



THE UNIVERSITY *of* EDINBURGH

This thesis has been submitted in fulfilment of the requirements for a postgraduate degree (e. g. PhD, MPhil, DClinPsychol) at the University of Edinburgh. Please note the following terms and conditions of use:

- This work is protected by copyright and other intellectual property rights, which are retained by the thesis author, unless otherwise stated.
- A copy can be downloaded for personal non-commercial research or study, without prior permission or charge.
- This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author.
- The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author.
- When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given.

Mixed Discrete Continuum Modelling of Dense Granular Flow

Akhil Mathews



Doctor of Philosophy
The University of Edinburgh

2026

To my friends, family and teachers

Abstract

Granular materials are central to many natural and industrial processes, yet faithfully simulating their behaviour at full scale remains challenging. The Discrete Element Method (DEM) resolves grain-scale physics but becomes costly for large assemblies, while the Finite Element Method (FEM) accelerates computation by treating the material as a continuum at the expense of particle-level detail. Coupling DEM at the particle scale with FEM at the macro scale offers a practical route for large-scale problems. This thesis introduces a hybrid DEM–FEM framework implemented in the open-source KRATOS MULTIPHYSICS platform, in which the two descriptions overlap in a hybrid zone; displacement compatibility is enforced via a penalty, and smoothly varying weights enable seamless transfer of stress and strain across the interface.

Rather than focusing solely on bulk metrics, the verification campaign resolves the mechanics within the overlap and quantifies the influence of coupling parameters. The framework is first exercised on an analytically guided one-dimensional force-transfer problem, then on two benchmarks: uniaxial compression of a monodisperse column and compaction of a polydisperse assembly. A systematic study charts how penalty stiffness, weight-function shape, mesh-to-particle size ratio, and inter-particle friction affect stress/strain transfer and convergence, yielding practical guidelines for robust, resolution-aware coupling across a range of mesh–particle ratios and overlap thicknesses.

A compact, dimensionally consistent scaling for the penalty, $\varepsilon = CE/(dL_c)$, is proposed, where E is an effective modulus, d a characteristic particle size, L_c the hybrid-zone thickness, and C accounts for nonidealities due to weighting, mapping, and engagement. Choosing $C \approx 10$ reliably places the coupling in a high-fidelity regime over diverse mesh/particle ratios and overlap sizes. For the polydisperse system under cyclic loading, a calibrated elasto-plastic Drucker–Prager continuum reproduces DEM responses, capturing friction-dependent stiffening and the evolution of the lateral stress ratio K while maintaining stable overlap behaviour.

Industrial relevance is demonstrated through a large silo-discharge simulation in which flowing regions are resolved with DEM and quasi-static zones are represented by an elasto-plastic FEM. Hybrid predictions of vertical stress, lateral stress ratio, and grain-scale kinematics are benchmarked against a high-resolution DEM reference, confirming process-scale fidelity and indicating where continuum constitutive enrichment (e.g., dilatancy-evolving or nonlocal models) can further reduce residuals. By releasing the implementation through Kratos’s CoSimulation module, the work provides a flexible, open-source foundation for adaptive model switching, multi-rate co-simulation, uncertainty quantification, and multiphysics extensions, alongside new insight into the mechanics of DEM–FEM overlap zones.

Lay Summary

Granular materials—think sand, coffee beans, or wheat kernels—behave in surprisingly complex ways as they pour, compact, or flow. Engineers can predict that motion on a computer in two very different ways. One option follows every single grain, giving a crystal-clear picture but demanding enormous computing power. The other treats the whole mass as a smooth continuum, which is far quicker but misses the small-scale details that often control clogging, wear, and load transfer. This thesis introduces a middle path: a hybrid model that zooms in on individual grains only where needed and zooms out elsewhere, keeping simulations practical without losing the physics that matter.

The framework joins two established tools. The Discrete Element Method (DEM) tracks each grain; the Finite Element Method (FEM) represents the bulk as a continuous solid. They overlap in a small region where the two descriptions are gently “glued” together so that forces and movements pass smoothly between them. Everything is implemented in Kratos Multiphysics—free, open-source software—so researchers and industry can adopt the approach readily.

Earlier hybrid studies mostly reported big-picture outcomes (such as total wall forces) and left open a key question: how well do the two halves actually blend inside the overlap? To answer that, the work first tests a simple one-dimensional column and then a confined compaction of mixed-size grains, letting us see exactly how information passes across the interface. These controlled studies show how different parameters affect the coupling: the strength of the glue, how quickly one model fades into the other, how fine the FEM mesh and the grain size are relative to the overlap length, and the grains’ internal friction. From this map, we derive insight into the overlap zone and practical parameter choices that maintain smooth stress transfer and prevent local stress spikes across a range of mesh discretisations.

A central outcome is a compact rule for choosing the glue strength (the “penalty”): set it in proportion to the material stiffness and inversely to both grain size and overlap thickness. In simple terms, stronger material and thinner overlap require a stiffer glue, while larger grains call for a softer glue. A single constant captures the remaining details of weighting, mapping, and how many layers are engaged; choosing that constant close to 10 consistently lands the method in a high-fidelity regime without over-tightening the coupling.

The approach is then exercised under more realistic, cyclic loading of the mixed-size system. Here an elasto-plastic continuum model (Drucker–Prager) is calibrated against DEM so the continuum can mimic permanent deformation and friction-driven stiffening. The hybrid captures the observed changes in the lateral-to-vertical stress ratio and maintains a stable, well-behaved overlap, indicating that the same recipe scales beyond purely elastic tests.

Finally, a full-scale silo-discharge simulation demonstrates industrial relevance. Fast-moving flow regions are handled grain-by-grain with DEM, while the largely stationary bulk is covered by the calibrated continuum. The hybrid reproduces key trends from a much more expensive, grain-resolved reference and remains computationally manageable, highlighting where modest upgrades to the continuum law could trim the remaining differences.

In short, the thesis delivers both an open-source tool and the know-how to tune it: clear benchmarks, a simple penalty-selection rule, and tested defaults that make resolution-aware granular simulations faster, steadier, and easier to trust. That combination opens the door to better designs and safer operations in agriculture, mining, pharmaceuticals, construction, and even planetary science.

Acknowledgements

First and foremost, I am deeply grateful to my supervisors and collaborators. Prof. Jin Ooi set the intellectual standard for this work while giving me direction and clarity; I am especially thankful for his technical insight—particularly in granular mechanics and discrete-to-continuum modelling—which guided the project’s conceptual design and methodological integrity, keeping the study rigorous and well grounded throughout. I am sincerely thankful to Dr Hongyang Cheng for timely guidance, generous advice, and many thoughtful discussions; his patience, attention to detail, and steady support were invaluable throughout the project. I also wish to acknowledge that the core implementation followed his earlier DEM–FEM coupling framework, which provided a solid foundation for this work and was subsequently adapted and extended. I am also grateful to Dr Miguel Angel Celigueta for thoughtful discussions that ranged from modelling choices to interpreting results; he helped develop this coupling as an independent application in KratosMultiphysics, assisted immensely with debugging, and improved my coding practices. Finally, I thank Dr Stefanos Papanicolopoulos for his useful suggestions that refined the study at important stages.

My appreciation extends to Dr. Antonia Larese, Laura M, Klaus, Pooyan, Philip, Veronika, and wider Kratos communities for sharing their Kratos Multiphysics expertise. I also thank Vanessa Magnanimo and Marina Susani for their support and thanks to the broader TUSAIL family for training events, secondment support, and a generous exchange of ideas. I would like to express my particular gratitude to the examiners and reviewers, Prof. Filipe Teixeira-Dias, Dr Daniel Barreto, and Dr Athanasios Angeloudis. Their insightful suggestions and thorough review of the thesis have significantly improved the technical quality and the overall presentation of this work. I thank the technical and administrative staff in the School of Engineering for their steady, behind-the-scenes support. Thanks to the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 955661 for funding this project. I also acknowledge the practical help of generative-AI tools in assisting with Python scripting for routine data processing.

Beyond academia, I’m grateful to the people who made Edinburgh feel like home. To my first-year flatmates—Shuyin, Dani, Grainne, Pan, and Pushali—thank you for a household of warmth and laughter. Rafael, Asaad, and Thomas for deep conversations, countless pints, and gentle kindness. Thiva, Rahela, Oguzhan, Afhsin, Behrad, Balagopal, Magda, Shon, Caterina and Aasheesh, for a simple, good connection. To my football and badminton friends—thank

you for the mid-week resets. And to all my Latin dance club members, thank you for keeping the weeks bright. My gratitude extends to my family for constant support—and to my mother for sending kisses every day. I also remain grateful to my physics teacher, Rajan, whose creative teaching first sparked my curiosity in science.

Akhil Mathews

Contents

Abstract	iii
Lay Summary	iv
Acknowledgements	vi
Figures and Tables	xii
1 Introduction	1
1.1 General Background	1
1.2 Objectives and Scope of Research	3
1.3 Thesis Structure	4
2 Literature review	6
2.1 Coupling Frameworks for Discrete and Continuum Models	6
2.2 Evolution of Discrete–Continuum Volume Coupling (1998–2020)	13
2.2.1 Foundations and Consolidation (1998–2005)	14
2.2.2 From Theory to Particles: Atomistic/Particle–Continuum Coupling (2007–2010)	14
2.2.3 First Overlap DEM–FEM for Granular Media and Mapping Choices (2007–2012)	16
2.2.4 Stabilisation, Large Deformations, and Operator Design (2010–2012)	17
2.2.5 Fracture, Damage, and Advanced Hybrids (2011–2020)	18
2.2.6 Recent Advances (2021–2024)	18
2.3 Research gaps	26
2.4 Contributions of the Present Work	26
2.5 Summary	28
3 Formulation of DEM, FEM, and coupled interaction laws	30
3.1 Introduction	30
3.2 Discrete Element Method - linear spring-dashpot contact model	30
3.2.1 Constant directional torque (CDT) model	32

3.2.2	Numerical integration of forces	34
3.3	Finite Element Method (Continuum)	35
3.3.1	Weak form of the Cauchy momentum balance	35
3.3.2	Semi-discrete equation of motion	36
3.3.3	Linear Elastic Constitutive Model	38
3.3.4	Drucker–Prager constitutive model (non-associated)	40
3.4	Implementation of DEM-FEM coupling with overlapping domains	42
3.4.1	Coupling approach	42
3.4.2	Mapping	46
3.4.3	Effects of weighting in the hybrid zone	47
3.4.4	DEM-FEM Interaction Mechanisms	49
3.5	Summary	50
4	Implementation in Kratos Multiphysics	52
4.1	Structure of implementation	52
4.2	Coupling utilities	54
4.3	Finite Element Method	56
4.4	Discrete Element Method	57
4.5	Cosimulation Application	58
4.6	Case setup and input files	60
4.7	Summary	63
5	Verification of the hybrid model	65
5.1	Methodology for verification of coupling method	65
5.2	Granular column verification test	66
5.2.1	Model setup and parameters	67
5.2.2	Results	69
5.3	Polydisperse particle compaction test	74
5.3.1	Model setup and parameters	74
5.3.2	Results	77
5.4	Conclusions	84
6	Dimensionless scaling of the penalty parameter	86
6.1	Introduction	86
6.2	Methods in brief	89
6.2.1	Coupling summary	89
6.2.2	Error metric and setups	89

6.3	Coupling penalty: scaling law and empirical behaviour	90
6.3.1	Derivation	90
6.3.2	Verification across configurations	91
6.3.3	Polydisperse particle compaction in a rigid container: penalty sweep	93
6.4	Scope and caveats	97
6.5	Conclusions	97
7	Elasto-plastic hybrid modelling of granular media	98
7.1	Uniaxial compaction of an elasto-plastic polydisperse granular column	98
7.2	Problem definition and numerical setup	99
7.3	Determination of elastic- and elasto-plastic material parameters	99
7.4	Micro-slip statistics	100
7.5	Loading–unloading response for COF = 0.8	103
7.6	Influence of particle friction on macroscopic behaviour	104
7.6.1	Stress and strain distribution in the hybrid domain	106
7.7	Conclusions	116
8	Silo Discharge of a polydisperse granular system	117
8.1	Introduction	117
8.2	Problem setup and domain decomposition	118
8.2.1	Geometry and modelling context	118
8.2.2	Particle generation, stress–free initialisation, and settling protocol	120
8.2.3	Hybrid overlap topology and weighting	120
8.2.4	Sampling, coarse graining, and wall statistics	122
8.2.5	Sections and reported fields	122
8.2.6	Error measures and presentation	122
8.3	Settled-state silo comparison	123
8.3.1	Comparison of stresses between pure FEM and pure DEM models	123
8.3.2	Comparison of stresses between Hybrid and pure DEM models	124
8.3.3	Comparison of wall stresses between Pure DEM, Pure FEM, and Hybrid models	126
8.4	Silo discharge analysis	130
8.4.1	Analysis along lateral cross sections	130
8.4.2	Analysis along vertical cross sections	136
8.4.3	Analysis of wall stresses	142

8.5	Conclusions	144
9	Conclusions	145
9.1	Conclusions	145
9.2	Limitations	147
9.3	Challenges faced	147
9.4	Future work	148
9.5	Closing statement	148
Appendices		
A	Derivation of volume coupling terms equation.	150
B	Algorithms for implementation	152
C	Core utility functions	157
D	Sample co-simulation input scripts	165
D.1	CoSimulationParameters.json	165
D.2	ProjectParametersDEM.json	170
D.3	ProjectParametersFEM.json	172
	Bibliography	178

Figures and Tables

Figures

1.1	DEM simulation showing funnel-flow pattern in a discharging silo—downward core flow surrounded by quasi-static particles adjacent to the bottom corner walls. Image adapted from [1]	2
2.1	Illustration showing DEM–continuum blending via partition-of-unity weights. In the hybrid region the discrete weight w varies from 0 to 1, while the continuum weight $w_c = 1 - w$; thus $w + w_c = 1$ everywhere.	9
2.2	Schematic variation of the weighting function $\alpha(\mathbf{x})$ across the hybrid overlap region Ω_o , shown for three common choices: constant, linear, and cubic. The adjoining subdomains are the continuum region Ω_c and the discrete (DEM) region Ω_d . Adapted from [2].	9
2.3	Domain of the hybrid coupled cubical system with one element layer as coupling domain. The weight w assigned to each particle is shown in colour. Adapted from [3].	12
2.4	Illustrations of displacement projection between discrete and continuum models.	17
2.5	Kinematic-training workflow for adaptive hybridisation. (a) The simulation begins with a continuum model over the full domain. (b) A shear band is detected by the oracle. (c) The corresponding reference discrete packing is identified. (d) The packing is kinematically trained by imposing the continuum motion on its boundary. (e) The trained DEM patch replaces the continuum region and the simulation proceeds. Figure is adapted from [4]	20
2.6	Domain of the hybrid coupled system with DEM particles in the middle, coarse grained DEM particles in the outer region coupled to the continuum. Adapted from [5]	23
2.7	Coarse-grained density field of 4^3 particles coupled with the cubic finite element using different CG widths. Adapted from [6]	23
3.1	Schematic of the rolling-resistance torque represented as a normal force acting at an effective offset e from the sphere centre.	33

3.2	The schematic of the discrete continuum coupling domain.	42
3.3	Hybrid DEM–FEM analogy with linear weight w . A specimen split along the dashed diagonal (DEM→FEM) shows that under compression the stress/strain fields coincide across all regions—weighting partitions energy only. The interface element ds transmits equal-and-opposite, unweighted coupling traction $\mathbf{t} = \boldsymbol{\sigma}\mathbf{n}$, giving $\mathbf{f}_{\text{cpl}} = \mathbf{t}ds$	49
3.4	Overview of the algorithm used in DEM-FEM volume coupling	51
4.1	Kratos Multiphysics solver structure for DEM–FEM volume coupling	59
5.1	Reference simulation domain and the equivalent hybrid domain.	66
5.2	Vertical stress (σ_{yy}) profile in the hybrid model showing variation of stress profile for different penalty ε (Pa/m^2) distribution. A weighting of 0.01 to 0.99 is used across the hybrid domain.	70
5.3	Vertical stress(σ_{yy}) profile in the hybrid model showing variation of stress profile for different weightage distribution. $\varepsilon = 1 \times 10^{10} \text{Pa}/\text{m}^2$ was used.	71
5.4	Vertical strain (ε_{yy}) profile in the hybrid model showing variation of strain profile for different weightage distribution. $\varepsilon = 1 \times 10^{10} \text{Pa}/\text{m}^2$ was used.	72
5.5	Vertical stress (σ_{yy}) and strain (ε_{yy}) profile in the hybrid model. The x-axis shows the vertical distance of the column starting from 0 m at the bottom. A penalty $\varepsilon = 1 \times 10^{10} \text{Pa}/\text{m}^2$ and weighting of 0 to 1 is used.	73
5.6	Reference simulation domain containing polydisperse particles and the equivalent hybrid domain.	76
5.7	The lateral stress ratio (k) of 0.27 is derived from the horizontal and vertical stress data obtained via reference DEM simulations, using the last 60% of the data points from the plot for calculation.	76
5.8	The stress-strain ratio, calculated as 81,233.16 Pa, was determined from reference DEM simulations by analysing the last 60 % of the data points from the plot.	77
5.9	Vertical stress (σ_{yy}) profile in the hybrid model. The x-axis shows the vertical distance of the column starting from 0 m at the bottom. A penalty $\varepsilon = 10^{11} \text{Pa}/\text{m}^2$ and a weighting of 0 to 1.0 are used.	78
5.10	Horizontal stress (σ_{xx})profile in the hybrid model. The x-axis shows the vertical distance of the column starting from 0 m at the bottom. A penalty $\varepsilon = 10^{11} \text{Pa}/\text{m}^2$ and a weighting of 0 to 1.0 are used.	79

5.11	Lateral stress ratio for the hybrid simulation compared with reference DEM simulation results. A penalty $\varepsilon = 10^{11}$ Pa/m ² and a weighting of 0 to 1.0 are used.	80
5.12	Vertical stresses (σ_{yy}) for different penalties. A weighting of 0.01 to 0.99 are used.	82
5.13	Horizontal stresses (σ_{xx}) for different penalties. A weighting of 0.01 to 0.99 are used.	82
5.14	Vertical stresses (σ_{yy}) for different weighting distribution. A penalty of $\varepsilon = 10^9$ Pa/m ² is used.	83
5.15	Vertical stresses (σ_{yy}) for different mesh sizes. A penalty $\varepsilon = 10^9$ Pa/m ² and a weighting of 0.01 to 0.99 are used.	85
6.1	Observation: with a fixed penalty $\varepsilon = 1 \times 10^8$ Pa/m ² and linear weights (0→1) across the overlap, increasing the particle normal stiffness K_{pp}^N progressively increases the mean-stress deviation (PD) between the hybrid and DEM stress profiles. This motivates scaling ε with particle stiffness.	88
6.2	Axes and layout for Case 3 ($d = 0.01$ m, element 0.02^3 m ³ , $n = 8$).	91
6.3	Penalty scaling across Cases 1–5 for target PD = –5% : C clusters near 0.77 when overlap engagement is uniform; deviations reflect small nonuniformities.	92
6.4	Case 3 penalties tuned for target PD = –0.01% : the inferred constant is $C \approx 9.23$, i.e. a high-fidelity saturation regime with $C \sim 10$	93
6.5	Mean stress deviation PD (%) versus the dimensionless penalty C (bottom axis), with the corresponding $\varepsilon = CE/(dL_c)$ shown on the top secondary axis for Case 3. The dashed vertical line at $C=10$ marks the onset of the high-fidelity regime; for $C > 10$ the curve plateaus with PD $\approx 0\%$	94
6.6	Polydisperse system, lower contact stiffness: penalty sweep around the reference $C = 10$ level. Parameters: $K_{pp}^N = K_{pp}^T = 500$ N/m; COF = 2; imposed displacement $\Delta y = 0.0045$ m.	95
6.7	Polydisperse system, higher contact stiffness: penalty sweep around the reference $C = 10$ level. Parameters: $K_{pp}^N = K_{pp}^T = 2000$ N/m; COF = 10; imposed displacement $\Delta y = 0.015$ m.	96

