



Berg Huettenmaenn Monatsh (2023) Vol. 168 (6): 287–293
<https://doi.org/10.1007/s00501-023-01363-2>
 © The Author(s) 2023

BHM Berg- und
 Hüttenmännische
 Monatshefte

Speed-up Measures and Associated Problems for Large-scale DEM-Models

Andreas Brandner and Lukas Alexander Koch

Lehrstuhl für Bergbaukunde, Bergtechnik und Bergwirtschaft, Montanuniversität Leoben, Leoben, Austria

Received May 2, 2023; accepted May 16, 2023; published online June 19, 2023

Abstract: This contribution discusses the possibilities to increase the efficiency of large Discrete Element Method (DEM) simulations. Simulations were conducted to test particle upscaling, decreasing shear modulus and using GPU instead of CPU for the computation. The conducted simulations modelled a simple ore extraction process from a defined outlet of a stope in a cave mine. The analysis is based on the influence of the approaches on the computation speed-up and the accuracy of the result. It was found that the real shear modulus could be decreased in the simulation by a factor of 10^3 without interfering strongly with the result, provided that the decreased shear modulus is not smaller than 10^8 Pa. However, a reduction of 10^2 was found to bring the highest speed-up to the present application. For the particle upscaling in this contribution without parameter calibration, the upscaling factor should not exceed 3. Higher upscaling factors affect the results significantly. However, the flow dynamics was already differing for an upscaling factor of 2, even though the flow zones were comparable. Using a GPU instead of a CPU is only recommended if the simulation contains a high number of particles. The speed-up for a simulation with almost 250,000 particles was over 8, but the advantage diminished for less particles. Furthermore, differences in the results between CPU and GPU computation could be found, which could be a starting point for future work. In summary, this research significantly aids in the development of more efficient DEM simulations for large-scale applications, such as cave mining.

Keywords: Discrete Element Method, Ore flow, Particle upscaling, Shear modulus reduction, GPU computation, Large DEM simulations

A. Brandner (✉)
 Lehrstuhl für Bergbaukunde, Bergtechnik und Bergwirtschaft,
 Montanuniversität Leoben,
 Erzherzog-Johann-Straße 3,
 8700 Leoben, Austria
andreas.brandner@stud.unileoben.ac.at

Beschleunigungsmöglichkeiten und einhergehende Probleme für große DEM-Modelle

Zusammenfassung: In diesem Beitrag werden die Möglichkeiten zur Effizienzsteigerung großer Simulationen mit der Diskrete-Elemente-Methode (DEM) diskutiert. Es wurden Simulationen durchgeführt, um die Hochskalierung von Partikeln, die Verringerung des Schermoduls und die Verwendung von GPUs anstelle von CPUs für die Berechnung zu testen. Die durchgeführten Simulationen modellierten einen einfachen Erzabzugsprozess aus einem definierten Abzugspunkt einer Abbaukammer mit kontrolliertem Verbrauch der Hangendschichten. Analysiert wurde der Einfluss der Ansätze auf die Berechnungsgeschwindigkeit und die Genauigkeit der Ergebnisse. Es wurde festgestellt, dass der reale Schermodul in der Simulation um den Faktor 10^3 verringert werden kann, ohne das Ergebnis stark zu beeinträchtigen, wenn der verringerte Schermodul nicht kleiner als 10^8 Pa ist. Es wurde jedoch festgestellt, dass eine Verringerung um 10^2 die höchste Beschleunigung für die vorliegende Anwendung bringt.

