (1991) A mixed formulation for frictional contact problems prone to newton like solution methods. *Comput Methods Appl Mech Eng* 92(3):353–375
13. Laursen T, Chawla V (1997) Design of energy conserving algorithms for frictionless dynamic contact problems. *Int J Numer Methods Eng* 40(5):863–886
14. Wriggers P (2006) *Computational contact mechanics*. Springer, Berlin
15. Kaneko K, Terada K, Kyoya T, Kishino Y (2003) Globallocal analysis of granular media in quasi-static equilibrium. *Int J Solids Struct* 40(15):4043–4069
16. Luding S (2004) Micro-macro models for anisotropic granular media. In: Vermeer PA, Ehlers W, Hermann HJ, Ramm E (eds) *Modelling of cohesive-frictional materials*. Proceedings of second international symposium on continuous and discontinuous modelling of cohesive-frictional materials (CDM 2004), CRC Press, Stuttgart, 27–28 Sept 2004, pp 195–206
17. Roux JN, Combe G (2002) Quasistatic rheology and the origins of strain. *C R Phys* 3(2):131–140
18. Luding S (2004) Micro-macro transition for an-isotropic, frictional granular packings. *Int J Solids Struct* 41(21):5821–5836
19. Nocedal J, Wright S (1999) *Numerical optimization*. Springer, New York
20. Steihaug T (1983) The conjugate gradient method and trust regions in large scale optimization. *SIAM J Numer Anal* 20(3):626–637
21. Krijgsman D, Ogarko V, Luding S (2014) Optimal parameters for a hierarchical grid data structure for contact detection in arbitrarily polydisperse particle systems. *Comput Part Mech* 1(3):357–372
22. Thornton A, Krijgsman D, te Voortwis A, Ogarko V, Luding S, Fransen R, González S, Bokhove O, Imole O, Weinhart T (2013) A review of recent work on the discrete particle method at the university of twente: an introduction to the open-source package mercurydpm. In: *Proceedings of the 6th international conference on discrete element methods, DEM-6*, Golden, CO, USA, pp 50–56

23. Thornton A, Weinhart T, Luding S, Bokhove O (2012) Modeling of particle size segregation: calibration using the discrete particle method. *Int J Mod Phys C* 23(08):1240,014
24. Weinhart T, Thornton A, Luding S, Bokhove O (2012) From discrete particles to continuum fields near a boundary. *Granular Matter* 14(2):289–294
25. O’Hern C, Silbert L, Liu A, Nagel S (2003) Jamming at zero temperature and zero applied stress: the epitome of disorder. *Phys Rev E* 68:011,306
26. Luding S (2009) Towards dense, realistic granular media in 2d. *Nonlinearity* 22(12):R101–R106
27. Kumar N, Luding S, Magnanimo V (2014) Macroscopic model with anisotropy based on micromacro information. *Acta Mech* 225(8):2319–2343
28. Krijgsman D, Luding S (2013) 2D cyclic pure shear of granular materials, simulations and model. *AIP Conf Proc* 1542(1):1226–1229
29. Luding S (2008) Introduction to discrete element methods: basic of contact force models and how to perform the micro-macro transition to continuum theory. *Eur J Environ Civil Eng* 12(7–8):785–826