7.1	Schematic of the iterative Bayesian calibration procedure in GrainLearning used to identify the Drucker–Prager plasticity angles (ϕ, χ) . DEM target curves define the prior; batches of Kratos FEM simulations are assimilated via sequential Monte Carlo (SMC) with DPGMM-adapted proposals; parameters are resampled and iterated until the posterior coefficient of variation falls below 2%.	101
7.2	Distribution of contact-level normalised mobilised friction η/COF for the three friction coefficients.	103
7.3	Axial response for COF = 0.8: vertical stress σ_{yy} versus vertical strain ϵ_{yy} during a complete loading–unloading cycle. Comparison between Pure DEM, Pure FEM and Hybrid models.	104
7.4	Lateral response for COF = 0.8: lateral stress σ_{xx} versus vertical stress σ_{yy} over a complete loading–unloading cycle. Comparison between Pure DEM, Pure FEM and Hybrid models.	105
7.5	Hybrid DEM–FEM simulations: vertical stress σ_{yy} versus vertical strain ϵ_{yy} for different coefficients of friction. Comparison between Pure DEM and Hybrid models.	107
7.6	Hybrid DEM–FEM simulations: lateral stress σ_{xx} versus vertical stress σ_{yy} for different coefficients of friction. Comparison between Pure DEM and Hybrid models.	108
7.7	Domain of different zones in the hybrid model which are split into Hybrid zone 1 , Hybrid zone 2, and DEM zone based on the vertical y-coordinates.	111
7.8	Stress distributions in the hybrid model compared to ground truth DEM at three different loading levels Midway loading (strain = 0.021) , Full load (strain = 0.042), and Midway unloading (strain = 0.021) for COF = 0.4.	112
7.9	Stress distributions in the hybrid model compared to ground truth DEM at three different loading levels Midway loading (strain = 0.021) , Full load (strain = 0.042), and Midway unloading (strain = 0.021) for COF = 0.8.	113
7.10	Stress distributions in the hybrid model compared to ground truth DEM at three different loading levels Midway loading (strain = 0.021) , Full load (strain = 0.042), and Midway unloading (strain = 0.021) for COF = 10.	114
7.11	Consolidated micro slip statistics (η/COF) for the three friction cases. Each panel compares the hybrid subzones (Hybrid zone 1, Hybrid zone 2, and the upper DEM layer) against pure DEM at the three stages of the loading cycle.	115

8.1	Particle size distributions: silo vs RVE. Bin width = 1.22×10^{-4} m.	118
8.2	Problem setup and domain decomposition. (a) pure DEM configuration. (b) Hybrid configuration obtained by embedding a finite-element region.	121
8.3	Settled silo comparison: pure FEM vs pure DEM stresses along various vertical cross sections at A1,A2, and A3.	125
8.4	Settled silo comparison: Hybrid vs pure DEM stresses along various vertical cross sections at A1,A2, and A3.	127
8.5	Settled silo: wall normal and wall shear stress comparisons.	129
8.6	Time evolution of vertical stress σ_{yy} vs X at three different lateral cross sections - comparison between Hybrid and pure DEM model.	131
8.7	Time evolution of lateral stress σ_{xx} vs X at three different lateral cross sections - comparison between Hybrid and pure DEM model.	133
8.8	Time evolution of lateral stress ratio k vs X at three different lateral cross sections - comparison between Hybrid and pure DEM model.	135
8.9	Time evolution of vertical stress σ_{yy} vs y at three different vertical cross sections - comparison between Hybrid and pure DEM model.	137
8.10	Time evolution of vertical stress σ_{xx} vs y at three different vertical cross sections - comparison between Hybrid and pure DEM model.	139
8.11	Time evolution of lateral stress ratio k vs y at three different vertical cross sections - comparison between Hybrid and pure DEM model.	141
8.12	Time evolution of wall stress (normal and shear) during discharge - comparison between Hybrid and pure DEM model.	143

Tables

2.1	Penalty parameter (ε) vs. Lagrange multiplier (λ) for kinematic constraint enforcement for DEM–continuum volume coupling. Here $\Delta \mathbf{u}$ is the discrete–continuum displacement mismatch.	11
2.2	Earliest defensible instances by aspect (with concise context/enforcement notes).	19
2.3	Key recent advances (2021–2024) in discrete–continuum <i>volume</i> coupling for granular materials.	25
5.1	DEM specifications for granular column verification test	68

5.2	FEM specifications for granular column verification test	68
5.3	DEM specifications for polydisperse particle compaction test.	75
5.4	FEM specifications for polydisperse particle compaction test.	81
6.1	Configurations used for scaling verification: particle diameter d , hybrid-zone thickness L_c , element size, and particles per element n	91
7.1	DEM specifications for uniaxial compaction of an elasto-plastic polydisperse granular column	102
7.2	FEM and coupling specifications for uniaxial compaction of an elasto-plastic polydisperse granular column	102
7.3	Calibrated Drucker–Prager parameters for each DEM friction case.	102
8.1	DEM specifications for silo particles	119
8.2	FEM and coupling specifications for silo discharge problem	119
8.3	Node-layer weighting scheme across the hybrid FEM–DEM overlap region	120

Chapter 1

Introduction

1.1 General Background

More than half of all global sales are estimated to involve products that rely on granular materials at some point in their production process, ranking granular media as the second most extensively used material in industry, surpassed only by water [7, 8]. Handling and processing these materials is central to a wide range of applications, including food and beverage production, pharmaceuticals, paints, coatings, and adhesives [9]. In addition to these manufacturing sectors, detailed simulation studies of granular behaviour yield enhanced design strategies and improved predictive capabilities for large-scale industrial challenges in fields such as excavation, mineral processing, and underground construction [10], to name a few. Beyond industrial applications, fundamental research into granular dynamics advances our understanding of natural hazards—including earthquakes, avalanches, landslides, and soil erosion—while enabling more effective risk assessment and mitigation [11]. These studies also provide crucial insights into astrophysical and extraterrestrial phenomena, from planetary formation and solar system evolution to lunar regolith dynamics and surface processes on celestial bodies [12].

Efficient and accurate modelling of granular materials is crucial for the optimisation of a wide variety of industrial and geotechnical processes. Numerical methods such as the Discrete Element Method (DEM) [13] offer particle-level accuracy by explicitly resolving individual particle interactions. However, due to the extremely small time-step requirements associated with particle stiffness and collision dynamics, DEM becomes computationally expensive and often impractical for industrial-scale applications involving large particle assemblies.

With the advent of GPU-accelerated Discrete Element Method (DEM) solvers, simulations involving on the order of 10^9 particles can now be executed routinely, frequently on cloud-based high-performance computing platforms. Yet industrial applications continue to demand even larger particle assemblies, and the quest for greater scalability shows no sign of abating. This persistent need highlights the scope for further algorithmic and hardware optimisations that could push the practical limits of DEM still higher. In contrast to DEM, continuum methods like the Finite Element Method (FEM) significantly improve computational efficiency by utilising larger time steps and employing continuum constitutive laws to describe granular materials [14? –17]. Continuum methods are particularly suitable for static or quasi-static regions where interparticle interactions exhibit negligible fluctuations, rendering detailed particle-level resolution unnecessary. However, their accuracy diminishes substantially in dynamic, highly sheared, or rapidly flowing regions because of their inability to adequately capture localised particle-level phenomena.

A promising approach to overcome these limitations is the development of hybrid computational frameworks coupling DEM and FEM. Such coupling leverages the accuracy of DEM in dynamically evolving regions characterised by significant particle interactions and the computational efficiency of FEM in static or slowly deforming areas. Figure 1.1 depicts an illustration of silo discharge with a central funnel-flow core surrounded by quasi-static zones adjacent to the bottom corner walls.




Figure 1.1: DEM simulation showing funnel-flow pattern in a discharging silo—downward core flow surrounded by quasi-static particles adjacent to the bottom corner walls. Image adapted from [1]

By assigning discrete and continuum models based on local granular flow conditions, hybrid frameworks can substantially reduce computational time while maintaining high accuracy [2–4, 6, 18]. Alternative scaling methodologies such as coarse-graining [1] may augment DEM-continuum coupling frameworks, enabling greater computational efficiency while preserving key mechanical behaviours at appropriate length scales.

1.2 Objectives and Scope of Research

This study presents the implementation and comprehensive analysis of a hybrid Discrete Element Method (DEM)–Finite Element Method (FEM) coupling framework, including systematic verification and validation through benchmark problems and parametric studies. The coupling enables seamless transitions between discrete and continuum representations through an overlapping domain and is exemplified by its application to large-scale silo problems. The proposed methodology employs displacement penalisation techniques to enforce mechanical compatibility across overlapping computational domains, drawing on established theoretical foundations [2, 19]. Implemented within the open-source Kratos Multiphysics platform, the framework is designed to significantly reduce computational costs compared to pure DEM simulations while maintaining high fidelity in numerical accuracy.

The research involved extensive analysis of the hybrid model's response, validated against ground-truth DEM simulations. A verification problem was designed to enforce one-dimensional force transfer in the axial direction, achieved through cubic packing of spherical particles within a rigid, frictionless container. This setup provides an exact analytical solution, allowing rigorous evaluation of the penalty parameter and weighting function's influence on computed stresses and strains.

Through systematic sweeps of penalty stiffness, weighting profiles, inter-particle friction and mesh-to-particle-size ratios in mono- and polydisperse compaction benchmarks, the stress and strain fields within the overlap are mapped, conditions leading to spurious spikes or localisation are identified, and quantitative guidelines for rational parameter selection are established. Crucially, this work establishes and validates a practical scaling framework for the penalty parameter, revealing its governing dependence on particle stiffness (or the Young's modulus of the continuum), particle

diameter, and the characteristic length of the hybrid zone. This scaling consolidates results across disparate test cases, clarifies how to choose penalties that are neither under- nor over-constraining, and provides a transferable recipe for tuning hybrid DEM–FEM couplings.

Finally, the framework is assessed on a gravity-driven silo-discharge benchmark. Using the same Drucker–Prager parameters as in the verification studies, discharge is simulated with the hybrid DEM–FEM model and, for reference, with a high-fidelity pure DEM under identical geometry, wall friction, material properties, and orifice size. The comparison focuses on the evolution of (i) vertical and horizontal Cauchy stresses and (ii) the lateral stress ratio reported as functions of height and time during discharge, with profiles extracted at selected elevations along the width of the silo.

1.3 Thesis Structure

The thesis is organised to systematically develop, implement, and validate a hybrid DEM–FEM coupling framework for granular–continuum interaction problems.

Chapter 2 reviews the literature on discrete element modelling of granular materials, continuum approaches, and hybrid DEM–FEM coupling strategies. Key limitations of existing methods are identified, particularly with respect to discretisation and parameter selection, thereby motivating the present work.

Chapter 3 presents the theoretical formulation of the proposed framework, including the governing equations for the discrete and continuum domains, the coupling strategy, and the definition of penalty-based interaction terms.

Chapter 4 describes the implementation of the methodology within the open-source *Kratos Multiphysics* framework, outlining the coupling algorithms and structure.

Chapter 5 presents verification studies used to assess the accuracy, consistency, and stability of the coupled framework through a set of benchmark problems.

Chapter 6 derives dimensionless scaling laws for the penalty parameters governing the DEM–FEM interaction and demonstrates their role in improving robustness across different discretisations.

Chapter 7 extends the framework to elasto–plastic continuum behaviour and analyses the difference between discrete particles modelled with DEM and an elasto–plastic finite element domain.

Chapter 8 applies the proposed methodology to a silo-discharge case study, demonstrating its performance in an engineering-relevant granular flow problem.

Chapter 9 concludes the thesis by summarizing the main contributions, discussing limitations, and outlining directions for future research.

Chapter 2

Literature review

2.1 Coupling Frameworks for Discrete and Continuum Models

Hybrid discrete–continuum coupling methods aim to combine the strengths of discrete and continuum descriptions by allowing both representations to coexist within a single computational framework. In contrast to sharp interface approaches, overlapping-domain techniques introduce a transition region in which the discrete and continuum models are simultaneously active. Within this hybrid zone, the contribution of each model is blended in a controlled manner to ensure mechanical consistency while avoiding artificial reflections or spurious force transmission at a prescribed interface.

The seamless integration of discrete and continuum models is a central theme in multiscale computational mechanics. Various coupling paradigms have emerged, with Domain Decomposition Methods (DDMs) being particularly prominent. DDMs conceptually divide a problem domain into subdomains, each solved by a potentially different numerical method, with information exchanged across their interfaces. These methods can broadly be categorised into non-overlapping [20, 21] and overlapping approaches [3, 22].

Coupling discrete and continuum models serves multiple purposes—chief among them, accurately capturing interactions between materials with fundamentally different mechanical descriptions. A common scenario involves particles exerting pressure on a wall of another material; here, a surface coupling approach is often adopted, with the particles represented by the discrete element method (DEM) and the structural wall by a continuum scheme such as the finite element method (FEM) or the material point method (MPM) [21, 23, 24]. Kinematic constraints imposed along the interacting surfaces generate the required coupling forces. Surface couplings are therefore well suited to heterogeneous systems.

When the same material is represented on both sides of the interface, however, surface coupling can introduce spurious wave reflections, owing to the mismatch between discrete and continuum descriptions of identical physics. In addition, enforcing precise kinematic constraints in directions tangential to the interface becomes challenging.

For such homogeneous cases, a volume coupling strategy provides a more robust and computationally efficient alternative. By gradually blending the discrete and continuum domains — typically via smoothly varying weighting functions—volume coupling mitigates artificial wave reflections [25, 26]. The use of overlapping domains is particularly advantageous because it establishes the transition over a finite region rather than enforcing coupling at a single surface, overlap-based approaches mitigate stress discontinuities and improve numerical robustness, especially in problems involving large deformations, granular flow, or evolving contact networks. Under quasi-static loading, the progressive transition enabled by weighting promotes a consistent stress–strain distribution across the overlap region. In the present work, we demonstrate the critical role of selecting appropriate weights within the coupling zone and quantify their influence on the accuracy of stress and strain transfer.

Many overlap-based coupling strategies are formulated within a variational or virtual-work framework, in which both the discrete and continuum subdomains are expressed using compatible work statements. This common variational setting enables a direct comparison and combination of contributions from particle forces, continuum stresses, inertia, and external loading, forming the basis for consistent hybridisation of the governing equations. Volume coupling is most commonly formulated within the Arlequin framework. A wealth of prior studies has exploited this strategy for multiscale modelling [19, 22, 27–29], superposing dissimilar physical models—e.g. molecular dynamics, finite-element discretisations of various resolutions, finite-difference schemes—inside an overlap zone where the different descriptions coexist. In that zone, a partition-of-unity representation expresses the global solution as weighted contributions from each sub-model, providing a seamless hand-over (see Fig. 2.1). The formulation is energy-based: a single functional aggregates the energetic terms of all sub-domains, while interfacial continuity (displacement, strain, etc.) is enforced either with Lagrange multipliers [2, 18] or with penalty terms [3, 6].

A key ingredient in such formulations is a spatially varying weighting function that regulates the relative contribution of the discrete and continuum descriptions within the hybrid region. The weight typically varies smoothly across the overlap, ensuring a gradual transfer of stiffness, inertia, and internal forces from one model to the

other. Enforcing that the weights sum to unity at every point guarantees that neither energy nor mass is artificially created or lost during the coupling process. Within the hybrid overlap region, the weight function $w(x)$ may be prescribed with arbitrary spatial variation (e.g., constant, linear, or cubic). As demonstrated by Bauman et al. [2], the use of a constant weight introduces abrupt changes at the overlap boundaries, which manifest as discontinuities in the coupled solution. In contrast, linear and cubic ramp functions provide a gradual transition between the DEM and FEM descriptions, resulting in smooth and well-behaved primary fields. While both ramped profiles perform comparably in terms of displacement continuity, the cubic weighting offers an additional advantage in that its slope vanishes at the overlap boundaries. This higher-order smoothness ensures not only continuity of the primary variable (e.g., displacement), but also improved continuity of its secondary variables, such as strain and stress, which are spatial derivatives of the displacement field.

It should be noted, however, that although the cubic profile theoretically provides smoother and more continuous secondary variables, the effectiveness with which this smooth transition is realised in practice strongly depends on the discretisation of both the finite and discrete element domains. In particular, the mesh resolution in the FEM region and the particle density and contact resolution in the DEM region govern how accurately the spatial variation of the weighting function can be represented and transferred across the overlap. The influence of discretisation on the quality of the coupling and on the capture of the weighting function profile is examined in detail in Chapter 5. . An illustrative profile across the overlap is shown in Fig. 2.2. Outside the overlap, each model evolves independently, which allows efficient resolution transitions (for example, from solid-like to fluid-like behaviour) within one computation.

In the penalty formulation, the displacement mismatch between the discrete and continuum fields is penalised through an additional coupling body force that is proportional to the mismatch, with the strength of enforcement governed by a penalty parameter. This approach is straightforward to implement and integrates naturally into force-based coupling frameworks; however, exact constraint satisfaction is achieved only in the limit of an infinitely large penalty parameter, which may lead to numerical ill-conditioning. In contrast, the Lagrange multiplier formulation enforces the kinematic constraint explicitly by introducing additional unknowns into the weak form, resulting in a constrained saddle-point problem. While this approach allows the constraint to be satisfied to solver

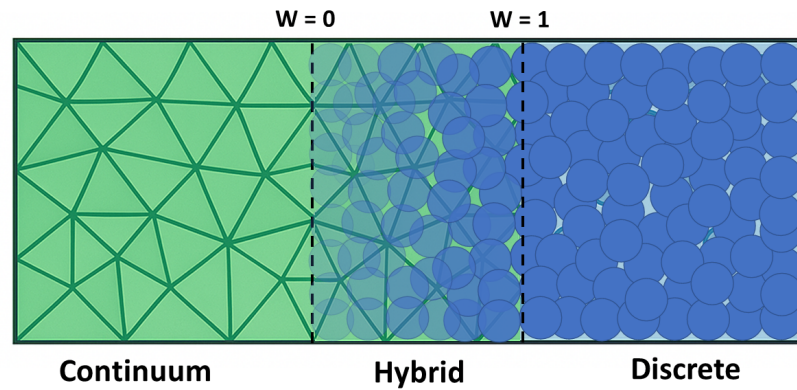


Figure 2.1: Illustration showing DEM–continuum blending via partition-of-unity weights. In the hybrid region the discrete weight w varies from 0 to 1, while the continuum weight $w_c = 1 - w$; thus $w + w_c = 1$ everywhere.




Figure 2.2: Schematic variation of the weighting function $\alpha(\mathbf{x})$ across the hybrid overlap region Ω_o , shown for three common choices: constant, linear, and cubic. The adjoining subdomains are the continuum region Ω_c and the discrete (DEM) region Ω_d . Adapted from [2].

tolerance without introducing systematic bias, it requires careful numerical treatment to ensure numerical stability and robust convergence. The differences between the penalty and Lagrange multiplier formulations for volume coupling are summarised in Table 2.1.

The concepts outlined above provide the theoretical background for the coupling strategy adopted in this work. In Chapter 3, these ideas are specialised to a DEM–FEM framework with overlapping domains, where both submodels are expressed in a common virtual-work form, weighted through a smooth spatial function, and coupled via a penalty-based displacement constraint. The detailed governing equations, implementation aspects, and numerical considerations are presented therein.

Among continuum techniques, Smoothed Particle Hydrodynamics (SPH) and the Material Point Method (MPM) each have particular strengths when coupled to particle systems. SPH interpolates field variables with a kernel over neighbouring particles and excels at fluid-style deformations [30, 31], albeit with intricate boundary treatment. MPM, by contrast, tracks material points on a background grid, avoiding mesh tangling and simplifying large strains [32]. For quasi-static granular flow, however, the Finite Element Method (FEM) remains the most extensively validated continuum option. Combining these merits leads naturally to a DEM–FEM hybrid: FEM supplies the stable quasi-static response, while DEM captures non-quasi-static shear and collision phenomena, furnishing a balanced description of systems that experience both static and dynamic particle behaviour.

Pioneering this direction, [3] extended the Arlequin framework to a DEM–FEM hybrid using a penalty-based coupling with spatial weighting in an overlap domain, and demonstrated the approach on triaxial compression of assemblies of ellipsoidal particles. This appears to be the first granular application featuring an overlap region with weighting functions. The granular assembly considered is shown in Fig. 2.3. At the mesoscopic (representative-volume) scale, they examined principal stress responses, and later applied the scheme to pile installation to illustrate deformation and velocity fields of the hybrid formulation, albeit without quantitative ground-truth validation against full-DEM simulations or experiments. [6] then introduced coarse-graining to convert discrete particle data into differentiable continuum fields before mapping them to the FEM, markedly improving momentum- and energy-conservation benchmarks.

Table 2.1: Penalty parameter (ε) vs. Lagrange multiplier (λ) for kinematic constraint enforcement for DEM–continuum volume coupling. Here $\Delta\mathbf{u}$ is the discrete–continuum displacement mismatch.