Die Hochskalierung der Partikelgröße war in der vorliegenden Anwendung ohne Parameterkalibrierung mit einem Skalierungsfaktor von 3 begrenzt. Höhere Skalierungsfaktoren beeinflussten die Ergebnisse erheblich. Allerdings war die Ausflusssdynamik bereits bei einem Skalierungsfaktor von 2 unterschiedlich, obwohl die Strömungszonen vergleichbar waren. Die Verwendung einer GPU anstelle einer CPU wird nur empfohlen, wenn die Simulation eine große Anzahl von Partikeln enthält. Bei einer Simulation mit fast 250.000 Partikeln betrug der Geschwindigkeitszuwachs mehr als 8, bei weniger Partikeln verringerte sich der Vorteil jedoch. Weiters ließen sich Unterschiede in den Ergebnissen zwischen CPU und GPU-Berechnung feststellen, was ein Anknüpfungspunkt für künftige Forschungsarbeiten sein könnte. Dieser Artikel trägt zur Entwicklung effizienterer DEM-Simulationen für groß angelegte Anwendungen wie den Bruchbau bei.

TABLE 1
Used parameters for the conducted simulations

Poisson's Ratio ν	Density ρ	Coefficient of Restitution	Coefficient of Static Friction	Coefficient of Rolling Friction
0.325 [-]	4700 [kg/m ³]	0.5 [-]	0.75 [-]	0.12 [-]

Schlüsselwörter: Diskrete Elemente Methode, Erzfluss, Partikel-Hochskalierung, Verringerung des Schermoduls, GPU Berechnung, Große DEM Simulationen

1. Introduction

When doing large DEM simulations with millions of particles, either the computation time is usually very high or it is not feasible to simulate the models at all. The increase in computational power, both hard- and software, is helpful, but large DEM simulations still take a considerable amount of time. Three other ways to reduce the computation time are to decrease the shear modulus in the model, to reduce the number of particles by upscaling their size, and to use a graphics card (GPU) instead of an ordinary processor (CPU) to conduct the computation. The general target is to maintain realistic simulations even though the model is simplified in one way or another [1].

The first method of decreasing the shear modulus leads to a reduction of the particle stiffness. A shear modulus lower than the original permits larger particle overlaps in the simulation without creating excessive forces. Therefore, larger time steps can be used, resulting in fewer iterations and a lower computation time. It was seen in literature that the shear modulus could be reduced by a factor of 10^3 without affecting the DEM results strongly. However, the shear modulus should always be kept above 10^8 Pa [2].

The second approach of particle upscaling refers to the practice of replacing the original particles with larger particles. This results in a lower number of particles in the model [3]. This method needs to be used carefully as a too large upscaling factor can affect the results strongly and can no longer model the flow process realistically. This is especially the case when the process is controlled by the particle size. It was found in literature that the maximum upscaling factor is dependent on the application and values of up to 4 are possible without falsifying the results [4].

The third method to increase the speed of DEM simulations is to use a GPU instead of CPU. This approach is especially useful when computing large numbers of particles as a GPU can perform more calculations at the same time than a CPU. Depending on the model and the utilized hardware, speed-up factors of up to 35 could be achieved in the past [5]. A precondition to use GPU calculation is that the software facilitates this massive parallel computing. As a result, software companies are putting effort in developing GPU solver engines for their simulation programs. An example is Altair with the software EDEM, which was used for the simulations in this contribution [6]. The models testing the ways of speeding up DEM simulations mentioned above are discussed in the next section. The focus was on the limits and the effect on the accuracy.

2. Experimental Setup

To investigate the most promising way to speed up computation times, a relatively simple numerical experiment has been set up. The analysis focuses on the effect of particle upscaling, reducing the shear modulus and GPU computation on the solution time and the plausibility of the results. To enable a clean comparison of the different approaches mentioned above, a basic model is established and simulated in the beginning. In the subsequent simulations, only one parameter is changed at a time, e.g. the particle diameter or the shear modulus. This approach facilitates finding the maximum speed-up factor for each method and tests the results from the research performed in the chapters above. The properties of the ore and the equipment, summarized in Table 1, are equivalent to the properties calibrated by Koch [7] for iron ore particles of the mine in Kiruna, Sweden. It should be noted here that the shape of the particles in this contribution are different to the originally calibrated material and the behavior therefore might deviate to some extent. For an easier comparison, all simulations use spheres as particles to minimize the impact of the particles' shape.