Characteristics	Kinematic constraint via penalty enforcement	Kinematic constraint via Lagrange multipliers (LM)
Core mechanism	Applies a coupling <i>body force density</i> proportional to the displacement mismatch: $\mathbf{b}_c = \varepsilon \Delta\mathbf{u}$.	Introduces multipliers to enforce $\Delta\mathbf{u} = 0$ in the weak form.
Constraint satisfaction	Residual gap $\ \Delta\mathbf{u}\ \sim \mathcal{O}(1/\varepsilon)$; vanishes only as $\varepsilon \rightarrow \infty$.	Enforced to solver tolerance; no systematic bias if converged.
Calibration / tuning	Choose ε : too small \Rightarrow visible gap, too large \Rightarrow ill-conditioning.	Scale/tolerances for λ : too large can destabilize Newton iterations, too small slows convergence.
Relation to material stiffness	Typically scale ε with a representative stiffness (e.g., E/h); material-coupled.	Not stiffness-linked; λ has reaction units—scaling is numerical, not constitutive.
Conditioning / robustness	Large ε degrades conditioning; easy to implement, usually robust to start.	Constrained saddle-point system; needs stable discretisation and good preconditioning; sensitive to λ scaling.
Convergence behaviour	Often quick per iteration, may need ε escalation to reduce the coupling gap.	Accurate when converged; under-scaled $\lambda \Rightarrow$ extra iterations, over-scaled $\lambda \Rightarrow$ instability.
Relation to pure DEM stresses (bias)	For finite ε , transmitted stresses are systematically <i>lower</i> than pure DEM.	Converges to pure DEM without one-sided bias; error is \pm solver tolerance.
Convergence to pure DEM	As $\varepsilon \rightarrow \infty$, the solution approaches pure DEM (from below).	Converges to pure DEM within prescribed tolerances when the KKT system is solved.

Chen et al. [4, 33] adopted Lagrange multipliers within a DEM–MPM framework, using a uniform weight of 0.5 in the overlap zone; they further presented an adaptive switch that transforms material-point regions into discrete-element regions (and vice versa), giving the solver added flexibility.




Figure 2.3: Domain of the hybrid coupled cubical system with one element layer as coupling domain. The weight w assigned to each particle is shown in colour. Adapted from [3].

A surge of recent studies applies bridging or Arlequin-style overlap techniques to complex multiphysics problems [26, 34, 35]. Voisin-Leprince et al. [26, 34] blended FEM and DEM in an overlap to reconcile discrete crack growth with the surrounding continuum, capturing highly localised damage under confining pressure. Bai et al. [35] coupled the three-dimensional Particle Flow Code (PFC3D), a discrete element method (DEM) solver, with FLAC3D, a finite-difference-based continuum geomechanics code, to study tunnel–rock interaction under high in-situ stress, demonstrating the adaptability of the method to complex geotechnical scenarios. Chaudry et al. [5] used a two-scale DEM–FEM approach to analyse ship-hull crashworthiness, distributing granular crush material inside the continuum hull. Ruysen and Ben Dhia [36] extended Arlequin to thermo-mechanical selective-laser-melting, embedding high-fidelity sub-models precisely where fine resolution was essential. Collectively these contributions demonstrate that bridging lowers computational cost while enhancing realism whenever localised cracks, thermal gradients or granular flows critically influence global response.

Further evidence of the method's versatility comes from additional domain-decomposition studies. Zhang et al. [37] proposed a superposition-based concurrent scheme for poro-dynamic simulations, merging models of different fidelity to capture both discontinuities and large-scale fluid flow in porous media. Müller et al. [38] developed a damage-aware multiscale framework that preserves microstructural information as regions transition from fine to coarse resolution. Zhang et al. [39] examined wet-type precast beam–column joints for seismic loading, integrating frame and solid elements while maintaining displacement compatibility across the interface.

Bridging ideas have also expanded into the time domain and wave-propagation contexts. Rojek et al. [25] coupled the boundary-element method (BEM) with DEM for three-dimensional wave transmission, employing a staggered temporal algorithm to resolve transients in unbounded media. Similarly, Zhang et al. [40] introduced a spatial–temporal discrete–continuum overlap in which the two sub-problems proceed at distinct time steps, accelerating granular-flow simulations that alternate between slow quasi-static motion and rapid free-flow phases.

2.2 Evolution of Discrete–Continuum Volume Coupling (1998–2020)

The transition of coupling models from purely local numerical schemes to multiscale frameworks has been driven by the need to balance computational efficiency with the high-fidelity resolution of localised phenomena. Central to this evolution is the development of volume-coupling methods, which allow for a seamless transition between discrete (particle-based) and continuum (mesh-based) descriptions. The evolution of volume coupling is characterised by a transition from simple kinematic constraints to sophisticated variational frameworks that partition energy and introduce versatile ways to synchronize multiple degrees of freedom, including displacement, velocity, and rotational fields.

2.2.1 Foundations and Consolidation (1998–2005)

Ben Dhia introduced the Arlequin idea as a unified way to superpose mechanical models in an overlapping zone and to partition the total potential energy with complementary weight functions while enforcing weak kinematic compatibility via Lagrange multipliers or augmented Lagrangians [22]. The central tenet of this framework is the definition of a total virtual work δW for the coupled system ($\Omega_{FE} \cup \Omega_{DE}$) with an overlap hybrid zone $\Omega^H = \Omega_{FE} \cap \Omega_{DE}$. The weighted virtual work is typically expressed as:

$$\delta W = \delta W^{FE}(1 - w) + \delta W^{DE}(w) + \delta W_C \quad (2.1)$$

where $1 - w$ and w are weight functions satisfying $(1 - w) + w = 1$ in the overlap, and δW_C represents the virtual work of the coupling term.

Early engineering studies demonstrated how Arlequin could act as a flexible global–local tool for mesh enrichment and defect insertion [29], formalised its use as an engineering design framework with guidance on suppressing spurious reflections [19], and extended it to dynamic impact problems on thin structures [41]. These studies proved that by choosing a smooth ramp for w , one could significantly reduce the numerical "glitch" or reflection of waves at the interface of different mesh resolutions.

2.2.2 From Theory to Particles: Atomistic/Particle–Continuum Coupling (2007–2010)

Theoretical developments clarified well-posed choices of weights and coupling operators and analysed stability of overlapping formulations [42]. In parallel, Arlequin was generalised to couple particle descriptions to continua: Bauman *et al.* compared L^2 , H^1 -seminorm and H^1 couplings and provided model problems that quantify consistency and convergence [2]. Specifically, they analysed the enforcement of the displacement constraint between discrete and continuum parts $\mathbf{u}^{DE} - \mathbf{u}^C = 0$ in the overlap using the L^2 scalar product:

$$\int_{\Omega^H} \lambda \cdot (\mathbf{u}^{DE} - \mathbf{u}^C) d\Omega = 0 \quad (2.2)$$

where λ is the Lagrange multiplier field. Their findings suggested that H^1 coupling, which penalises gradients, provides superior stability for high-gradient problems like crack tips.

These particle/continuum variants revealed artefacts familiar from atomistic-to-continuum methods, notably ghost forces. In granular systems, these forces are spurious numerical imbalances that emerge because the discrete force laws are inherently truncated at the coupling boundary. For a particle α governed by linear spring-dashpot dynamics, the total contact force is:

$$\mathbf{F}_\alpha^{contact} = \sum_{\beta} \mathbf{f}_{\alpha\beta} \quad (2.3)$$

where $\mathbf{f}_{\alpha\beta}$ is the sum of the normal and tangential forces. Ghost forces arise because a particle at the edge of the discrete domain "misses" potential neighbours β that reside just beyond the coupling interface. While the continuum model represents the bulk material properties in that space, it does not provide the individual contact forces required to satisfy the discrete equilibrium of boundary particles. Consequently, even in a state of uniform compression where particle overlap δ_n should be constant across all contacts, the boundary particles experience a net non-zero "ghost force" due to this missing coordination.

This artifact extends to rotational dynamics as well. The rolling-friction torque \mathbf{T}_r , which depends on the normal force F_n and the coefficient μ_r , becomes unbalanced at the interface. Without a discrete neighbor to provide the reciprocal torque, the particle may undergo unphysical self-rotation. To mitigate these effects, [43] proposed post-processing dead forces and blending-zone design rules to ensure traction jumps and energy bias are minimised. In practice, this involves choosing a weight w that varies from 0 to 1 ensuring minimal effects at boundaries and treating the coupling force as an external force \mathbf{f}_α^c in the numerical integration, mitigating the effects of the ghost force.

Closely related overlapping concepts developed in the bridging domain community provided interchangeable kinematic/energetic coupling operators and weightings that later informed granular couplings [44, 45]. The "Bridging Domain" method (BDM) effectively uses a Lagrange multiplier to constrain the FE nodes and particles to move in a synchronised manner, effectively "averaging" the discrete fluctuations of V_α and ω_α against the smooth continuum displacement field. This prevents high-frequency force oscillations, stemming from the sudden making and breaking of contacts in the overlap, from destabilizing the global FE solution.

2.2.3 First Overlap DEM–FEM for Granular Media and Mapping Choices (2007–2012)

For granular media, a key checkpoint in volume coupling is the coupled DEM–FEM formulation of Rojek and Oñate [46]. In their overlap, the finite element displacement field is projected to particle locations to prescribe particle kinematics, while weak compatibility is enforced via Lagrange multipliers or a penalty. The interaction force \mathbf{F}_p on a particle p within the overlap is derived from the penalty term and effectively acts as the external body force F_α^b in the numerical integration:

$$\mathbf{F}_p = F_\alpha^b = k_{penalty}(\mathbf{u}_{FE}(\mathbf{x}_p) - \mathbf{u}_p) \quad (2.4)$$


where the continuum displacement at the particle location is evaluated as $\mathbf{u}_{FE}(\mathbf{x}_p) = \psi(\mathbf{x}_p) \mathbf{u}^{FE}$, with ψ representing the FE shape-function vector and \mathbf{u}^{FE} the nodal displacements. This FEM→DEM mapping became a template for many subsequent granular hybrids, particularly for enforcing boundary conditions on dense assemblies. However, the formulation did not employ weighting functions, lacked strict energetic consistency, and did not ensure a smooth transition between the discrete and continuum descriptions.

A second milestone is the two-scale overlap model of Wellmann and Wriggers [3], where DEM responses are homogenised in a representative volume to calibrate an equivalent continuum law. In this framework, the overlap enforces weighted energetic consistency with a penalty; field transfer includes projecting particle displacements to FEM nodes using the shape functions. This "Discrete-to-Continuum" (D2C) approach defines the homogenised nodal displacement \mathbf{u}^{DE} via a weighted aggregation of surrounding particles:

$$\mathbf{u}_I^{DE} = \frac{\sum_p \omega_p \psi_I(\mathbf{x}_p) \mathbf{u}_p}{\sum_p \omega_p \psi_I(\mathbf{x}_p)} \quad (2.5)$$

where ω_p is a particle-specific weight which is typically equated to the particle mass m_p to ensure a momentum-preserving transfer. The difference between the two projection methods: discrete to continuum and continuum to discrete is illustrated in Figs. 2.4a and 2.4b.

In practice, two families of projection operators emerged inside the overlap: (i) direct evaluation of FE shape functions ψ at particle positions (as above), and (ii) nodal aggregation (weighted, lumped, or coarse grained deposits) that satisfy partition of unity and in general, assign higher weights to particles closer to a node. Both enable



(a) Discrete-to-continuum projection (D2C). Δu denotes the displacement mismatch to be penalised, u^{DE} the homogenised particle displacements at the nodes, and u^{FE} the nodal displacement values.

(b) Continuum-to-discrete projection (C2D). Δu denotes the displacement mismatch to be penalised, u_p the particle displacements, and u^{FE} the continuum displacement field projected at the particle location.

Figure 2.4: Illustrations of displacement projection between discrete and continuum models.

reciprocal particle \leftrightarrow node transfers used for kinematics and coarse-grained fields. In this thesis, the projection of discrete displacements onto FEM nodes is used; while node to particle projection helps in capturing discrete details of the ground truth, particle to node projections helps in generating a more uniform coupling force field in the hybrid zone that may reduce stress spikes in the hybrid zone. Determining which strategy most accurately reproduces ground-truth DEM remains an open question and motivates future comparative studies for granular media.

2.2.4 Stabilisation, Large Deformations, and Operator Design (2010–2012)

Beyond ghost-force cures [43], penalty-only formulations were analysed and implemented for geometrically nonlinear problems, providing guidance on penalty magnitudes, operator choices, and finite element realisations [47]. A critical finding in this era was the "modelling error" associated with the choice of coupling operator; for instance, using a purely local penalty can lead to spurious oscillations if the particle spacing is not much smaller than the element size.

Nonlocal–local couplings compatible with Arlequin/bridging ideas (for example, non-local elasticity blended with classical continua) refined weight construction and operator consistency under strong gradients and long-range interactions [48]. This was particularly important for granular materials where the "influence radius" of a particle (its coordination number) dictates the material response, requiring a coupling that respects the inherent length scale of the discrete phase.

2.2.5 Fracture, Damage, and Advanced Hybrids (2011–2020)

During the 2010s, overlap couplings matured into robust tools for multiscale fracture and discontinuities. Open-source implementations (for example, PERMIX) showcased concurrent MD–XFEM examples using overlap energy partitioning and weak compatibility to drive transient crack growth [49]. Arlequin-style blends appeared in diverse settings: coupling of nonlocal/peridynamic regions to classical continua or XFEM to confine high-fidelity physics to small subdomains while preserving global efficiency and accuracy [50].

Summary of evolution of DEM–FEM volume coupling (1998–2022)

By 2022, granular volume-coupled hybrids typically featured: (i) a band-type overlap, (ii) displacement projection from FE nodes to DEM particles to prescribe particle motion or vice versa (with reciprocal coarse-graining for continuum fields), (iii) partition-of-unity weighting of the superposed energies/governing equations, and (iv) weak coupling via penalty or Lagrange multipliers.

Representative developments [3, 35, 46] exemplify this evolution, grounded in Arlequin /bridging principles [2, 22, 44]. The key advancements are summarised in Table 2.2.

2.2.6 Recent Advances (2021–2024)

Research over the past five years has focused on adaptive, conservative, and application-driven extensions of discrete–continuum *volume* coupling for granular media. The main directions are: (i) on-the-fly model switching to concentrate DEM resolution where needed; (ii) systematic choices of projection operators, specifically the comparison between direct shape-function ψ evaluation at particle positions and coarse-grained nodal aggregation; (iii) constraint strategies involving both velocity- and displacement-

Table 2.2: Earliest defensible instances by aspect (with concise context/enforcement notes).

Aspect	First paper	Setting & enforcement notes
Overlap + weighting (Arlequin)	Ben Dhia; 1998 [22]	FE–FE (non-granular); energetic weighting; LMs (weak kinematic constraint).
Granular + overlap + node→particle projection	Rojek & Oñate; 2007 [46]	DEM–FEM; no weighting; LMs/penalty in overlap.
Overlap + weighting + particle→node projection	Bauman et al.; 2008 [2]	Atomistic–continuum; Arlequin; LMs; L^2/H^1 couplings.
Overlap + weighting + node→particle projection	Xiao & Belytschko; 2004 [44]	MD–FE; bridging-domain; LMs/augmented Lagrangian.
Granular + overlap + weighting + particle→node projection	Wellmann & Wriggers; 2012 [3]	DEM–FEM; energetic blending; penalty/variational; homogenised transfer.
Granular + overlap + weighting + node→particle projection	Bai et al.; 2022 [35]	DEM–FDM; contribution functions α, β ; LMs at particle DOFs.

LM = Lagrange multiplier; MD = molecular dynamics.

based formulations, as well as the balance between strong pointwise and weak integral enforcement; and (iv) the rigorous conservation of linear and angular momentum alongside total energy in transient dynamics. The key advances are discussed in detail below.

The demand for computational efficiency in large-scale simulations has led to the development of adaptive strategies, such as the approach by Chen et al. (2021) [4]. This study implements *on-the-fly* switching between DEM and continuum, concentrating DEM where the response is highly non-linear while retaining a continuum model elsewhere. A compact workflow of the kinematic training used to replace a continuum patch with DEM is shown in Fig. 2.5.

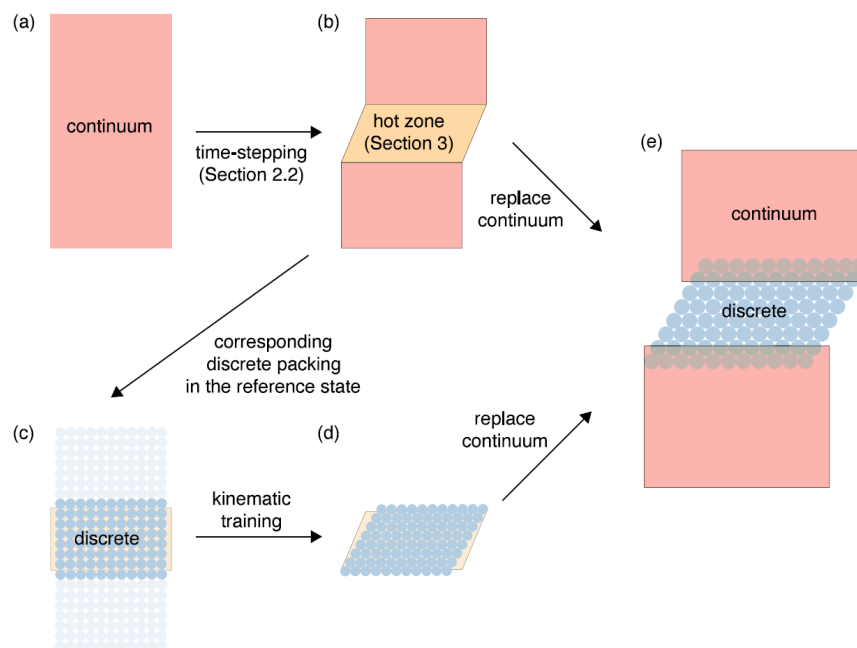


Figure 2.5: Kinematic-training workflow for adaptive hybridisation. (a) The simulation begins with a continuum model over the full domain. (b) A shear band is detected by the oracle. (c) The corresponding reference discrete packing is identified. (d) The packing is kinematically trained by imposing the continuum motion on its boundary. (e) The trained DEM patch replaces the continuum region and the simulation proceeds. Figure is adapted from [4]

The overlap employs a constant weight $w = 0.5$, and kinematic compatibility is enforced via Lagrange multipliers on *velocities*. Velocity constraints are advantageous for explicit time integration (reduced drift relative to displacement constraints in strongly dynamic phases) and help limit transient locking during resolution exchange. Although constant weighting simplifies implementation, it does not yield a spatially smooth energy

partition; accuracy therefore hinges on the chosen overlap width and stabilisation of the multiplier system. A key outcome is that activation/deactivation of DEM can be performed stably when the switching criterion and overlap update are synchronised with the time step. The oracle then operates in the following concise form:

$$s(\mathbf{x}, t) = \|\dot{\epsilon}_p(\mathbf{x}, t)\|, \quad \mathcal{O}(\mathbf{x}, t) = \frac{|\nabla^2 s(\mathbf{x}, t)|}{\max_{\mathbf{y} \in B_r(\mathbf{x})} s(\mathbf{y}, t) + \varepsilon},$$

$$\text{model}(\mathbf{x}, t) = \begin{cases} \text{DEM}, & \mathcal{O}(\mathbf{x}, t) \geq \theta_{\text{on}} \\ \text{Continuum}, & \mathcal{O}(\mathbf{x}, t) \leq \theta_{\text{off}}, \quad \text{with } \theta_{\text{on}} > \theta_{\text{off}}. \\ \text{no change}, & \text{otherwise} \end{cases}$$

where: s is the plastic strain-rate magnitude; ∇^2 is the spatial Laplacian (curvature) operator; $B_r(\mathbf{x})$ is a neighborhood of radius r for local normalisation; $\varepsilon > 0$ avoids division by zero; θ_{on} and θ_{off} are activation/deactivation thresholds providing hysteresis ($\theta_{\text{on}} > \theta_{\text{off}}$).

A notable application of these principles is found in the study of high-stress tunnel–rock interactions, where Bai et al. (2022) [35] implemented a DEM–FDM volume-coupling model. This work targets tunnel–rock interaction at high *in situ* stress using a DEM–FDM pairing and enforces both translational *and* rotational compatibility across the overlap. Continuum grid values are mapped to particles via shape functions (grid \rightarrow particles), and Lagrange multipliers enforce displacement and an effective rotational measure (e.g., vorticity) at particle DOFs. Capturing grain rotations is crucial for shear-band localization and anisotropic dilatancy; explicitly constraining rotations markedly improves the transmission of couple stresses and reduces artificial slip along the interface. Practical considerations include (i) selecting overlap width relative to grain size to avoid over-constraining rotations, and (ii) tuning the LM solver for the mixed translational/rotational block to maintain conditioning at high confinement. Although the efficacy of coupling particle rotations to continuum vorticity remains uncertain.