The motivation behind this research is the development of the novel mining method called raise caving. In the joint work with LKAB Kiruna, particle flow was identified as a critical point for cave mining methods [8, 9]. The DEM model replicates the material flow in the stope of a cave mine when material is extracted. Thereby, a stope is filled with broken material and the extraction is performed from the draw-point at the bottom of the structure. The dimensions of the cuboid stope are $10 \times 10 \times 20$ m and the outlet drift at the bottom possesses a cross-section of 5×5 m (see Fig. 1). The continuous extraction of the material is conducted 1.5 m outside the brow of the draw-point, which is also the length

Fig. 1: CAD geometry of the stope for the DEM flow simulations


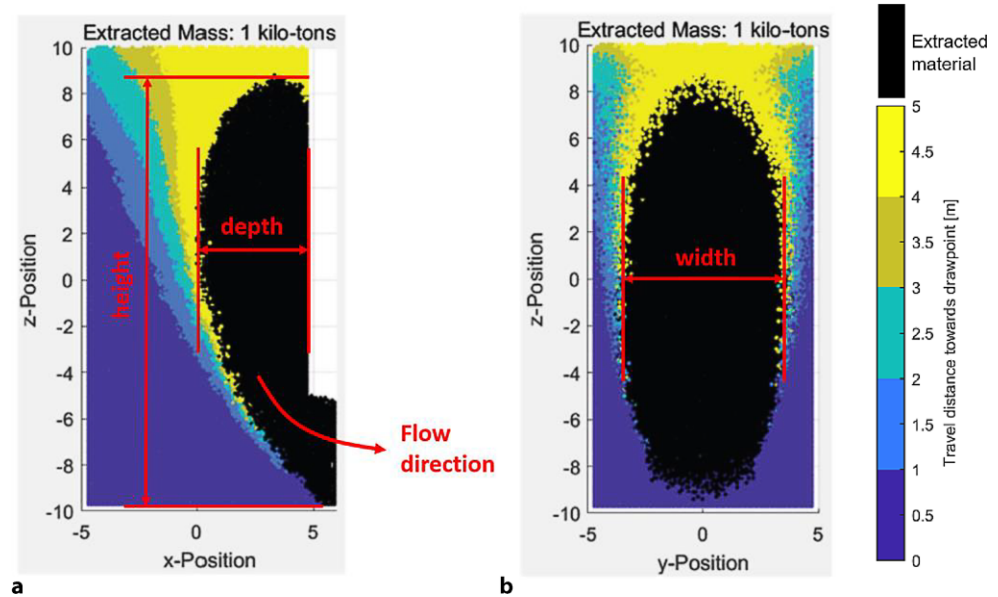


Fig. 2: Measuring of the depth, height and width with the (a) sideview and (b) front view of the flow zone



of the attached drift. The CAD file of the stope is imported to EDEM as a geometry with the same interaction parameters with the particles as for the particles among themselves.

The simulations used 16 cores of an Intel® Core™ i7-5960X [10] Processor with 3.00GHz. For a better understanding of the results, a distinction must be made between model time and computation time. The model time refers to the real duration which is represented in the model. This is used to compare the flow dynamics, meaning the mass flowing out of the stope in a given amount of model time. In contrast to this, the computation time refers to the time it takes to run the simulation. This is the significant time measured and compared to find the speed-up ratios. The speed-up value is determined as computation time of the original model divided by the computation time of the model with the applied speed-up measure.

The so-called Rayleigh time step can be calculated automatically by the software and serves as the basis for the calculation time step [11]. The Rayleigh time step is influenced by the particle radius and the shear modulus among other factors. The calculation time step can be given as

a specific fraction of the Rayleigh time step and usually ranges around 20% or below [12].