The continuum nodal *displacements* and *vorticity* are projected to particle locations in the overlap, and the DEM particle fields are constrained to match them. Concretely, the node \rightarrow particle constraints read

$$\Delta \mathbf{u}_p = \mathbf{u}_p - \sum_{i \in N_p} k_{pi} \mathbf{u}_i^{\text{FE}} = \mathbf{0}$$

$$\Delta\omega_p = \omega_p - \sum_{i \in N_p} h_{pi} \omega_i^{\text{FE}} = \mathbf{0}$$

where \mathbf{u}_p and ω_p are the DEM particle displacement and spin; \mathbf{u}_i^{FE} are continuum nodal displacements; ω_i^{FE} denotes the continuum vorticity contribution associated with node i (evaluated at the particle location via standard shape-function gradients); k_{pi} and h_{pi} are the coupling (projection) matrices from nodes to particle p . For formulations implemented in rate form, the continuum rotational measure ω^{FE} is computed using the curl of the velocity field, $\frac{1}{2}(\nabla \times \mathbf{v}^{\text{FE}})$, ensuring the constraint is dimensionally and kinematically consistent with the particle angular velocities. This approach allows the coupling to directly synchronize the rate of rotation between the discrete and continuum descriptions during dynamic events.

A practical application of these frameworks in industrial safety is provided by the study of ship-hull crashworthiness involving crushable granular cores, as investigated by Chaudry et al. (2022) [5]. Here, a coarse-grained DEM (CG-DEM) is coupled to a continuum for ship-hull crashworthiness with a crushable granular core, see Fig. 2.6. Particles project to the continuum field (particles \rightarrow continuum) using shape-function averaging, and compatibility is weakly enforced via a *penalty* on displacement. The coarse-graining reduces DOF counts dramatically, enabling large-scale parametric runs. However, the authors note that penalty-based enforcement requires careful calibration. An insufficiently small ε under-enforces the interface and leads to unphysical slippage, while an excessively large ε induces artificial stiffness. Because the effective coupling stiffness is a direct function of the penalty parameter ε , an over-penalization may significantly reduce the stable time step required for numerical integration. The authors' findings point to application-specific penalty tuning and to the utility of CG-DEM when material calibration (crushability, rate effects) is dominant over micromechanical fidelity.

A significant refinement in the spatial mapping of these hybrid systems is the development of a coarse-grained (CG) enriched volume-coupling scheme by Cheng et al. (2023) [6]. An Arlequin-style scheme that coarse-grains DEM particle velocities to a continuous field and penalises the macroscopic velocity/incremental-displacement mismatch in the overlap. This field coarse-graining is distinct from DEM “coarse-grained particles” (super-particles): the full particle assembly and contacts are retained; only kinematic fields are smoothed with a kernel of width c . In a beam wave-propagation test with a one-third-volume overlap and smoothly ramped weights, DEM micro-parameters are calibrated to match the continuum; CG homogenisation also removes the costly global–local mapping of the classical projection. Relative to $c = 0$, CG better enforces




Figure 2.6: Domain of the hybrid coupled system with DEM particles in the middle, coarse grained DEM particles in the outer region coupled to the continuum. Adapted from [5]

kinematics, tracks the full-DEM reference more closely, and reduces reflections and numerical dissipation—especially for high-frequency content—while strictly conserving mass and keeping total momentum nearly constant. A moderate width performs best (optimal $c \approx 1.5 \Delta X$); larger c over-smooths, and $c \rightarrow 0$ recovers the standard projection. See Fig. 2.7 for particle density fields with $c/d = 1$ and $c/d = 2$.

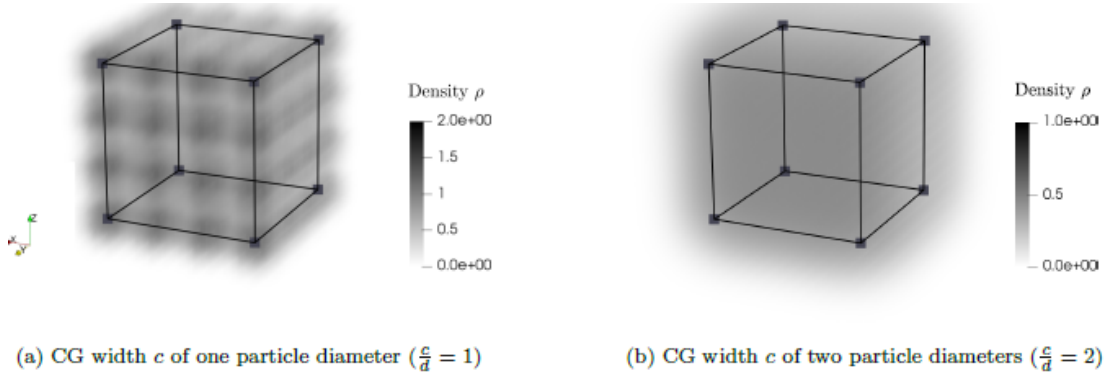


Figure 2.7: Coarse-grained density field of 4^3 particles coupled with the cubic finite element using different CG widths. Adapted from [6]

The fidelity of dynamic signal propagation is further explored by Voisin et al. [26]. This study probes how enforcement strength affects compressional and shear wave transmission across the interface. With shape-function projection from grid nodes to particles (nodes \rightarrow particles), two alternatives are contrasted. *Strong* constraints enforce a pointwise match,

$$(S) \quad u_p - \psi(x_p) u^{FE} = \mathbf{0} \quad \forall p \in \Omega_H,$$

whereas *weak* coupling enforces it in an averaged/penalised sense,

$$(W) \quad \int_{\Omega_c} w(x) [u_p(x) - u_{FE}(x)] dV = \mathbf{0}$$

Here, u^{FE} are FE nodal unknowns, $u_{FE}(x) = \psi(x)u^{FE}$ is the FE field, u_p are particle DOFs at locations x_p , $\psi(\cdot)$ are FE shape functions, and $w(x)$ is a coupling weight. Strong enforcement provides excellent high-frequency transmission and minimal phase change error but increases stiffness and time-step sensitivity; weak enforcement is cheaper and more robust but filters high frequencies. This leads to the practical recommendation of employing strong constraints for high-fidelity transients, such as impact and wave diagnostics, where the preservation of high-frequency signal content is paramount. Conversely, weak coupling is preferred for quasi-static or slowly varying processes, as it offers enhanced numerical robustness and allows for larger stable time steps by mitigating the individually varying coupling forces associated with pointwise enforcement.

The research landscape between 2021 and 2024 represents a shift from foundational proof-of-concept couplings toward physically consistent, application-specific frameworks. These advances can be broadly categorised into three methodological pillars.

First, the move toward adaptive hybridisation allows for the dynamic transition between discrete and continuum descriptions based on local state variables, significantly optimizing computational resources. Second, there is an increased emphasis on rotational and multi-field consistency, where the coupling of displacement is augmented by vorticity or angular velocity to capture the inherent microrotations in granular shear bands. Finally, the development of Coarse-Graining (CG) based volume coupling has addressed traditional projection errors. By smoothing particle kinematics through a kernel before mapping them to the continuum, CG techniques reduce stress spikes and artificial signal reflections that typically occur at the discrete-continuum interface. To provide a comparative overview of these strategies, the projection operators (mapping direction) and kinematic constraint enforcement methods (Lagrange multipliers vs. penalty) for each milestone work are cross-referenced in Table 2.3.

Table 2.3: Key recent advances (2021–2024) in discrete–continuum *volume* coupling for granular materials.

Work & scope	Projection of displacement / velocity (source → target)	Kinematic constraint (enforced quantity)
[4] Adaptive on-the-fly switching between DEM and continuum; overlap uses constant weight $w=0.5$	Shape-function averaging: particles → background grid	Lagrange multipliers enforce <i>velocity</i> continuity
[35] Tunnel–rock interaction at high in-situ stress with DEM–FDM coupling	Shape functions: grid nodes → particles	Lagrange multipliers enforce <i>displacement</i> and <i>rotation</i> (via vorticity)
[5] Ship-hull crashworthiness with crushable granular core; coarse-grained DEM	Shape functions: particles → continuum field	Penalty method on <i>displacement</i> continuity across the interface
[6] Momentum- and energy-conserving concurrent coupling	Coarse-grained DEM → FEM nodes via shape functions	Projection-based constraint on the <i>displacement</i> field
[26] Transmission of pressure/shear waves across the FEM/DEM interface; weak vs. strong coupling compared	Shape functions: grid nodes → particles	Lagrange multipliers on <i>displacement</i> ; weak (integral) vs. strong imposition

2.3 Research gaps

Although discrete–continuum coupling techniques for granular materials have advanced markedly in recent years, existing work has concentrated almost exclusively on bulk-scale responses. The local behaviour inside the discrete–continuum overlap region—crucial for the fidelity of DEM–FEM couplings—remains largely unexplored. As a result, the quantitative influence of key coupling parameters, including the penalty coefficient, weighting functions, inter particle friction, and mesh resolution, is still not well characterised.

A further gap concerns how the penalty parameter should be set in DEM–FEM couplings. Existing studies typically choose the penalty by ad-hoc tuning for each case, with little consensus—or evidence-based guidance—on how it ought to scale with problem characteristics (e.g., geometry, particle size distribution, mesh resolution, hybrid-zone thickness). As a result, reported settings are difficult to transfer across configurations, and the sensitivity of predictions to the chosen penalty remains poorly quantified. There is a clear need for systematic studies that establish defensible selection rules, assess robustness across regimes, and provide reproducible defaults practitioners can adopt without case-by-case trial and error.

In addition, most large-scale industrial simulations are assessed only through domain-level metrics, without systematic cross-validation against high-resolution discrete models. This practice weakens confidence in their predictive capability. To address these limitations, the present study performs a large-scale silo-discharge simulation and benchmarks the hybrid solution against a reference pure-DEM model. The comparison establishes rigorous verification procedures, thereby enhancing the robustness of DEM–FEM analyses for industrial applications.

2.4 Contributions of the Present Work

This thesis advances the state of discrete–continuum coupling for granular media in five tightly linked directions:

- i) **Verification under quasi-static loading.** The solver is rigorously verified against analytical solutions and high-resolution reference DEM calculations for linear-elastic granular assemblies subjected to one-dimensional force transfer. This exercise establishes baseline accuracy for stress transmission, displacement compatibility, and energy consistency in the overlap zone.

ii) **Comprehensive hybrid-zone analysis.** A systematic parametric sweep quantifies the influence of (i) penalty stiffness, (ii) weighting-function shape, and (iii) mesh-to-particle size ratio on overlap-zone performance. Two canonical benchmarks are employed:

- a monodisperse granular column under uniaxial compression;
- a polydisperse compaction test.

Resulting maps of stress error, strain localisation, and convergence behaviour provide practical guidelines for robust parameter selection. These insights, unavailable in earlier granular-coupling studies, equip practitioners to deploy hybrid DEM–FEM models with confidence and efficiency.

iii) **Penalty-parameter scaling law.** A compact dimensionally consistent rule that ties the coupling penalty to material stiffness and hybrid-zone geometry is derived and validated,

$$\varepsilon = C \cdot \frac{E}{dL_c},$$

where the single, dimensionless constant C cleanly gathers the algorithmic details of mapping and weighting. Recasting our data through this relation collapses results from all test families—mono/polydisperse, varied mesh–particle ratios, and size of the elements—onto narrow trends, revealing a near invariance of C across configurations. The outcome is a practical recipe for selecting ε *a priori* (by choosing C) together with a natural way to report uncertainty. For the poly-disperse systems, setting $C = 10$ delivers the closest hybrid response attainable relative to the DEM reference, with negligible benefit from further increasing ε .

iv) **Elasto-plastic benchmarking and industrial validation.** The framework is extended to an *elasto-plastic* FEM region governed by a Drucker–Prager law. First, the same polydisperse configuration is used to build DP benchmarks by *varying the (interparticle/continuum) coefficient of friction* and assessing sensitivity of yield onset, lateral stress ratio, and stress transfer across the overlap. Next, a realistic silo-discharge problem containing $\mathcal{O}(10^5)$ particles is simulated: the discrete bulk is coupled to the elasto-plastic continuum, and hybrid predictions of vertical stress, lateral stress ratio, and particle-scale kinematics are benchmarked against a high-resolution pure-DEM reference.

- v) **Open-source implementation.** The coupling has been realised through the *CoSimulation* module of *Kratos Multiphysics*. Because it interfaces directly with the existing *DEM* and *Structural Mechanics* applications—both of which already offer a rich palette of contact and constitutive laws, and allow user-defined models to be plugged in seamlessly—the present implementation provides a flexible foundation onto which additional physics (e.g. thermal, fluid, or chemo-mechanical effects) can be incorporated in future work.

Collectively, these contributions deliver both a transferable software asset and new physical insight into the mechanics of DEM–FEM overlap regions, paving the way for extensions to elasto-plastic constitutive laws, adaptive model switching, and coupled multi-physics applications.

2.5 Summary

Hybrid DEM–continuum modelling has matured into a versatile toolkit for problems that mix particle-scale mechanics with bulk continuum behaviour. The prevailing strategy is domain decomposition where, discrete and continuum subdomains are advanced concurrently and coupled either across a sharp interface (surface coupling) or within an overlap (volume coupling). Arlequin-style volume coupling has proven especially effective for granular media because smoothly varying weights mitigate spurious reflections, promote compatible stress–strain fields under quasi-static loading, and enable resolution transitions without intrusive remeshing. Within this family, two enforcement mechanisms dominate: Lagrange multipliers, which are energetically consistent but numerically more involved (saddle-point systems), and penalty methods, which are simple and robust to implement but require careful parameter selection. Mapping choices (shape-function projections, lumped/partition-of-unity nodal deposits, or coarse-grained fields) and the design of overlap weights materially influence stability, conservation, and accuracy.

Recent advances have broadened both scope and reliability. Coarse-graining pipelines that convert particle data into differentiable continuum fields improve momentum- and energy-consistent transfer from DEM to FEM. Dedicated interface studies on the transmission of pressure and shear waves compare weak (integral/penalised) and strong (pointwise) enforcement; strong constraints better preserve high-frequency content but increase stiffness and time-step sensitivity, whereas weak enforcement

is cheaper and more forgiving while naturally filtering high frequencies. Adaptive and concurrent schemes now switch resolution *on the fly*, and overlap formulations extend to elasto-plastic solids, thermo-mechanical processes, and transient wave propagation, with demonstrations at engineering scale (e.g. tunnels, crashworthiness, additive manufacturing).

Despite this progress, important gaps remain. The local mechanics of the overlap zone—how engagement varies layerwise, how weights and mapping interact with mesh-to-particle ratios, and how these choices redistribute stress—has not been characterised with the same rigour as bulk responses. Reported penalty choices are typically case specific and tuned *post hoc*; there is little consensus on how the penalty should scale with geometry, particle size distribution, or overlap thickness, which hampers transferability and reproducibility. Moreover, many industrial studies still rely on domain-level metrics without systematic cross-validation against high-resolution DEM references.

This thesis advances discrete–continuum volume coupling through four primary contributions. First, it establishes a verification baseline for DEM–FEM hybrids under quasi-static loading against analytical and high-resolution DEM references. Second, it systematically maps the influence of overlap parameters—including penalty stiffness, weighting functions, and spatial scales—using mono- and polydisperse benchmarks.

Third, the work introduces a dimensionally consistent guideline for selecting the penalty parameter based on material stiffness and hybrid-zone geometry. By utilising a dimensionless constant to capture algorithmic mapping and weighting, this relation enables *a priori* settings and collapses disparate datasets into a unified framework for transparent uncertainty reporting. Finally, the methodology is implemented within *Kratos Multiphysics*, extending the scope to elasto-plastic FEM regions governed by Drucker–Prager plasticity, which are validated against industrial-scale silo-discharge problems.

Together, these developments provide a robust methodological foundation and actionable defaults for DEM–continuum coupling, enhancing reproducibility and facilitating large-scale engineering analyses.

Chapter 3

Formulation of DEM, FEM, and coupled interaction laws

3.1 Introduction

This chapter presents the governing formulations used throughout the coupling framework developed in this work. It includes the detailed balance laws and variational formulations for both the Finite Element Method (FEM) and the Discrete Element Method (DEM), along with the constitutive models employed to represent material behaviour.

While the general formulations are comprehensive, the specific choice of constitutive law may vary depending on the problem under consideration. For instance, different verification cases or the silo discharge scenario may use distinct material models (e.g., linear elasticity, Drucker–Prager plasticity). All such models used in this thesis are outlined in this section for completeness.

For clarity, the particular constitutive model adopted in each numerical example will be explicitly identified in the corresponding summary tables that precede the presentation of each problem.

3.2 Discrete Element Method - linear spring-dashpot contact model

In this study, all the simulations are performed with linear spring-dashpot contact model. Under this model, when two granular particles are in contact Fig.??, the normal force F_n can be formulated using Eq.(3.1) and tangential force F_t using Eq.(3.2). The force acts only when the distance r between the centres of two particles of radii R_i and R_j is

less than their contact distance $d' = R_i + R_j$. The governing equations in this model is as shown below:

$$F_n = K_n \delta_n + F_n^{\text{damp}} \quad (3.1)$$

$$F_t = \min \left(|K_t \delta_t + F_t^{\text{damp}}|, \mu_f |F_n| \right) \quad (3.2)$$

Here, the normal overlap is $\delta_n = d' - r$ and δ_t is the tangential displacement. The normal force has two terms, a linearly varying force with spring coefficient (K_n) and a damping force with damping coefficient (γ). The tangential force also has two terms: a linearly varying shear force with spring coefficient (K_t) and a damping force with damping coefficient (γ). μ_f is the coefficient of Coulomb friction while subscripts n and t denote the normal and tangential components respectively.

The damping parameters for both the normal and tangential directions are derived from the coefficient of restitution following the methodology outlined by Schwager and Pöschel [51]. In this approach, the coefficient of restitution (COR), which quantifies the kinetic energy loss during collisions, is analytically related to the dimensionless damping ratio γ . This ratio is then used to compute the damping forces in both the normal and tangential directions as follows:

$$F_n^{\text{damp}} = -2\gamma\sqrt{K_n m_{\text{red}}} v_n^{\text{rel}}, \quad F_t^{\text{damp}} = -2\gamma\sqrt{K_t m_{\text{red}}} v_t^{\text{rel}} \quad (3.3)$$

where m_{red} is the reduced mass of the two contacting particles, and v_n^{rel} , v_t^{rel} are the relative velocities in the normal and tangential directions, respectively. The stiffness coefficients K_n and K_t are specified directly as constants for each contact pair.

The relationship between the damping ratio γ and COR depends on the damping regime, and is defined in piecewise fashion as:

$$\text{COR}(\gamma) = \begin{cases} \exp\left(-\frac{\gamma}{\sqrt{1-\gamma^2}}\left(\pi - \tan^{-1}\left(\frac{2\gamma\sqrt{1-\gamma^2}}{1-2\gamma^2}\right)\right)\right), & \text{for } \gamma \leq \frac{1}{\sqrt{2}} \\ \exp\left(-\frac{\gamma}{\sqrt{1-\gamma^2}}\tan^{-1}\left(\frac{2\gamma\sqrt{1-\gamma^2}}{2\gamma^2-1}\right)\right), & \text{for } \frac{1}{\sqrt{2}} < \gamma < 1 \\ \exp(-2), & \text{for } \gamma = 1 \\ \exp\left(-\frac{\gamma}{\sqrt{\gamma^2-1}}\ln\left(\frac{\gamma/\sqrt{\gamma^2-1}+1}{\gamma/\sqrt{\gamma^2-1}-1}\right)\right), & \text{for } \gamma > 1 \end{cases} \quad (3.4)$$

In this study, it is assumed that the same value of γ applies to both the normal and tangential directions. When the user specifies only the coefficient of restitution e , the corresponding γ is computed numerically (e.g., via a bisection root-finding method) using the above expressions. This is consistent with the implementation used in Kratos Multiphysics [51] and ensures physically meaningful damping in all contact scenarios.

The tangential force is also subject to the Coulomb friction constraint, ensuring that:

$$|F_t| \leq \mu_f |F_n|, \quad (3.5)$$

This condition limits the tangential force to prevent unphysical buildup of sliding forces, thereby accurately modelling sticking-sliding transitions in granular contacts.