To compare the quality of the results, the flow zones inside the stope of each model is analyzed after an extracted mass of 1 kt (see Fig. 2). The non-extracted particles are colored according to the traveled distance towards the drawpoint at the time of plotting the image. The black particles describe the extraction zone, which is the original area of the extracted material. This was the reference area whose width, height, and depth are compared between the models. This was done to see if the results of the simulations are still comparable to the original and therefore the speed-up measures are valid and reliable. Furthermore, the flow dynamics is compared by the discharge rate, meaning the model time it took to extract 1 kt of material.

2.1 Decreasing Shear Modulus

The original shear modulus of the investigated material is $3.7e+10$ Pa. In five different models, this value was reduced

TABLE 2
Parameters and results of the shear modulus reduction simulations

Shear modulus [Pa]	$3.7e+10$	$3.7e+09$	$3.7e+08$	$3.7e+07$	$3.7e+06$
Particle number [-]	246,176	246,880	248,840	271,230	410,100
Bulk density [kg/m^3]	2423.3	2430.2	2449.5	2669.9	4036.9
Rayleigh time step [s]	$1.21e-04$	$3.81e-04$	$1.21e-03$	$3.81e-03$	$1.21e-02$
Percentage used [%]	15	15	15	5	3
Computation time step [s]	$1.81e-05$	$5.71e-05$	$1.81e-04$	$1.91e-04$	$3.61e-04$
Flow depth [m]	4.4	4.3	4.8	4.4	3.6
Flow height [m]	20	19.9	17.2	18	12.3
Flow width [m]	7.2	7.2	7.3	7.2	6.3
Model time [s]	24	24	23.25	21.5	16.5
Computation time [s]	79,860	26,400	8580	8540	9480
Speed-up [-]	1	3.02	9.31	9.35	8.42




Fig. 3: Speed-up for decreasing shear modulus method

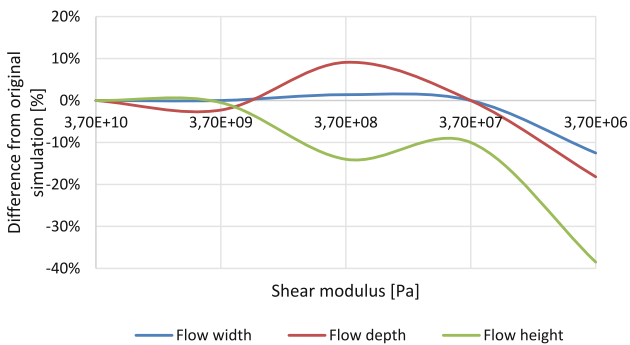


Fig. 4: Decreasing shear modulus accuracy

to $3.7e+06$ Pa. The particle diameter was set to 0.2 m. For the models, 15% of the Rayleigh time step was chosen for the computation. However, the models with a very low shear modulus contained a higher number of particles due to a larger overlap. Thus, they did not stabilize even without external interference, which forced a smaller time step for these models. This and the increased number of particles nullified the effect of the shear modulus decrease. The simulation results including the used time steps, speed-up factors, and flow dimensions can be seen in Table 2. The achieved speed-up for every shear modulus is shown graphically in Fig. 3. The deviation of the results from the original model is shown in Fig. 4.

The results summarized in Table 2 and Fig. 4 highlight the problems with the low shear modulus, see the column on the right-hand side of Table 2. It needs to be mentioned that,

especially for the model with $3.7e+06$, the flow situation is not comparable anymore due to an excessive overlap of the particles and hence a much higher bulk density. This means that the reasons for the problems mentioned require further investigation.

The simulations demonstrated that reducing the shear modulus can accelerate the computation time. The achieved speed-up factors are significant and reach up to 10. A reduction of more than 10^2 could not bring further improvements in the present application. This was mainly due to the weight acting from the particle pile on the lower particles. This led to excessive overlaps of the particles and demanded a decrease of the time step for a small shear modulus.