3.2.1 Constant directional torque (CDT) model

The motion of a particle interacting with another particle or wall can be segmented into sliding and rolling phases. During the sliding phase, tangential forces are determined using Eq. (3.2). In the rolling phase, a *constant directional torque* (CDT) model is employed to account for rolling friction.

The implemented CDT model [52] computes an opposing rolling-friction torque proportional to the normal contact force and directed opposite to the relative angular velocity. Specifically, the rolling-friction torque \mathbf{T}_r is calculated as

$$\mathbf{T}_r = -\mu_r F_n \ell_{\text{eff}} \frac{\boldsymbol{\omega}_{\text{rel}}}{\|\boldsymbol{\omega}_{\text{rel}}\|} \quad (3.6)$$




Figure 3.1: Schematic of the rolling-resistance torque represented as a normal force acting at an effective offset e from the sphere centre.

where μ_r is the rolling-friction coefficient defined for the particle–particle or particle–wall interaction, F_n is the magnitude of the normal contact force, $\ell_{\text{eff}} = R - \delta_{\text{eff}}$ is the effective arm length, and ω_{rel} is the relative angular-velocity vector between contacting entities. Equivalently, the schematic offset in Fig. 3.1 satisfies

$$\|\mathbf{T}_r\| = F_n e \quad \Rightarrow \quad e = \mu_r \ell_{\text{eff}} \quad (3.7)$$

For particle–particle contacts, the indentation-adjusted term is

$$\delta_{\text{eff}} = \frac{E_{\text{other}}}{E_{\text{self}} + E_{\text{other}}} \delta \quad (3.8)$$

where E_{self} and E_{other} are the respective Young’s moduli and δ is the indentation depth. For particle–wall contacts this simplifies to

$$\delta_{\text{eff}} = \delta \quad (3.9)$$

The rolling-friction torque always opposes rolling, dissipating rotational energy and slowing the particle’s spin. The Kratos implementation computes the torque using normalised relative angular velocities to ensure correct directionality, and accumulates rolling-dissipation energy each time step, with separate bookkeeping for particle–particle and particle–wall interactions.

Mechanically, this coupling between normal compression and rotational resistance prevents the excessive 'rolling-ball' effect common in idealised spherical assemblies. By penalizing relative angular motion more heavily in high-pressure zones, the model effectively captures the increased interlocking and rotational hindrance observed in real-world granular media under stress. Consequently, the CDT model is essential for accurately predicting shear strength and energy dissipation in regions where high normal stresses would otherwise lead to non-physical particle spinning. Even with small rolling-friction coefficients μ_r (typically 10^{-2} – 10^{-1}), it can noticeably influence the macroscopic response of dense assemblies. This sensitivity arises because the resistive torque \mathbf{T}_r is directly proportional to the normal contact force F_n . In scenarios involving high confinement or dense packing, the magnitude of F_n increases significantly, thereby amplifying the rolling resistance even if μ_r remains low.

3.2.2 Numerical integration of forces

When a particle α of mass m and moment of inertia I comes in contact with another particle β , they exert forces on each other. These forces $F_{\alpha\beta}^t$, the resulting moments $M_{\alpha\beta}^t$ and the external body forces F_{α}^b such as gravity and coupling forces, are calculated at each time-step. The resulting linear and angular accelerations are multiplied by the time step dt to obtain the new velocity V_{α}^{t+1} and angular velocity ω_{α}^{t+1} , as shown in Eqs. 3.11–3.12. The equations below illustrate simple procedural integration to update displacement and velocity:

$$\mathbf{x}_{\alpha}^{t+1} = \mathbf{x}_{\alpha}^t + \mathbf{v}_{\alpha}^t dt \quad (3.10)$$

$$\mathbf{v}_{\alpha}^{t+1} = \mathbf{v}_{\alpha}^t + \frac{dt}{m_{\alpha}} \left[\sum_{\beta \neq \alpha} \mathbf{f}_{\alpha\beta}^t + \mathbf{f}_{\alpha}^b \right] \quad (3.11)$$

$$\omega_{\alpha}^{t+1} = \omega_{\alpha}^t + \frac{dt}{I_{\alpha}} \sum_{\beta \neq \alpha} \mathbf{m}_{\alpha\beta}^t \quad (3.12)$$

3.3 Finite Element Method (Continuum)

This section outlines the continuum mechanics formulation used in the finite element (FE) representation of the solid domain. Starting from the strong form of the balance of linear momentum, the weak form is derived and subsequently discretised in space to obtain the semi-discrete equations of motion. The formulation is general and accommodates various constitutive laws, which are introduced in subsequent subsections. This provides the foundation for the continuum part of the coupled DEM–FEM simulations [53–56].

3.3.1 Weak form of the Cauchy momentum balance

A deformable solid is defined over a region $\Omega_{\text{FE}} \subset \mathbb{R}^3$ with boundary $\partial\Omega_{\text{FE}}$, split into two nonoverlapping parts: Γ_u , where displacements are prescribed, and Γ_t , where surface tractions (forces per unit area) are prescribed, so that $\partial\Omega_{\text{FE}} = \Gamma_u \cup \Gamma_t$ and $\Gamma_u \cap \Gamma_t = \emptyset$. The unknown field is the displacement $\mathbf{X}(\mathbf{x}, t)$; material density is $\rho(\mathbf{x})$; body forces per unit mass (e.g. gravity) are $\mathbf{b}(\mathbf{x}, t)$; and \mathbf{n} denotes the outward unit normal on the boundary. The boundary data are

$$\mathbf{X} = \bar{\mathbf{X}} \text{ on } \Gamma_u, \quad \boldsymbol{\sigma} \mathbf{n} = \mathbf{t} \text{ on } \Gamma_t,$$

where $\boldsymbol{\sigma}$ is the (symmetric) Cauchy stress tensor, $\bar{\mathbf{X}}$ is the prescribed displacement, and \mathbf{t} the prescribed traction vector.

The motion is governed pointwise by the local balance of linear momentum: the divergence of the stress plus body forces equals mass density times acceleration,

$$\nabla \cdot \boldsymbol{\sigma} + \rho \mathbf{b} = \rho \ddot{\mathbf{X}} \quad \text{in } \Omega_{\text{FE}}. \quad (3.13)$$

Equation (3.13) together with the boundary conditions is called the *strong form*, because it must hold at every point and requires relatively smooth (differentiable) fields.

To construct a formulation well suited to finite elements, the strong form is converted into an integral statement by multiplying (3.13) by an arbitrary *virtual displacement* $\delta \mathbf{X}$ that vanishes on Γ_u and integrating over Ω_{FE} :

$$\int_{\Omega_{\text{FE}}} \delta \mathbf{X} \cdot (\nabla \cdot \boldsymbol{\sigma}) dV + \int_{\Omega_{\text{FE}}} \rho \mathbf{b} \cdot \delta \mathbf{X} dV - \int_{\Omega_{\text{FE}}} \rho \ddot{\mathbf{X}} \cdot \delta \mathbf{X} dV = 0. \quad (3.14)$$

This is the principle of virtual work: for all admissible (kinematically consistent) variations, the work of internal stresses balances the work of external actions (body forces and inertia) [53, 56].

The first integral in (3.14) contains the divergence of the stress. Applying the divergence theorem (the three-dimensional version of integration by parts) replaces that volume term by a boundary term and an integral involving the gradient of $\delta\mathbf{X}$:

$$\int_{\Omega_{\text{FE}}} \delta\mathbf{X} \cdot (\nabla \cdot \boldsymbol{\sigma}) dV = \int_{\Gamma_{\mathbf{t}}} \delta\mathbf{X} \cdot \boldsymbol{\sigma} \mathbf{n} dA - \int_{\Omega_{\text{FE}}} \boldsymbol{\sigma} : \delta\boldsymbol{\varepsilon} dV,$$

where the virtual strain tensor is $\delta\boldsymbol{\varepsilon} = \frac{1}{2}(\nabla\delta\mathbf{X} + \nabla\delta\mathbf{X}^\top)$. On $\Gamma_{\mathbf{t}}$ the traction condition $\boldsymbol{\sigma} \mathbf{n} = \mathbf{t}$ is introduced, and on Γ_u the variation $\delta\mathbf{X}$ vanishes by construction.

Collecting terms leads to the *weak (variational) form* of the momentum balance: find \mathbf{X} such that, for all virtual displacements $\delta\mathbf{X}$ that are square-integrable together with their first derivatives and vanish on Γ_u ,

$$\int_{\Omega_{\text{FE}}} \boldsymbol{\sigma} : \delta\boldsymbol{\varepsilon} dV - \int_{\Omega_{\text{FE}}} \rho(\mathbf{b} - \ddot{\mathbf{X}}) \cdot \delta\mathbf{X} dV - \int_{\Gamma_{\mathbf{t}}} \mathbf{t} \cdot \delta\mathbf{X} dA = 0, \quad \forall \delta\mathbf{X} \in \mathcal{V}_0, \quad (3.15)$$

with the test space

$$\mathcal{V}_0 := \{ \delta\mathbf{X} \in [H^1(\Omega_{\text{FE}})]^3 : \delta\mathbf{X} = \mathbf{0} \text{ on } \Gamma_u \}.$$

Here H^1 denotes the set of square-integrable functions whose first derivatives are also square-integrable. This functional setting matches the continuity provided by standard (C^0) finite element shape functions, and the weak form (3.15) is equivalent to the strong form under mild regularity assumptions while being directly amenable to finite element discretisation [53, 54, 56].

3.3.2 Semi-discrete equation of motion

To convert the variational principle (3.15) into algebraic equations suitable for computation, the displacement field is interpolated using a set of basis functions organised in the matrix $\Psi(\mathbf{x})$ and time-dependent nodal values collected in the vector $\mathbf{d}(t)$. This gives

$$\mathbf{X}_h(\mathbf{x}, t) = \Psi(\mathbf{x}) \mathbf{d}(t), \quad \delta\mathbf{X} = \Psi(\mathbf{x}) \delta\mathbf{d},$$

so that the corresponding strain fields follow from the symmetric gradient as

$$\boldsymbol{\varepsilon}_h = \mathbf{B}\mathbf{d}, \quad \delta\boldsymbol{\varepsilon} = \mathbf{B}\delta\mathbf{d},$$

where $\mathbf{B}(\mathbf{x})$ is the strain–displacement matrix obtained from $\boldsymbol{\Psi}$. Substituting these interpolations into Eq.(3.15), using the constitutive law through the elasticity (or algorithmic) matrix \mathbf{D} , and collecting all terms that multiply the arbitrary nodal variations $\delta\mathbf{d}$ yields, at the element level,

$$\delta\mathbf{d}^\top (\mathbf{K}_e \mathbf{d} + \mathbf{C}_e \dot{\mathbf{d}} + \mathbf{M}_e \ddot{\mathbf{d}} - \mathbf{f}_e^{\text{ext}}) = 0,$$

with the element stiffness, mass, damping and external-force contributions defined by

$$\begin{aligned} \mathbf{K}_e &= \int_{\Omega_e} \mathbf{B}^\top \mathbf{D} \mathbf{B} dV, & \mathbf{M}_e &= \int_{\Omega_e} \rho \boldsymbol{\Psi}^\top \boldsymbol{\Psi} dV, & \mathbf{C}_e &= \alpha \mathbf{M}_e + \beta \mathbf{K}_e, \\ \mathbf{f}_e^{\text{ext}} &= \int_{\Omega_e} \rho \boldsymbol{\Psi}^\top \mathbf{b} dV + \int_{\Gamma_{t,e}} \boldsymbol{\Psi}^\top \mathbf{t} dA. \end{aligned}$$

Here, \mathbf{K}_e captures how internal stresses oppose deformation (material stiffness and geometry), \mathbf{M}_e encodes inertia (how mass distributes with the chosen interpolation), and \mathbf{C}_e represents viscous-like damping when Rayleigh coefficients (α, β) are used. Because $\delta\mathbf{d}$ is arbitrary, the bracketed factor must vanish, giving the element equilibrium equations

$$\mathbf{K}_e \mathbf{d} + \mathbf{C}_e \dot{\mathbf{d}} + \mathbf{M}_e \ddot{\mathbf{d}} = \mathbf{f}_e^{\text{ext}}.$$

Numerically, this is the core FEM result: once $\boldsymbol{\Psi}$ and \mathbf{B} are fixed by the chosen element type and mesh, all matrices are obtained by standard quadrature and depend only on material properties and the current configuration (for linear kinematics, on geometry and material parameters).

Element equations are assembled by summing contributions from all elements and enforcing essential (Dirichlet) boundary conditions, which couples shared nodal degrees of freedom and yields the global semi-discrete second-order system

$$\mathbf{K}\mathbf{d}(t) + \mathbf{C}\dot{\mathbf{d}}(t) + \mathbf{M}\ddot{\mathbf{d}}(t) = \mathbf{f}_{\text{ext}}(t), \quad (3.16)$$

with $\mathbf{K} = \sum_e \mathbf{K}_e$, $\mathbf{M} = \sum_e \mathbf{M}_e$, $\mathbf{C} = \sum_e \mathbf{C}_e$ and $\mathbf{f}_{\text{ext}} = \sum_e \mathbf{f}_e^{\text{ext}}$ after applying boundary conditions. The matrix \mathbf{M} is symmetric positive definite for standard choices, \mathbf{K} is symmetric positive semidefinite (positive definite after constraints for stable linear elasticity), and \mathbf{C} inherits symmetry from its Rayleigh construction; these properties underpin stability and efficiency of the linear algebra solvers used in practice.

Equation (3.16) is a system of ordinary differential equations in time. To advance the solution, a time-integration method (for example, the Newmark family) recasts (3.16) at each step into a linear algebraic problem of the form

$$\mathbf{K}^{\text{eff}} \Delta \mathbf{d} = \mathbf{r},$$

where \mathbf{K}^{eff} combines \mathbf{K} , \mathbf{C} , and \mathbf{M} with scheme-dependent coefficients and \mathbf{r} collects external forces together with inertia and damping contributions from the previous step. This algebraic form is the basis for all subsequent numerical procedures: preconditioning and solving the linear system, applying contact and coupling forces, and updating kinematic quantities consistently. In the coupled DEM–FEM context used later, weighting or partitioning of DEM-to-FEM interface forces is naturally interpreted at the level of Eq.(3.16) and its time-discrete counterpart, making the numerical implementation of the weighting transparent in terms of contributions to \mathbf{K} , \mathbf{C} , and \mathbf{M} .

3.3.3 Linear Elastic Constitutive Model

The formulation presented thus far requires a constitutive model to relate the strain tensor $\boldsymbol{\varepsilon}$ to the Cauchy stress tensor $\boldsymbol{\sigma}$. In the context of linear elasticity, applicable for small deformations and small strains, the constitutive relation takes the form

$$\boldsymbol{\sigma} = \mathbf{D}\boldsymbol{\varepsilon}, \quad (3.17)$$

where \mathbf{D} denotes the fourth-order elasticity tensor, often represented in matrix form under Voigt notation for computational efficiency.

In this work, the constitutive behaviour of the solid domain Ω_{FE} is governed by an isotropic linear elastic model, as implemented through the `LinearElastic3DLaw` in *Kratos Multiphysics*. For such materials, the elasticity matrix \mathbf{D} is defined by the Lamé parameters λ and μ , which are functions of the Young's modulus E and Poisson's ratio ν :

$$\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}, \quad \mu = \frac{E}{2(1+\nu)}.$$

By mapping the stress and strain tensors into sixth-order vectors under the convention $\boldsymbol{\sigma} = [\sigma_{xx}, \sigma_{yy}, \sigma_{zz}, \tau_{xy}, \tau_{yz}, \tau_{zx}]^\top$ and employing engineering shear strains, the elasticity matrix \mathbf{D} is expressed as:

$$\mathbf{D} = \begin{bmatrix} \lambda + 2\mu & \lambda & \lambda & 0 & 0 & 0 \\ \lambda & \lambda + 2\mu & \lambda & 0 & 0 & 0 \\ \lambda & \lambda & \lambda + 2\mu & 0 & 0 & 0 \\ 0 & 0 & 0 & \mu & 0 & 0 \\ 0 & 0 & 0 & 0 & \mu & 0 \\ 0 & 0 & 0 & 0 & 0 & \mu \end{bmatrix}. \quad (3.18)$$

This vector representation and the corresponding matrix formulation follow the standard treatments for linear elasticity [53, 55].

The constitutive relation (3.17) enters the weak form of the momentum balance through the internal virtual work term. Substituting the stress–strain relation into the weak form (3.15) yields

$$\int_{\Omega_{FE}} (\mathbf{D}\boldsymbol{\varepsilon}) : \delta\boldsymbol{\varepsilon} dV = \int_{\Omega_{FE}} \boldsymbol{\sigma} : \delta\boldsymbol{\varepsilon} dV. \quad (3.19)$$

Upon spatial discretisation using the standard small-strain formulation $\boldsymbol{\varepsilon}_h = \mathbf{B}\mathbf{d}$ and $\delta\boldsymbol{\varepsilon} = \mathbf{B}\delta\mathbf{d}$, the above becomes

$$\delta\mathbf{d}^\top \left(\int_{\Omega_e} \mathbf{B}^\top \mathbf{D} \mathbf{B} dV \right) \mathbf{d} = \delta\mathbf{d}^\top \mathbf{K}_e \mathbf{d}, \quad (3.20)$$

where the element stiffness matrix is explicitly given by

$$\mathbf{K}_e = \int_{\Omega_e} \mathbf{B}^\top \mathbf{D} \mathbf{B} dV. \quad (3.21)$$

The constitutive matrix \mathbf{D} is thus a critical component in the computation of both the internal forces and the stiffness matrix. Within *Kratos Multiphysics*, the material law defined by `LinearElastic3DLaw` provides the required stress and tangent stiffness tensor at each integration point during assembly. Consequently, the selection of the constitutive law directly impacts the internal force vector and the overall dynamic response through the term $\mathbf{K}\mathbf{d}(t)$ in the semi-discrete equation of motion Eq. (3.16).

3.3.4 Drucker–Prager constitutive model (non-associated)

The Drucker–Prager (DP) model provides a smooth, convex approximation to the Mohr–Coulomb criterion commonly used for frictional and cohesive–frictional geomaterials, eliminating corner singularities and enabling robust return-mapping algorithms in finite elements [57, 58]. To set notation, the Cauchy stress is decomposed into volumetric and deviatoric parts,

$$\boldsymbol{\sigma} = \mathbf{s} + p\mathbf{I}, \quad p = \frac{1}{3} \text{tr} \boldsymbol{\sigma}, \quad \mathbf{s} = \boldsymbol{\sigma} - p\mathbf{I},$$

and the invariants used by pressure–sensitive plasticity are

$$I_1 = \text{tr} \boldsymbol{\sigma} = 3p, \quad J_2 = \frac{1}{2} \mathbf{s} : \mathbf{s}.$$

With these, the classical DP yield function is often written

$$f(\boldsymbol{\sigma}, \boldsymbol{\kappa}) = \alpha I_1 + \sqrt{J_2} - k(\boldsymbol{\kappa}) \leq 0, \quad (3.22)$$

where α controls pressure sensitivity and k plays the role of a cohesion-like threshold [58]. In the implementation used here (as in KRATOS), it is convenient to scale the same cone into a uniaxial-equivalent form,

$$F(\boldsymbol{\sigma}, \boldsymbol{\kappa}) = \sigma_{\text{eq}}(\boldsymbol{\sigma}; \phi) - \sigma_y(\boldsymbol{\kappa}) \leq 0, \quad (3.23)$$

with an equivalent stress

$$\sigma_{\text{eq}}(\boldsymbol{\sigma}; \phi) = \frac{2 \sin \phi}{3(1 - \sin \phi)} I_1 + \frac{\sqrt{3}(3 - \sin \phi)}{3(1 - \sin \phi)} \sqrt{J_2}.$$

The two normalisations (3.22)–(3.23) describe the same conical surface up to a positive factor; the latter is convenient because σ_y has units of stress and directly matches uniaxial strength [56]. For calibration against Mohr–Coulomb parameters under the compression–positive convention one may use

$$\alpha = \frac{2 \sin \phi}{\sqrt{3}(3 - \sin \phi)}, \quad k = \frac{6c \cos \phi}{\sqrt{3}(3 - \sin \phi)}, \quad (3.24)$$

and, if the plastic potential g is chosen independently (nonassociated flow), the same mapping with the dilation angle χ to define

$$g(\boldsymbol{\sigma}) = \alpha_\chi I_1 + \sqrt{J_2} - k_\chi, \quad \alpha_\chi = \frac{2 \sin \chi}{\sqrt{3}(3 - \sin \chi)}, \quad k_\chi = \frac{6c \cos \chi}{\sqrt{3}(3 - \sin \chi)}.$$

Using $g \neq f$ decouples dilatancy from the yield surface so that volumetric plastic strain is governed by χ rather than ϕ [56, 58].