2.2 Particle Upscaling

Using the findings from the previous section, a shear modulus of $3.7e+08$ is used in the following simulations. The particle diameter was 0.2 m for the basic model and was multiplied with upscaling factors of 2, 3, 4, and 5 for the subsequent simulations. All the simulations are done using a time step which is 10% of the Rayleigh time step. The simulation parameters and results of the simulations can be found in Table 3. It can be seen that the flow dynamics changes as the particles are upscaled. The larger the particles, the longer it takes for 1 kt of the material to flow out of the outlet. In the meantime, the flow zone dimensions are comparable for different upscaling factors after 1 kt of extracted material. This effect has a geometrical background. The opening size is the same for all models, but the particles get larger and hinder each other from flowing out. For applications where the flow dynamics are important, the upscaling approach is therefore inadequate. If only the flow zones are of interest, this approach might still be valid. This indicates that particle upscaling should be evaluated carefully, and a subsequent parameter calibration can be beneficial.

As the reference value for comparison was the extracted mass of 1 kt and the larger particles had a slower outflow, the upscaled particle simulations had to be run for a longer

TABLE 3 Particle upscaling simulation parameters and results					
Upscaling factor [-]	1	2	3	4	5
Particle diameter [m]	0.2	0.4	0.6	0.8	1.0
Particle number [-]	248,840	32,543	9751	3729	1866
Bulk density [kg/m^3]	2495	2562	2591	2349	2296
Rayleigh time step [s]	$1.21e-03$	$2.41e-03$	$3.61e-03$	$4.82e-03$	$6.02e-03$
Percentage used [%]	10	10	10	10	10
Computation time step [s]	$1.21e-04$	$2.41e-04$	$3.61e-04$	$4.82e-04$	$6.02e-04$
Flow width [m]	7	6.7	6.8	6.9	7.5
Flow depth [m]	4.6	4.4	4.4	4.4	4
Flow height [m]	18.7	19.7	19.5	20	17
Model time [s]	23.1	38	192.25	810	4455
Computation time [s]	13,260	1380	1500	2220	6420
Speed-up [-]	1	9.61	8.84	5.97	2.07




Fig. 5: Achievable speed-up values with particle upscaling factors of 1–5




Fig. 6: Particle upscaling accuracy

time than the original model. Hence, not only the flow dynamics was disturbed in these simulations, but also the advantage of upscaling was reduced. The achieved speed-up for every upscaling factor as well as the difference of the results are shown in Figs. 5 and 6. The maximum speed-up value is 9.61 and can be achieved with an upscaling factor of 2. When the upscaling factor is increased to 3, the long model time for the flow out results in a higher computation time, which only achieves a speed-up factor of 8.84. The trend of longer simulations continues when using larger upscaling factors of 4 and 5.

2.3 GPU Calculation

The same upscaling simulations as before were also simulated with a GPU to see the effect on the computation time. The used GPU was the AMD Radeon VII [13]. The results of the GPU simulations can be seen in Table 4. The comparison of the speed-up factors for the particle upscaling simulations on GPU can be seen in Fig. 7. The resulting accuracy is shown in Fig. 8.

The upscaling factor of 2 resulted in a speed-up factor of almost 2 compared to the original simulation. Like in the CPU calculations, a disturbed flow dynamics was observable. This means that 1 kt of mass needed more time for larger particles to flow out than for smaller ones. Hence, no other upscaling factor could achieve a faster computation than the original model. A second reason for the decrease in performance is inherent to the GPU, which holds advantages mainly for a large number of particles. As there are fewer particles included in the model with upscaled particles, the advantage of the fast GPU vanishes.

Comparing the GPU speed to the one of the CPU, the advantage is also mainly present for the original model with a high number of particles. This can be seen in Fig. 9, in which the speed-up of GPU over CPU is shown. Higher upscaling factors of 3 and 4 with fewer particles involved in the model show a superiority of the CPU. That underlines the research, which shows that cost intensive GPUs are only effective for a high number of particles. At an upscaling factor of 3, the simulation time is even increasing for the GPU, making the particle upscaling method inadequate in this application.