Given a strain increment, the constitutive update proceeds by an elastic predictor–plastic corrector (closest-point) return mapping: starting from the trial stress $\boldsymbol{\sigma}^{\text{tr}} = \mathbf{D} : (\boldsymbol{\varepsilon} - \boldsymbol{\varepsilon}^p)$, the algorithm enforces $F(\boldsymbol{\sigma}, \boldsymbol{\kappa}) = 0$ at the end of the step, updates the plastic strain via $\dot{\boldsymbol{\varepsilon}}^p = \dot{\lambda} \partial g / \partial \boldsymbol{\sigma}$ with $\dot{\lambda} \geq 0$, and delivers the algorithmic elasto–plastic tangent required for quadratic convergence of the global Newton iterations. In compact form,

$$\mathbf{D}_{\text{ep}} = \mathbf{D} - \frac{(\mathbf{D} : \partial g / \partial \boldsymbol{\sigma}) \otimes (\mathbf{D} : \partial f / \partial \boldsymbol{\sigma})}{(\partial f / \partial \boldsymbol{\sigma}) : \mathbf{D} : (\partial g / \partial \boldsymbol{\sigma}) + \mathcal{H}}, \quad (3.25)$$

where \mathcal{H} collects any isotropic hardening contribution; for perfect plasticity $\mathcal{H} = 0$ [56, 59]. Substituting the updated stress into the internal virtual work and linearising yields the consistent finite element tangent

$$\mathbf{K}_T = \int_{\Omega_{\text{FE}}} \mathbf{B}^\top \mathbf{D}_{\text{ep}} \mathbf{B} \, dV,$$

which is the quantity assembled at the global level and used by the solver.

In all simulations reported here a perfect-plastic variant is used, so no hardening or softening evolution is introduced and the fracture-energy regularisation mechanisms available in KRATOS are inactive by design; consequently, to maintain solver consistency, specific parameters such as the tensile and compressive yield stresses are set to a nominal value of 10 Pa, which is several orders of magnitude smaller than the operational stress levels encountered in the elasto-plastic analysis. Other specified inputs, including the hardening-curve selector `HARDENING_CURVE=3` (PerfectPlasticity), and the fracture energy $G_f = 5.0 \times 10^7$ N/m are specified but do not influence the response in the present settings. General background on the DP model, its nonassociated extension, and return-mapping algorithms can be found in Simo and Hughes [56], Drucker and Prager [57], Chen and Han [58], and de Souza Neto et al. [59].

3.4 Implementation of DEM-FEM coupling with overlapping domains

To couple DEM and FEM, the two domains are superimposed as shown in Figure 3.2. The hybrid domain separating the DEM and FEM regions is modelled using a combined virtual work equation that weights the energy and mass distribution in the hybrid domain. To weight the virtual work performed in the DEM model δW_{DE} , a smooth and continuous function $w(\mathbf{x})$ of the position of the particles $\mathbf{x} \in \Omega_H$ is used. Ideally, this function increases monotonically from zero on $\partial\Omega_{FE} \cap \Omega_H$ to one on $\partial\Omega_{DE} \cap \Omega_H$. The virtual work performed in the FEM model, δW_{FE} , is weighted by the function $1 - w(\mathbf{X})$, where \mathbf{X} represents the position of interest, typically the integration points within the FEM region of the hybrid domain, $\mathbf{X} \in \Omega_H$. This weighting ensures that the coupling weights on the FEM and DEM sides sum to unity at any given location, ensuring that the total energy and mass do not exceed the reference DEM model.




Figure 3.2: The schematic of the discrete continuum coupling domain.

3.4.1 Coupling approach

Starting from the continuum weak form in Eq. (3.15) and summing the contributions of all finite elements gives the *unweighted* virtual-work relation for the finite-element sub-domain

$$\delta w^{FE} = \sum_{e=1}^{N_e} \int_{\Omega_e^{FE}} \left[\boldsymbol{\sigma} : \delta \boldsymbol{\varepsilon} - (\mathbf{b} - \rho \ddot{\mathbf{X}}) \cdot \delta \mathbf{X} \right] dV + \sum_{r=1}^{N_r} \int_{\Gamma_t^{FE}} \mathbf{t} \cdot \delta \mathbf{X} dA, \quad (3.26)$$

Analogously, observing that mechanical work equals the product of force and displacement, the discrete-element equations of motion (Eq. (3.11)) are multiplied by the virtual particle displacement $\delta \mathbf{x}_\alpha$, yielding the discrete virtual-work statement

$$\delta W^{\text{DE}} = \sum_{\alpha=1}^{N_p} \left[\left(m_\alpha \ddot{\mathbf{x}}_\alpha - \mathbf{f}_\alpha^{\text{b}} \right) - \sum_{\beta=1}^{N_\alpha} \mathbf{f}_{\alpha\beta} \right] \cdot \delta \mathbf{x}_\alpha, \quad (3.27)$$

where N_e and N_r are the total numbers of volume and surface elements in Ω and Γ_t , respectively, and N_p is the total number of discrete particles; \mathbf{t} denotes a surface traction acting on a subset of the body's boundary, $\Gamma_t \subset \partial\Omega_{\text{FE}}$. Expressing both subdomains within this common variational framework facilitates the later introduction of a spatial weighting function that partitions the total virtual work smoothly between FEM and DEM while satisfying the normalisation constraint $w + (1 - w) = 1$ at every material point. This energy partitioning ensures that the coupled DEM–FEM formulation remains strictly energy-consistent and precludes any spurious numerical energy generation.

Multiplying the DEM and FEM governing equations with their respective coupling weights $1 - w$ and w for δW^{FE} and δW^{DE} , gives the weighted sum of the total virtual work $\delta W = \delta W^{\text{FE}} + \delta W^{\text{DE}}$, with

$$\delta W^{\text{FE}} = \sum_{e=1}^{N_e} \int_{\Omega_e^{\text{FE}}} (1 - w) \{ \boldsymbol{\sigma} : \delta \boldsymbol{\varepsilon} - (\mathbf{b} - \rho \ddot{\mathbf{X}}) \cdot \delta \mathbf{X} \} dV + \sum_{r=1}^{N_r} \int_{\Gamma_r^{\text{FE}}} \mathbf{t} \cdot \delta \mathbf{X} dA, \quad (3.28)$$

$$\delta W^{\text{DE}} = \sum_{\alpha=1}^{N_p} \left[w_\alpha (m_\alpha \ddot{\mathbf{x}}_\alpha - \mathbf{f}_\alpha^{\text{b}}) - \sum_{\beta=1}^{N_\alpha} w_{\alpha\beta} \mathbf{f}_{\alpha\beta} \right] \cdot \delta \mathbf{x}_\alpha, \quad (3.29)$$

The shorthand notation $w_\alpha = w(\mathbf{x}_\alpha)$ and $w_{\alpha\beta} = w(\mathbf{x}_{\alpha\beta})$ is used for the weights at the positions of the particles and the contact points. $\delta \mathbf{x}_\alpha$ is a variation of the position of a particle α .

The virtual work of the hybrid model is equal to the weighted sum of the work done by the submodels plus the work done by the penalty term, which can be expressed as: $\delta W = \delta W_{\text{FE}} + \delta W_{\text{DE}} + \delta W_{\text{C}}$, where δW_{C} is expressed as:

$$\delta W_C = \int_{\Omega^H} \varepsilon(\mathbf{u}^{DE} - \mathbf{u}^{FE}) \cdot (\delta \mathbf{x} - \delta \mathbf{X}) dV, \quad (3.30)$$

$$= \int_{\Omega^H} \underbrace{\varepsilon(\mathbf{u}^{DE} - \mathbf{u}^{FE})}_{:=\mathbf{b}^C} \cdot \delta \mathbf{x} dV + \int_{\Omega^H} \varepsilon(\mathbf{u}^{FE} - \mathbf{u}^{DE}) \cdot \delta \mathbf{X} dV, \quad (3.31)$$

A penalty parameter, denoted as ε , is introduced to enforce the displacement continuity between the DEM and FEM regions within the hybrid zone. The contribution of the work done by the penalty term consists of two parts: one from the DEM domain and the other from the FEM domain as shown in Eq. (3.31). These two contributions appear with opposite signs in the DEM and FEM residuals, so that the coupling forces are equal and opposite and do not act as an external load on the combined system. In this way, the penalty term only stores internal energy associated with the displacement mismatch and does not introduce any spurious net energy into the DEM–FEM system. This ensures that the displacement difference between the DEM and FEM regions in the hybrid zone is penalised without altering the overall energy balance of the system. Adding this additional virtual work to Eq. (3.28) and Eq. (3.29) gives the coupled governing equations which are adjusted by the corresponding coupling weights as shown below.

$$\begin{aligned} \delta W = & \int_{\Omega^{FE}} (1-w) \{ \boldsymbol{\sigma} : \delta \boldsymbol{\varepsilon} - (\mathbf{b} - \rho \ddot{\mathbf{X}}) \cdot \delta \mathbf{X} \} dV + \int_{\Gamma_t^{FE}} \mathbf{t} \cdot \delta \mathbf{X} dA + \int_{\Omega^{FE}} \mathbf{b}^C \cdot \delta \mathbf{X} dV \\ & + \sum_{\alpha=1}^{N_p} \left[w_\alpha (m_\alpha \ddot{\mathbf{x}}_\alpha - \mathbf{f}_\alpha^b) \cdot \delta \mathbf{x}_\alpha - \sum_{\beta} w_{\alpha\beta} \mathbf{f}_{\alpha\beta} \cdot \delta \mathbf{x}_\alpha \right] - \sum_{\alpha=1}^{N_p} \mathbf{f}_\alpha^C \cdot \delta \mathbf{x}_\alpha, \end{aligned} \quad (3.32)$$

The coupling force density is calculated by penalising the displacement difference between the homogenised particle displacement (\mathbf{u}^{DE}) and the FEM nodal displacements (\mathbf{u}^{FE}) as shown in Eq. (3.33). This penalty-based coupling force density is subsequently integrated within each element to compute the nodal coupling forces, as expressed in Eq. (3.34). Here, N^E denotes the number of nodes in each FEM element, and ψ represents the shape function associated with each node, evaluated at the respective integration points. The term $\Pi_{i\alpha}$, defined in Eq. (3.35), serves to map the nodal coupling forces to the discrete DEM particles. N_p^E represents the number of DEM particles that are contained strictly within a particular FEM element. This

mapping ensures that the total force acting on the DEM particles and the FEM nodes remains equal in magnitude, but opposite in direction, satisfying Newton's third law and maintaining force balance during the coupling. This results in the DEM coupling force as shown in Eq. (3.36).

$$\mathbf{b}^C = \varepsilon(\mathbf{u}^{\text{DE}} - \mathbf{u}^{\text{FE}}), \quad (3.33)$$

$$\mathbf{f}_i^C = \varepsilon \sum_{i=1}^{N^E} \sum_{j=1}^{N^E} \psi_i \psi_j dV (\mathbf{u}_j^{\text{DE}} - \mathbf{u}_j^{\text{FE}}). \quad (3.34)$$

$$\Pi_{i\alpha} = \frac{m_\alpha \psi_i(\mathbf{x}_\alpha)}{\sum_{\alpha=1}^{N_p^E} m_\alpha \psi_i(\mathbf{x}_\alpha)} \quad (3.35)$$

$$\mathbf{f}_\alpha^C = -\varepsilon \sum_{i=1}^N \sum_{j=1}^N \Pi_{i\alpha} \int_{\Omega^H} \psi_i \psi_j dV (\mathbf{u}_j^{\text{DE}} - \mathbf{u}_j^{\text{FE}}). \quad (3.36)$$

Furthermore, Eq. (3.37) outlines the procedure for mapping the velocities from the DEM particles to the FEM nodes within a given element.

$$\mathbf{v}_i^{\text{DE}} = \frac{\sum_{\alpha=1}^{N_p^E} m_\alpha \mathbf{v}_\alpha \psi_i(\mathbf{x}_\alpha)}{\sum_{\alpha=1}^{N_p^E} m_\alpha \psi_i(\mathbf{x}_\alpha)} \quad (3.37)$$

Here, m_α is the mass of the α -th particle. \mathbf{v}_α represents the velocity of the α -th DEM particle. $\psi_i(\mathbf{x}_\alpha)$ is the shape function of node i , evaluated at the location \mathbf{x}_α of the α -th particle. The term \mathbf{v}_i^{DE} refers to the homogenised velocity at node i of the FEM element, which is derived by averaging the contributions of the DEM particles in that element.

The numerator of Eq. (3.37) sums the weighted contributions of the particle velocities to node i , considering both the particle mass and the shape function. The denominator normalises these contributions by summing the mass-weighted shape functions, ensuring that the resulting velocity is a consistent average across all particles in the FEM element.

This formulation allows the velocities from the particles to be mapped to FEM nodes in a way that reflects the distribution and influence of the particles within the FEM element. Consequently, the total displacement at node i is found as $\mathbf{u}_i^{\text{DE}}(t) = \mathbf{v}_i^{\text{DE}} \cdot \Delta t + \mathbf{u}_i^{\text{DE}}(t - 1)$. This displacement is then used to calculate the DEM and FEM displacement difference as shown in Eq. (3.33). For more detailed information on the coupling methodology, the reader is referred to the work of Cheng et al.[6].

3.4.2 Mapping

The 'MappingApplication' of Kratos allows for mapping any nodal data between (non-matching) meshes or particles. The process of mapping linear momentum from DEM particles to FEM nodes (Eq. (3.37)) is as follows. The shape function $\psi_{\alpha_j}^i$ associated with node i is evaluated at the position of particle α_j . The value of the shape function at the particle's position is the weight by which the momentum is multiplied before assigning it to the i -th node. Here, m_{α_j} and \mathbf{v}_{α_j} are the mass and velocity of particle α_j , and \mathbf{P}_i^{DE} represents the aggregated particle momentum at each node. This momentum is divided by the interpolated particle mass m_i^{DE} stored at each node to get the homogenised velocity. The following elaborates how the aggregated particle momentum is calculated on each FEM node — an operation shown in the numerator of Eq. (3.37).

$$\begin{aligned}
 \mathbf{P}_1^{\text{DE}} &\leftarrow \psi_{\alpha_1}^1 \cdot (m_{\alpha_1} \mathbf{v}_{\alpha_1}) + \psi_{\alpha_2}^1 \cdot (m_{\alpha_2} \mathbf{v}_{\alpha_2}) + \dots \\
 \mathbf{P}_2^{\text{DE}} &\leftarrow \psi_{\alpha_1}^2 \cdot (m_{\alpha_1} \mathbf{v}_{\alpha_1}) + \psi_{\alpha_2}^2 \cdot (m_{\alpha_2} \mathbf{v}_{\alpha_2}) + \dots \\
 \mathbf{P}_3^{\text{DE}} &\leftarrow \psi_{\alpha_1}^3 \cdot (m_{\alpha_1} \mathbf{v}_{\alpha_1}) + \psi_{\alpha_2}^3 \cdot (m_{\alpha_2} \mathbf{v}_{\alpha_2}) + \dots \\
 \mathbf{P}_4^{\text{DE}} &\leftarrow \psi_{\alpha_1}^4 \cdot (m_{\alpha_1} \mathbf{v}_{\alpha_1}) + \psi_{\alpha_2}^4 \cdot (m_{\alpha_2} \mathbf{v}_{\alpha_2}) + \dots \\
 &\vdots
 \end{aligned} \tag{3.38}$$

For the inverse mapping, updated coupling forces are transferred from the FEM nodes to the DEM particles using a conservative scheme. Force equilibrium is strictly maintained by ensuring the resultant of the nodal forces is identically equal to the sum of the forces mapped to the particle locations. This mapping is general in nature, formulated to preserve any conservative quantity during the transfer process and thereby ensuring consistency across the discrete–continuum interface. In the mapping shown below, \mathbf{f}_i^{C} represents the coupling force stored at the i -th node. The force on each particle is computed by evaluating the shape functions at the position of the

particle, similar to the previous particle-to-node mapping. The coupling force on the α -th particle, \mathbf{f}_α^C , is given by the following expressions, where each term on the right-hand side is divided by $m_i^{\text{DE}} = \sum_{\alpha_i} m_{\alpha_i} \psi_{\alpha_i}^i$, the interpolated mass of particles stored at the i -th node, and the entire right-hand side is multiplied by m_{α_i} , the mass of the particle:

$$\begin{aligned} \mathbf{f}_{\alpha_1}^C &\leftarrow m_{\alpha_1} \left(\frac{\psi_{\alpha_1}^1 \mathbf{f}_1^C}{m_1^{\text{DE}}} + \frac{\psi_{\alpha_1}^2 \mathbf{f}_2^C}{m_2^{\text{DE}}} + \frac{\psi_{\alpha_1}^3 \mathbf{f}_3^C}{m_3^{\text{DE}}} + \frac{\psi_{\alpha_1}^4 \mathbf{f}_4^C}{m_4^{\text{DE}}} + \dots \right) \\ \mathbf{f}_{\alpha_2}^C &\leftarrow m_{\alpha_2} \left(\frac{\psi_{\alpha_2}^1 \mathbf{f}_1^C}{m_1^{\text{DE}}} + \frac{\psi_{\alpha_2}^2 \mathbf{f}_2^C}{m_2^{\text{DE}}} + \frac{\psi_{\alpha_2}^3 \mathbf{f}_3^C}{m_3^{\text{DE}}} + \frac{\psi_{\alpha_2}^4 \mathbf{f}_4^C}{m_4^{\text{DE}}} + \dots \right) \\ &\vdots \end{aligned} \quad (3.39)$$

Note that the inverse mapping employs a local, lumped approximation of the projection matrix $\mathbf{\Pi}_{i\alpha}$ (defined in Eq. (3.35)) to facilitate an efficient nodal-to-particle transfer. For further details on the non-local, consistent versions of the projection rule—which involve the inversion of the global mass or coupling matrices—the reader is referred to Cheng et al. [6, Eq. 18] and Wellmann and Wriggers [3, Eq. 35].

3.4.3 Effects of weighting in the hybrid zone

The weighting modifies the FE equations of motion Eq. (3.16) in the hybrid domain by multiplying the stiffness and the mass matrix with the respective weights; the resulting formulation can be written as:

$$(1 - w)\mathbf{K}\mathbf{d}(t) + \mathbf{C}\dot{\mathbf{d}}(t) + (1 - w)\mathbf{M}\ddot{\mathbf{d}}(t) = \mathbf{f}_{\text{ext}}(t), \quad (3.40)$$

with the unaltered Rayleigh damping matrix $\mathbf{C} = \alpha\mathbf{M} + \beta\mathbf{K}$ inheriting the scaling implicitly through \mathbf{M} and \mathbf{K} .

On the DEM side of the hybrid domain, the particle masses and all contact forces are scaled by the local weighting factor. The guiding requirement is *model consistency*: at any spatial location within the hybrid region, the stress and strain computed by the DEM and FEM sub-models must coincide and must also match the solution that would be obtained were the hybrid domain replaced by a single (pure) model. Formally,

$$\boldsymbol{\sigma}^{\text{DEM}}(\mathbf{x}) = \boldsymbol{\sigma}^{\text{FEM}}(\mathbf{x}) = \boldsymbol{\sigma}^{\text{single}}(\mathbf{x}), \quad \boldsymbol{\varepsilon}^{\text{DEM}}(\mathbf{x}) = \boldsymbol{\varepsilon}^{\text{FEM}}(\mathbf{x}) = \boldsymbol{\varepsilon}^{\text{single}}(\mathbf{x}). \quad (3.41)$$

Figure 3.3 illustrates this idea using a physical analogy: a single, simple specimen is split into two parts that are perfectly glued along a dashed diagonal. One may identify the upper portion with the discrete (DEM) description and the lower portion with the continuum (FEM) description; the dashed diagonal represents a linear weighting function $w \in [0, 1]$ varying from one model to the other. Under a vertical compressive load, the resulting stress and strain are identical at every material point—whether that point lies in the discrete part, the continuum part, or anywhere in the hybrid zone. In other words, the weighting should only partition *energy* between the two sub-models; it should not alter the stress or strain fields in the hybrid region. Equivalently, at any point \mathbf{x} within the hybrid region, the description is a mixture of the two sub-models with participation factors $w(\mathbf{x})$ and $1 - w(\mathbf{x})$; numerically the DEM and FEM “diffusely” interpenetrate, even though in the physical analogy the parts are distinctively separated.

This behaviour follows from the structure of the hybrid governing equations: the weighting factor multiplies every internal term (e.g., inertia, constitutive response, DEM contact interactions, effective masses) but *does not* multiply the externally applied loads. In the coupling scheme, we therefore introduce explicit coupling forces as additional external loads, applied to the continuum nodes and to the DEM particle centres. These coupling forces are *unweighted* and, ideally, equal the equilibrium forces that each sub-model would experience from the other (weighted) part of the domain.

To make this precise, consider in Fig. 3.3 an infinitesimal horizontal slice of thickness dH within the hybrid region. Within this slice, the DEM and FEM sub-domains occupy fractions w and $(1 - w)$, respectively, and are perfectly bonded across an interface of area ds . The coupling force is the “glue” traction transmitted across this interface. By equilibrium,

$$\mathbf{t}^{\text{DEM}} ds + \mathbf{t}^{\text{FEM}} ds = \mathbf{0}, \quad \mathbf{f}_{\text{cpl}} = \mathbf{t} ds, \quad (3.42)$$

where \mathbf{t} is the (Cauchy) traction vector acting on the interface and ds is the interfacial measure (a line element in 2D, an area element in 3D). Hence the interface (coupling) forces are equal and opposite and must *not* be weighted.

The DEM discreteness and the use of finite-dimensional shape functions to project and map fields between DEM and FEM introduce small inconsistencies in the coupling forces that vary with position \mathbf{x} in the hybrid region. These discrepancies perturb the otherwise smooth stress and strain profiles in the hybrid zone. The magnitude and character of these effects will be evident in the verification problems presented in Chapter 5.

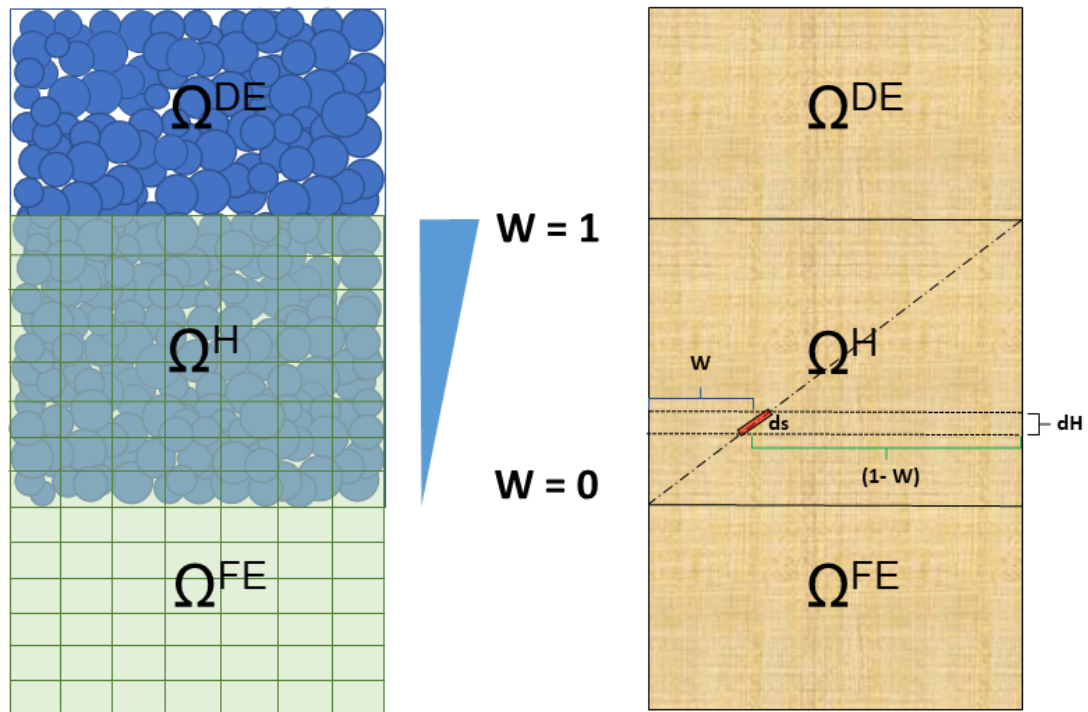


Figure 3.3: Hybrid DEM–FEM analogy with linear weight w . A specimen split along the dashed diagonal (DEM→FEM) shows that under compression the stress/strain fields coincide across all regions—weighting partitions energy only. The interface element ds transmits equal-and-opposite, unweighted coupling traction $\mathbf{t} = \boldsymbol{\sigma} \mathbf{n}$, giving $\mathbf{f}_{\text{cpl}} = \mathbf{t} ds$.

3.4.4 DEM-FEM Interaction Mechanisms

In the hybrid region, the FEM uses custom elements to integrate forces and displacements through interpolated coupling weights at the integration points. These weights are initially assigned to the nodes of the FEM mesh elements and are subsequently interpolated to the Gauss integration points for continuum-side calculations. Conversely, for the DEM sub-domain, the weights are interpolated to the particle centres to

govern discrete-side computations. When considering particle-particle interactions, the individual weights of each particle are used to determine the specific coupling weight at the contact point. This weighting scheme ensures smooth transitions of stresses and displacements between the FEM and DEM domains.

On the DEM side, particle data such as mass and contact forces are dynamically scaled based on their position relative to the hybrid region. This scaling facilitates consistent interactions both between DEM particles and across the DEM-FEM interface. The coupling weights, computed based on spatial positioning, ensure accurate interpolation across the hybrid domain boundaries. Additionally, force equilibrium is maintained by adjusting forces and displacements dynamically, enabling consistent integration of the two modelling approaches. The coupling process is performed using the *Kratos CoSimulation Application*, which facilitates data exchange, iterative updates, and synchronisation between the FEM and DEM solvers. This modular framework ensures that the core solvers remain independent while enabling efficient and flexible coupling for multiphysics simulations. A concise overview of the coupling algorithm is illustrated in Fig. 3.4.

3.5 Summary

This chapter derived the energy-consistent variational framework used for the hybrid DEM–FEM analyses. On the discrete side, contacts follow a linear spring–dashpot law with damping tied to the restitution coefficient, complemented by a constant directional torque model for rolling dissipation. On the continuum side, the FE formulation stems from the weak form of Cauchy momentum and accommodates both linear elasticity and a non-associated Drucker–Prager plasticity. The two descriptions are coupled in an overlapping region with a smooth spatial weight that partitions mass and internal work, while an unweighted penalty-based coupling force enforces kinematic compatibility. Conservative particle-to-node and node-to-particle mappings via FEM shape functions preserve action–reaction and prevent spurious energy. This formulation underpins all simulations and verifications in the remainder of this chapter and is applied consistently in the subsequent application chapters. Where required, the precise constitutive choices and parameter values (for both DEM and FEM) are stated explicitly in the corresponding solver-specifications tables.

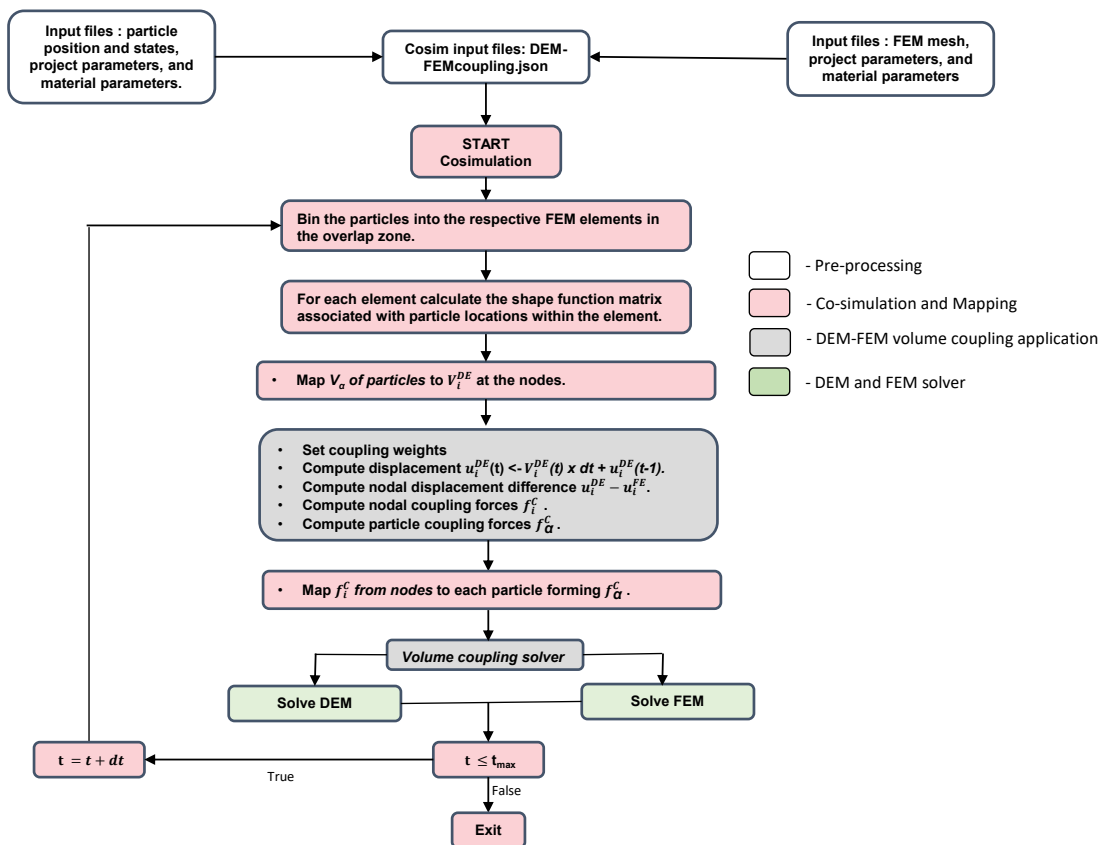


Figure 3.4: Overview of the algorithm used in DEM-FEM volume coupling

Implementation in Kratos Multiphysics

How to read this chapter

Readers who need nothing more than the high-level software layout can skip directly to the concise *Chapter Summary* in Section 4.7. That section distils all information required to compile, run, or extend the volume-coupled DEM–FEM solver. If instead you are interested in the full technical details, simply continue with the structured presentation that begins below.

4.1 Structure of implementation

All the computations presented in this study rely on the volume–coupling strategy described in Section 3.4.1. The entire source code has been developed inside the KRATOS MULTIPHYSICS ecosystem and is distributed publicly in the `volume_coupling` branch. This branch introduces a self-contained `DEM_FEM_VolumeCouplingApplication` that is built on top of these existing Kratos applications:

- **StructuralMechanicsApplication** (FEM),
- **DEMApplication** (DEM),
- **MappingApplication**,
- **ConstitutiveLawsApplication**,
- **CoSimulationApplication** (solver-agnostic coupling infrastructure).

Source-tree overview

The new functionality is realised in *three* locations:

applications/DEM_FEM_VolumeCouplingApplication

Three-level directory with FEM/DEM extensions:

- **Top level**

CMakeLists.txt, DEMFEM_volume_coupling_application.{cpp,h,py}

- **custom_elements**

volume_coupling_element.{cpp,h}, volume_coupling_particle.
{cpp,h}

- **custom_python**

add_custom_utilities_to_python.{cpp,h}, DEMFEM_volume_coupling_ python_application.cpp

- **custom_utilities**

utility_functions.h

applications/CoSimulationApplication

Thin layer that connects the new physics to the generic co-simulation framework:

- **python_scripts/coupled_solvers**

volume_coupling.py — implements DemFemVolumeCoupledSolver:

- (i) orchestrates DEM and FEM time stepping,
- (ii) executes the coupling operations,
- (iii) manages data exchange.

- **python_scripts/coupling_operations**

compute_dem_fem_volume_coupling_force.py — computes FEM nodal coupling forces via DEMFEMVolumeCouplingUtilities;

compute_dem_momentum.py — evaluates DEM momentum and contact forces.

- **python_scripts/solver_wrappers/kratos**

volume_coupling_structural_wrapper.py — specialised *structural wrapper* that toggles ACTIVATION_LEVEL in the ProcessInfo prior to and after output, enabling weight-dependent post-processing inside VolumeCouplingElement.

applications/DEMApplication

One-line patch in python_scripts/DEM_procedures.py:

```
import KratosMultiphysics.DEMFEMVolumeCouplingApplication as VCA
model_part.AddNodalSolutionStepVariable(VCA.PARTICLE_COUPLING_WEIGHT)
model_part.AddNodalSolutionStepVariable(VCA.DISPLACEMENT_MULTIPLIED_MASS)
model_part.AddNodalSolutionStepVariable(VCA.VELOCITY_MULTIPLIED_MASS)
model_part.AddNodalSolutionStepVariable(VCA.DEMFEM_VOLUME_COUPLING_FORCE)
```

These variables expose the DEM solver to the data required by the coupling utilities. No further core changes are necessary.

Build integration

The root-level `CMakeLists.txt` integrates the application seamlessly into the Kratos super-build; recompiling Kratos is sufficient to incorporate all additions.

Code modules

The new source files are grouped into three coherent packages:

- i) *Coupling utilities* — helper functions (`custom_utilities`) that handle geometry queries, search trees, and data exchange.
- ii) *VolumeCouplingElement* — continuum element (`volume_coupling_element`) derived from `SmallDisplacementElement`.
- iii) *VolumeCouplingParticle* — DEM sphere variant (`volume_coupling_particle`) with adaptive interaction radius.

Each component is registered in `DEM_FEM_VolumeCouplingApplication::Register()`; auxiliary Python scripts are auto-discovered by `CoSimulationApplication`. Compilation with the flag `-DENABLE_DEM_FEM_VOLUME_COUPLING=ON`, followed by selection of `DemFemVolumeCoupledSolver` in the CoSimulation JSON, grants immediate access to the new elements, particles, utilities, solver wrapper, and coupling operations.

4.2 Coupling utilities

The header file `utility_functions.h` (see Appendix C) provides an autonomous toolbox that is shared by both the DEM and FEM solvers. Every routine operates exclusively on the native Kratos data structures and MPI communicators, ensuring full thread-safety and perfect scalability on distributed meshes.

Available operations

- `AssignPointLoads` (optional) – applies frictional wall loads on selected FEM boundaries.
- `SetNodalCouplingWeightsOnFEMLinearly` – assigns hybrid weights by *linear* interpolation between two horizontal interfaces.
- `SetNodalCouplingWeightsFromLayers` – assigns hybrid weights from user defined *SubModelParts*, allowing hybrid layers of arbitrary geometry.

- `CalculateDisplacementDifference` – accumulates the DEM–FEM velocity mismatch over time.
- `CalculateNodalCouplingForces` – converts the accumulated mismatch into consistent nodal forces by means of a penalty parameter k_{\max} .
- `CalculateNodalDEMCouplingForces` – rescales the FEM penalty forces with the lumped mass m^* before they are sent back to the DEM solver.
- `CalculateMomentum` – stores $m \mathbf{u}$ and $m \mathbf{v}$ in the nodal variables `DISPLACEMENT_MULTIPLIED_MASS` and `VELOCITY_MULTIPLIED_MASS` for export to the FEM side.
- `CalculateDEMForces` – multiplies the external force `EXTERNAL_APPLIED_FORCE` by the nodal mass whenever the particle weight ϖ is non-zero, This is how the mapped mass averaged coupling forces are scaled up and applied to individual particles.

Hybrid–region weighting

Two complementary strategies translate the abstract concept of “hybrid region” into concrete nodal weights $\omega \in [0, 1]$, where $\omega = 0$ denotes a purely FEM node and $\omega = 1$ a purely DEM node:

- `SetNodalCouplingWeightsOnFEMLinearly` performs a one-dimensional interpolation between two horizontal planes located at $y = y_{\text{FEM}}$ and $y = y_{\text{DEM}}$. The routine is straightforward, robust, and therefore the default choice for the verification cases reported in Chapter 5.
- `SetNodalCouplingWeightsFromLayers` is completely geometry-agnostic. The user supplies a dictionary such as

```
{
  "DISPLACEMENT_layer1" : 0.0,
  "DISPLACEMENT_layer2" : 0.5,
  "DISPLACEMENT_layer3" : 1.0
}
```

where each key is the name of a `SubModelPart` and the value is the desired weight. The algorithm walks the Kratos model hierarchy in depth-first order and assigns the specified weight to all nodes that belong to a matching sub-part. Because selection is done by membership rather than by geometric checks, the same routine works for curved, cut-out, or dynamically changing hybrid zones without any additional tolerances.

Penalty enforcement and calculation of displacement difference

Coupling consistency is enforced through a velocity-based penalty term that penalises the relative motion between collocated DEM particles and FEM nodes. The procedure unfolds in three steps:

- i) `CalculateDisplacementDifference` integrates the instantaneous displacement mismatch $(\mathbf{v}_{\text{DEM}} - \mathbf{v}_{\text{FEM}}) \times \Delta \mathbf{t}$ into the history variable `PENALIZE_DISPLACEMENT`
 $\Delta \mathbf{u}+ = (\mathbf{v}_{\text{DEM}} - \mathbf{v}_{\text{FEM}}) \times \Delta \mathbf{t}$.
- ii) `CalculateNodalCouplingForces` multiplies that history variable by a user-defined maximum penalty ε and distributes the resulting force to the hybrid FEM nodes.
- iii) `CalculateNodalDEMcouplingForces` rescales the FEM forces with the lumped mass of DEM particles m_{DEM} and writes them to `DEM_FEM_VOLUME_COUPLING_FORCE`, from where they are picked up by the DEM solver and scaled to the mass of each DEM particle in the next coupling iteration.

4.3 Finite Element Method

Baseline solver

All finite-element computations are carried out with the unmodified **Structural Mechanics Application** supplied by Kratos Multiphysics. This application covers linear and non-linear kinematics, small- and large-strain formulations, contact, damping, and several related features. No changes are introduced in its source code; every coupling-specific addition resides in the `DEM_FEM_VolumeCouplingApplication`.

VolumeCouplingElement

Coupling on the FEM side is achieved through the class `VolumeCouplingElement`¹, which derives from `SmallDisplacement`. The first modification appears in `GetIntegrationWeight`: the standard Gauss weight W_0 is replaced by

$$W_{\text{FEM}}(\xi) = (1 - \bar{\omega}(\xi)) W_0, \quad \bar{\omega}(\xi) = \sum_{i=1}^n \psi_i(\xi) \omega_i,$$

1. File location: `applications / DEM_FEM_VolumeCouplingApplication / custom_elements / volume_coupling_element.cpp`

where $\psi_i(\xi)$ denotes the shape functions and $\omega_i \in [0, 1]$ the nodal coupling weights generated by the utilities in Section 4.2. This scaling factor progressively attenuates the element stiffness and mass matrices as the DEM contribution becomes dominant.

A second override, `CalculateOnIntegrationPoints`, post-processes Cauchy stresses and Green–Lagrange strains *only* while the analysis flag `ACTIVATION_LEVEL` equals 1. The flag is toggled automatically by the dedicated wrapper `volume_coupling_structural_wrapper.py`: it is set to 0 at the start of the simulation, switched to 1 immediately before `OutputSolutionStep`, and reset to 0 afterwards. During the brief interval in which the flag is active, the tensor returned by the base class is multiplied by $(1 - \bar{\omega})$ and stored in the output buffer; the numerical solution itself is not affected. This mechanism is used to weight stresses and add it with the DEM weighted stresses to form a combined Hybrid stress in the hybrid domain. This function override can be ignored if the user do not want to weight the stresses.

4.4 Discrete Element Method

Kratos provides an open-source discrete-element framework, the **DEMApplication** [60]. The volume–coupling workflow leaves the parent application untouched and adds a dedicated particle implementation:

VolumeCouplingParticle

The class `VolumeCouplingParticle`² derives from `SphericParticle`. Four member functions are overridden:

- i) `Initialize` – sets the particle coupling weight $\bar{\omega}$ to 1 at creation.
- ii) `GetMass` – returns $m^* = \bar{\omega} m$, i.e. the real mass scaled by the weight stored at the particle’s single node.
- iii) `EvaluateBallToBallForcesForPositiveIndentations` – applies a symmetric weight

$$\bar{\omega}_{12} = \frac{\bar{\omega}_1 a_2 + \bar{\omega}_2 a_1}{a_1 + a_2},$$

where a_1, a_2 are the arm lengths in the contact model, ensuring that weighting is done at the contact point.

2. Source file: `applications / DEMFEMVolumeCouplingApplication / custom_elements / volume_coupling_particle.cpp`

- iv) `EvaluateBallToRigidFaceForcesForPositiveIndentations` – rescales particle to wall forces in the same manner as for particle to particle contacts.