The upscaling factor of 5 could not be computed until the 1 kt extraction because a hang-up occurred, which was not the case for the CPU calculation. It becomes apparent that the GPU calculation gives slightly different results to the ones of the CPU. There are different reasons why CPU and GPU simulations might deliver slightly different results despite starting from the identical time step. When using the OpenCL solver, no simulation data is transferred between the GPU and the CPU. Therefore, the time steps that require a data save will be run fully on the CPU. The combination of

TABLE 4 Properties and results of particle upscaling simulations using GPU					
Upscaling factor [-]	1	2	3	4	5
Particle diameter [m]	0.2	0.4	0.6	0.8	1.0
Particle number [-]	248,840	32,543	9751	3729	1866
Bulk density [kg/m ³]	2495	2562	2591	2349	2296
Rayleigh time step [s]	1.21e-03	2.41e-03	3.61e-03	4.82e-03	6.02e-03
Percentage used [%]	10	10	10	10	10
Time step [s]	1.21e-04	2.41e-04	3.61e-04	4.82e-04	6.02e-04
Flow width [m]	7	7.3	7.4	7.5	–
Flow depth [m]	4.7	4.4	4.4	4	–
Flow height [m]	18.8	20	19.3	19.6	–
Model time [s]	23.25	38	140	810	–
Computation time [s]	1620	840	1740	8460	–
Speed-up [-]	1	1.93	0.93	0.19	–
Speed-up vs CPU [-]	8.19	1.64	0.86	0.26	–




Fig. 7: Speed-up for particle upscaling simulations with GPU




Fig. 8: GPU particle upscaling accuracy




Fig. 9: The comparison of the speed-up of GPU over CPU for different up-scaling factors

CPU and GPU may account for the observed differences in the results [14]. Additionally GPUs produce approximation and round-off errors, also called truncation errors, which may lead to less accurate results [15]. Generally, this finding of different results between CPU and GPU calculation should be a point for future investigations.

3. Summary and Discussion

The objective of this contribution is to improve the understanding of speed-up measures for large DEM models using the example of a cave mining simulation. The models used in this contribution recreated the extraction process from the bottom of a filled stope. Using simplified, spherical particles helped to ensure the comparability of the different simulations. Without a proper calibration, the results of flow zones should be seen qualitatively as the focus was rather on the speed-up. Three methods were tested, namely particle upscaling, decreasing shear modulus, and using GPU instead of CPU for computation.

It was found in past DEM research activities that using the correct shear modulus is not significant for realistic DEM results. However, a smaller shear modulus can decrease the computation time as it allows larger time steps. Literature however, showed that the results are highly questionable when the shear modulus of rocks becomes lower than 10^8 Pa. This was also supported by the simulations in this contribution. A reduction of the shear modulus larger than by the factor of 10^2 could not bring any improvement in computation time for the present application. This was due to the weight pressure of the particle pile over the lowest portion of particles. This led to high overlaps in the model of low shear modulus and demanded a lower time-step. Also, the bulk density changed strongly due to the overlapping particles when the shear modulus became too small.

The particle upscaling simulations showed that it is only reliable up to a certain point and requires parameter calibration beyond that. Upscaling factors of up to 3 delivered similar flow zone results as the original model. However, flow dynamics differed, and the outflow was slower for larger particles. For applications where the flow dynamics is of interest, the particle upscaling method is therefore inadequate, at least without subsequent parameter calibration. For the present application where just the flow zones were investigated, an upscaling factor of 2 was shown to be optimal. Further upscaling reduced the speed again due to longer outflow times.

Using a GPU instead of a CPU for the upscaling simulations was seen to only be advantageous with a high number of particles. While a simulation with almost 250,000 particles resulted in a speed-up of over 8, the advantage diminished for less particles. For smaller simulations with 10,000 particles or fewer the CPU showed a superiority. Furthermore, slight differences in results could be observed between CPU and GPU calculation, which should be a point of future investigation.

Overall, the contribution provides insights into the benefits and limits of different methods to improve the efficiency of large DEM models using the example of a cave mining stope. The contribution highlights the direction of speed-up methods and simulations and shows the potential for future research to optimize processes without sacrificing accuracy.

Funding. Open access funding provided by Montanuniversität Leoben.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Kruggel-Emden, H., Rickelt, S., Wirtz, S., Scherer, V.: A study on the validity of the multi-sphere Discrete Element Method. *Powder Technol.* **188**(2), 153–165 (2008). <https://doi.org/10.1016/j.powtec.2008.04.037>
- Lommen, S., Schott, D., Lodewijks, G.: DEM speedup: Stiffness effects on behavior of bulk material. *Particuology* **12**, 107–112 (2014). <https://doi.org/10.1016/j.partic.2013.03.006>
- Lommen, S., Mohajeri, M., Lodewijks, G., Schott, D.: DEM particle upscaling for large-scale bulk handling equipment and material interaction. *Powder Technol.* **352**, 273–282 (2019). <https://doi.org/10.1016/j.powtec.2019.04.034>
- Coetzee, C.J.: Particle upscaling: Calibration and validation of the discrete element method. *Powder Technol.* **344**, 487–503 (2019). <https://doi.org/10.1016/j.powtec.2018.12.022>
- Stamato, H., Vanarase, A., Pandey, P., Bharadwaj, R., Hack, M.: Advances in DEM computing which improve predictive capability for processes. Proceedings of AIChE Annual Meeting, San Francisco (2016)
- Curry, D.: 10 Common Questions about EDEM GPU Answered—EDEM—Altair Products—Altair Community. https://community.altair.com/community?id=community_blog&sys_id=c63459d6b0f2810cfd5f6a4e296192a. Accessed: 14 February 2023
- Koch, L.: Internal project report of DEM calibration of iron ore particles in the Kiruna iron ore mine, Internal Project Report (2023). Montanuniversität Leoben
- Ladinig, T., Wimmer, M., Wagner, H., Eds.: Raise caving: a novel mining method for (deep) mass mining, Australian Centre for Geomechanics, Perth (2022)
- Koch, L.A., Ladinig, T., Wimmer, M., Wagner, H., Grynienko, M.: Key issues of ore-flow in Raise Caving (2023) Caving 2022: Proceedings of the Fifth International Conference on Block and Sublevel Caving, Australian Centre for Geomechanics, Perth
- Intel: Intel® Core™ i7-5960X Processor Extreme Edition (20M Cache, up to 3.50 GHz)—Product Specifications | Intel. <https://www.intel.com/content/www/us/en/products/sku/82930/intel-core-i75960x-processor-extreme-edition-20m-cache-up-to-3-50-ghz/specifications.html>. Accessed 19 March 2023).
- DEM Solutions Ltd.: EDEM — Course, E2020: Introduction to Edem (2021)
- Chung, Y.C., Ooi, J.Y.: Influence of Discrete Element Model Parameters on Bulk Behavior of a Granular Solid under Confined Compression, *Particulate Sci. Technol.*, **26**(1), 83–96 (2007). <https://doi.org/10.1080/02726350701759381>
- Advanced Micro Devices, Inc.: AMD Radeon™ VII Drivers & Support | AMD. <https://www.amd.com/de/support/graphics/amd-radeon-2nd-generation-vega/amd-radeon-2nd-generation-vega/amd-radeon-vii>. Accessed 4 July 2023
- Altair Engineering Inc.: GPU OpenCL Guide. https://2022.help.altair.com/2022.1/EDEM/GPU_Guide.htm. Accessed 4 July 2023
- Linderman, M.D., Ho, M., Dill, D.L., Meng, T.H., Nolan, G.P.: Towards Program Optimization through Automated Analysis of Numerical Precision. Proceedings of the ... CGO: International Symposium on Code Generation and Optimization. *Int. Symp. Code Gener. Optim.* **2010**, 230–237 (2010). <https://doi.org/10.1145/1772954.1772987>

Publisher's Note. Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.