4.5 Cosimulation Application

The **CoSimulationApplication** controls the multi-physics time loop. The tools provided in the module include data transfer operators, coupling schemes (weak, iterative, etc.), relaxation methods, among others. The `CoSimulationApplication` is completely agnostic to which specific solvers the user chooses, but serves as an interface to ensure no steps are forgotten during the coupling operations. The flow chart showing the code structure for `CoSimulation` is shown in Fig. 4.1. For completeness, a step-by-step pseudo-code of the coupled time loop is given in Appendix B, which explains the complete sequence of the cosimulation.

For volume coupling a bespoke coupled solver, `coupled_solvers/volume_coupling.py`, derives from the generic `CoSimulationCoupledSolver`. Its core routine `DemFem-VolumeCoupledSolver::SolveSolutionStep` executes the sequence

- i) **DEM momentum**: `ComputeDemMomentum.InitializeCouplingIteration()` evaluates $m \mathbf{u}$ and $m \mathbf{v}$ on the DEM side.
- ii) **Synchronise** → **FEM**: `_SynchronizeInputData("structure")` transfers the momentum and mass fields to the structural solver.
- iii) **FEM coupling forces**: all remaining coupling operations call `InitializeCouplingIteration()` to assemble the penalty forces on hybrid FEM nodes.
- iv) **Synchronise** → **DEM**: `_SynchronizeOutputData` maps the FEM forces back to the DEM particles.
- v) **Finalise coupling**: `FinalizeCouplingIteration()` scales the forces with particle mass or applies additional loads, as required.
- vi) **Advance solvers**: `solver.SolveSolutionStep()` is invoked for both DEM and FEM wrappers.

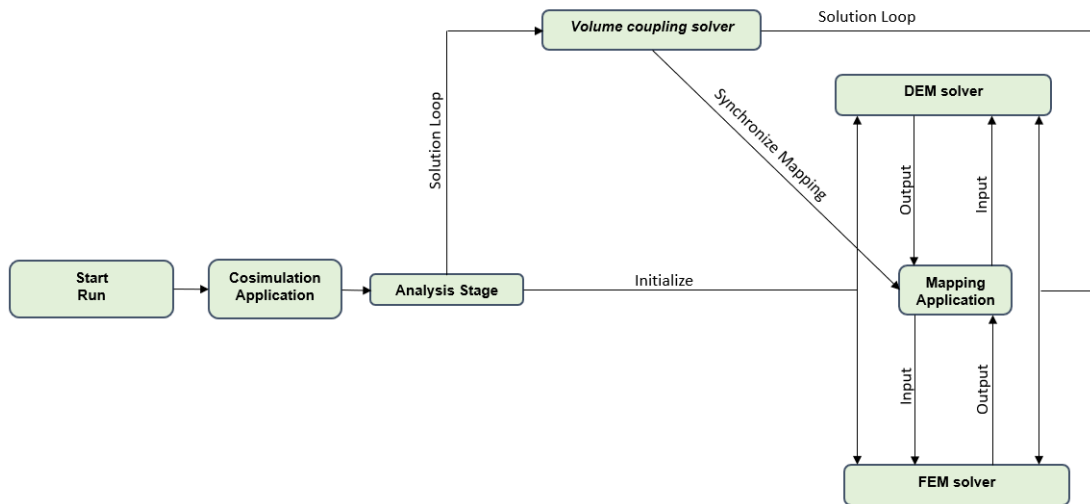


Figure 4.1: Kratos Multiphysics solver structure for DEM–FEM volume coupling

Modular Python components. The orchestration relies on three lightweight modules under `applications/CoSimulationApplication/python_scripts`:

- `coupling_operations/compute_dem_momentum.py` – gathers DEM momentum before every coupling iteration and scales the coupling forces applied to individual particles.
- `coupling_operations/compute_dem_fem_volume_coupling_force.py` – assembles and maps FEM penalty forces back to the DEM side.
- `coupled_solvers/volume_coupling.py` – implements `DemFemVolumeCoupled-Solver` and sequences the operations described above.

4.6 Case setup and input files

This section describes the input artefacts required to run the coupled DEM–FEM test case and how they relate to each other. A complete, working example is available under `applications/DEMFEVolumeCouplingApplication/tests`.

Overview

The case is organised around three groups of inputs:

- i) **DEM project parameters** (`.json`) controlling the discrete element simulation, plus a **DEM materials file** (e.g. `Materials.json`).
- ii) **FEM project parameters** (`.json`) controlling the finite element simulation, plus a **structural materials file** (`StructuralMaterials.json`).
- iii) **CoSimulation parameters** (`.json`) orchestrating the coupling (participants, data exchange, time integration, convergence).

On the mesh/geometry side, DEM typically uses two `.mdpa` files (one for particles and one for boundaries), while the FEM side uses a single `.mdpa` containing nodes/elements (and any boundary `SubModelParts`).

A typical folder layout is:

```
case/  
  CoSimulationParameters.json  
  dem/
```

```
ProjectParametersDEM.json
Materials.json
particles.mdp
boundaries.mdp
fem/
ProjectParametersFEM.json
StructuralMaterials.json
model.mdp
```

DEM inputs

ProjectParametersDEM.json Defines solver settings (strategy, scheme, contact search), material/contact laws, gravity and units, time stepping (e.g. `delta_time`, `end_time`), output, and process lists (e.g. injections, walls activation). See Appendix D.2 for the full JSON used in this work.

DEM Materials.json Collects particle and wall/boundary constitutive data and assigns them to `ModelParts/SubModelParts`. Typical entries include particle density, stiffness/Young's modulus, Poisson ratio, friction and rolling friction coefficients, restitution, and contact-law parameters.

DEM .mdp files

- *Particles .mdp*: provides particle `NODES` and optional `ELEMENTS` (for clusters), plus `PROPERTIES`.
- *Boundaries .mdp*: defines walls/obstacles (e.g. facet `CONDITIONS`) and grouping into `SubModelParts` for boundary processes.

FEM inputs

ProjectParametersFEM.json Specifies the FEM solver wrapper (application, solver type), linear solver and strategy settings, boundary/initial conditions, and output configuration. Time stepping may differ from DEM; the CoSimulation layer reconciles these. The exact JSON used here is given in Appendix D.3.

StructuralMaterials.json Defines material models and properties for the structural domain (e.g. Young's modulus, Poisson's ratio, density, damping/constitutive law options) and maps them to *ModelParts/SubModelParts*.

FEM model.mdpa Contains *NODES*, *ELEMENTS/CONDITIONS*, and *SubModelParts* for loads/BCs. Physical properties are assigned via the materials file and *PROPERTIES* blocks.

CoSimulation inputs

CoSimulationParameters.json This file declares:

- **Participants:** references to the DEM and FEM projects (paths to their *ProjectParameters*.json*), and their main *ModelPart* names.
- **Coupling sequence:** the ordered list of actions per (sub)time step (e.g. advance DEM → map field(s) to FEM → advance FEM → map back).
- **Data transfer operators:** mapping settings between non-conforming meshes (nearest-neighbour, projection/RBF/mortar) and variable pairs to exchange (e.g. displacements, tractions, contact pressures, porosity/volume fraction or source terms).
- **Coupling scheme and convergence:** weak (staggered) vs. strong coupling, optional Aitken under-relaxation, residual norms/tolerances, and maximum coupling iterations.
- **Time integration and control:** global start/end times, participant-specific *delta_time* handling, and whether one solver subcycles within another.
- **IO:** result exports and checkpoints (if enabled).

The exact configuration used in this thesis is provided in Appendix D.1.

Practical notes

- **Naming/paths:** ensure participant names in *CoSimulationParameters.json* match the keys used in the coupling sequence, that relative paths point correctly to the DEM/FEM project files, and that materials files are in the locations referenced by each participant.
- **Units:** DEM and FEM must be consistent (length, mass, time). Any intended scaling should be explicit in the project and materials files.

- **SubModelParts**: use named `SubModelParts` to target BCs, loads, mapped interfaces, and material assignments; reference these consistently in process lists and materials files.
- **Reproducibility**: keep solver versions and key parameter hashes in the case README for traceability (especially contact/friction laws and mapping settings).

4.7 Summary

This chapter presents the volume–coupled DEM–FEM implementation in KRATOS MULTIPHYSICS. The code is public in the `volume_coupling` branch and packaged as a dedicated *DEM FEM Volume Coupling Application* built on *Structural Mechanics Application*, *DEM Application*, *Mapping Application*, *Constitutive Laws Application*, and the solver-agnostic *CoSimulation Application*.

New functionality lives in `applications/DEM FEM Volume Coupling Application` (custom elements, utilities, Python bindings), with a thin `CoSimulation` integration (coupled solver and operations) and a minimal DEM patch to expose coupling variables.

The method comprises three building blocks. (i) *Coupling utilities* provide geometry/search and data-exchange helpers. (ii) *VolumeCouplingElement* (FEM) scales stiffness and mass—and optionally, for output only, stresses—by the local hybrid weight. (iii) *VolumeCouplingParticle* (DEM) applies weight-aware mass and contact evaluations. Hybrid weights $\omega \in [0, 1]$ are assigned either by linear interpolation between two planes or via geometry-agnostic `SubModelPart` layers.

Consistency is enforced with a velocity-based penalty: the DEM–FEM mismatch is time-integrated, converted to consistent FEM nodal forces through a user parameter ε , and mapped back to DEM particles with mass scaling. On the FEM side, `VolumeCouplingElement` derives from `SmallDisplacement` and scales Gauss weights by $(1 - \bar{\omega})$; an `ACTIVATION_LEVEL` flag gates optional stress weighting used only in post-processing. On the DEM side, `VolumeCouplingParticle` (from `SphericParticle`) returns a weighted mass $m^* = \bar{\omega}m$ and applies consistent weighting in ball–ball and ball–wall contact laws.

Time synchronisation is handled by `DemFemVolumeCoupledSolver`, which gathers DEM momentum per step, maps to FEM, assembles FEM penalty forces, maps forces back to DEM, and advances both solvers. The implementation is enabled at build time by setting the `-DENABLE_DEMFEM_VOLUME_COUPLING=ON` CMake flag. Once compiled, the workflow is activated by selecting `DemFemVolumeCoupledSolver` in the CoSimulation JSON configuration, which allows for the execution of the coupling sequence without modifying the core DEM or FEM solver modules.

A runnable case requires: (DEM) `ProjectParametersDEM.json` + `Materials.json` and `particle/boundary.mdpa`; (FEM) `ProjectParametersFEM.json` + `StructuralMaterials.json` and a single `model.mdpa`; and (coupling) `CoSimulationParameters.json` defining participants, mappings, scheme, and time control. A compact layout is:

```
case/  
  CoSimulationParameters.json  
  dem/ ProjectParametersDEM.json Materials.json particles.mdpa  
      boundaries.mdpa  
  fem/ ProjectParametersFEM.json StructuralMaterials.json model.mdpa
```

A complete, ready-to-run example is included under:
`applications/DEMFEMVolumeCouplingApplication/tests`.

Verification of the hybrid model

5.1 Methodology for verification of coupling method

The distinct natures of DEM and FEM raise questions regarding mechanical compatibility: DEM is a discontinuum method that admits particle rotations and sliding, whereas FEM represents a continuous displacement field between nodes. To detect and isolate strain–stress inconsistencies early, a progressive two-case plan is adopted—moving from an idealised setup to a more realistic configuration—while, in parallel, a principled rule for selecting the coupling penalty is developed:

- i) **Monodisperse 1D column.** A vertical column of frictionless DEM particles transitions through an overlap into a purely FEM region. This configuration tests vertical stress and strain transfer while exposing potential overlap artefacts (Section 5.2).
- ii) **Polydisperse uniaxial compaction.** A randomly packed, high-friction assembly, where slip and rotation are minimal, allows for the contrast of DEM lateral contact forces against the Poisson-driven lateral response in FEM. This probes the hybrid zone under conditions of realistic heterogeneity (Section 5.3).
- iii) **Penalty scaling law and verification.** A compact, discretisation-agnostic scaling for the coupling penalty is derived, namely $\varepsilon = CE/(dL_c)$ (or $\varepsilon = CK_{pp}^N/(d^2L_c)$), where E is the equivalent continuum modulus, d is a characteristic (or mean) particle diameter, and L_c is the overlap thickness. This relation is verified across multiple mesh, particle, and overlap configurations to establish practical choices for the constant C (Section 6.3).

Together, these tests probe the hybrid zone both as a physical model and as a transfer layer, identifying the limited circumstances under which the two descriptions may diverge, such as those due to weighting, mapping, or microstructural heterogeneity. The penalty-scaling analysis complements the verification by supplying an *a priori* recipe for ε that mitigates overlap artefacts and improves robustness across various materials, meshes, and overlap sizes. This chapter is based on a manuscript currently in preparation for submission to a peer-reviewed journal.

5.2 Granular column verification test

Figure 5.1 shows the verification test problem setup. The reference simulation comprises a granular column of particles in a rigid, frictionless container. To represent this system as a hybrid model combining FEM and DEM, it is essential to derive an equivalent continuum representation of the original reference domain. When modelling granular materials using FEM, it is necessary to calibrate the parameters in the constitutive relations based on reference (DEM) simulations.




Figure 5.1: Reference simulation domain and the equivalent hybrid domain.

5.2.1 Model setup and parameters

The hybrid model consists of DEM and FEM subdomains; while their material parameters can never be strictly identical owing to the discrete-versus-continuum nature of the two formulations, they must be calibrated to be as nearly equivalent as practicable to reproduce the material response of a purely DEM system. Assuming an isotropic linear elastic constitutive model for the FEM domain, two elastic constants require calibration: Poisson's ratio and Young's modulus. In this test problem, the particles are mono-disperse, frictionless, aligned vertically and subjected to purely uniaxial loading from the top, so the Poisson's ratio can be assumed to be zero in the continuum. Young's modulus is determined by equating the equivalent stiffness of the particles with that of the continuum. Additionally, the mass of the DEM particles from the reference simulation must be coarse-grained to match the density used in the FEM material model. A complete list of material and numerical properties are provided in Table 5.1 and Table 5.2.

In the reference discrete-element simulations, a lower contact stiffness is preferred to permit a larger stable time-step, thereby enhancing computational efficiency. Furthermore, Rayleigh damping is applied to the finite-element sub-domain to accelerate convergence toward a steady-state solution, as the primary focus of this study is the quasi-static response of the granular system. It should be noted that the damping coefficients in the DEM and FEM regions are independent and do not require numerical equivalence to maintain stability or accuracy. The damping parameters for both normal and tangential directions are derived from the coefficient of restitution following the methodology outlined by Schwager et al. [51], with identical values applied in both directions. The hybrid model is constructed with three zones: FEM, Hybrid, and DEM, to represent the reference domain, which consists entirely of DEM particles. The boundary conditions from the DEM reference solution are maintained in the hybrid model.

The discretisation is performed using eight-node cuboidal elements with linear shape functions, assuming a small-displacement formulation. Each element contains eight Gauss points, a configuration maintained consistently throughout this work. While the specific mesh dimensions for each simulation can be inferred from the domain figures provided in the respective chapters, the verification problem in this chapter employs a cubic mesh with an element size of 0.02 m, containing eight particles with a diameter of 0.01 m.

To facilitate post-processing along the vertical (y) axis, multiple planes perpendicular to the y -axis are defined at the vertical locations of the Gauss points. For the FEM sub-domain, each plane contains four Gauss points. Correspondingly, for the DEM sub-domain, binning is performed using sub-cuboids of size $0.02 \times 0.01 \times 0.02$ m. The resulting data points within the hybrid zone are generated by interpolating the discrete and continuum results to the FEM Gauss points and the DEM bin centres; this approach ensures that spatial variations are captured effectively, particularly near the boundaries of the hybrid domain.

The objective of this approach is to ensure that the response of the hybrid setup closely replicates that of the reference DEM setup, thereby validating the fidelity of the hybrid method. In the subsequent results, uniaxial compression is induced by imposing a downward velocity of 0.005 m/s on the upper boundary nodes for 0.025 seconds, after which they remain stationary for 0.2 seconds. Results are presented for the final timestep.

Table 5.1: DEM specifications for granular column verification test

Specification	Value
Diameter of particle	0.01 m
Solver type	Linear spring dashpot
Particle density	2500 kg/m ³
K_{normal} Particle-Particle	5000 N/m
K_{normal} Particle-Wall	10000 N/m
Coeff. of Restitution	0.1
Timestep	1×10^{-5} s

Table 5.2: FEM specifications for granular column verification test

Specification	Value
Solver type	Dynamic
Constitutive law	Linear Elastic
Rayleigh alpha damping	200 s ⁻¹
Young's modulus (E)	500 kPa
Poisson's ratio	0
Density	1312.5 kg/m ³
Analysis type	Non-linear
Timestep	1×10^{-5} s

5.2.2 Results

The compressive stress and strain behaviour of the granular column was analysed following the application of a small displacement at the top. In the hybrid domain, stress and strain computation was influenced by the magnitudes of weights applied at each location. In the formulation, the stresses were calculated by multiplying the stress values by the corresponding weights at the respective location. The stresses for the DEM part were calculated on the basis of the volume of the host element containing the particles; this procedure yields the weighted DEM stress. The total stress was then obtained by adding the weighted DEM and weighted FEM stresses. Consequently, the plots shown later in this chapter for the hybrid stress represent this combined total stress in the hybrid zone. Calculation of non-weighted DEM stress involves dividing the weighted particle contact forces by adequate areas according to the local weights. Ideally, the unweighted DEM and FEM stresses would coincide; however, weighting the individual contributions and summing them both simplifies stress evaluation on the DEM side and yields a single stress-field prediction that coherently integrates both DEM and FEM contributions within the hybrid zone. Hybrid stress fields can deviate locally from the exact solution, occasionally exhibiting stress spikes; this behaviour is examined further in Section 5.3.2.

Effect of the penalty parameter

The magnitude of the penalty parameter ε serves as an indicator of the stiffness of the coupling. Since the coupling force is introduced by penalizing the difference in displacements, a lower penalty value allows for greater displacement differences under a given coupling force, whereas a higher penalty value restricts the displacement differences more effectively. Consequently, when a force is applied to the top of the column, lower values of ε result in larger displacements at the top nodes compared to higher values of ε . Similarly, if a fixed displacement is applied at the top, the reaction force observed at the bottom decreases with lower penalty values. This implies that penalty parameter also significantly influences the stress distribution across the domain, as illustrated in Fig. 5.2. Lower penalty values yield a more uniform stress profile but tend to under-predict the overall stress magnitude relative to the reference solution provided by the discrete element method (DEM). In contrast, higher penalty values lead to a more non-uniform stress distribution in the hybrid region, often overshooting

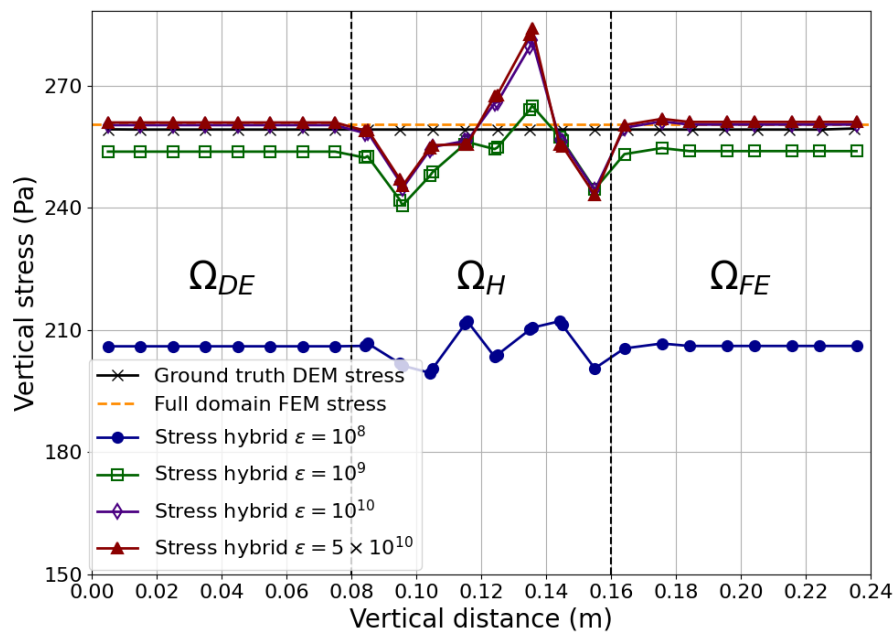


Figure 5.2: Vertical stress (σ_{yy}) profile in the hybrid model showing variation of stress profile for different penalty ϵ (Pa/m^2) distribution. A weighting of 0.01 to 0.99 is used across the hybrid domain.

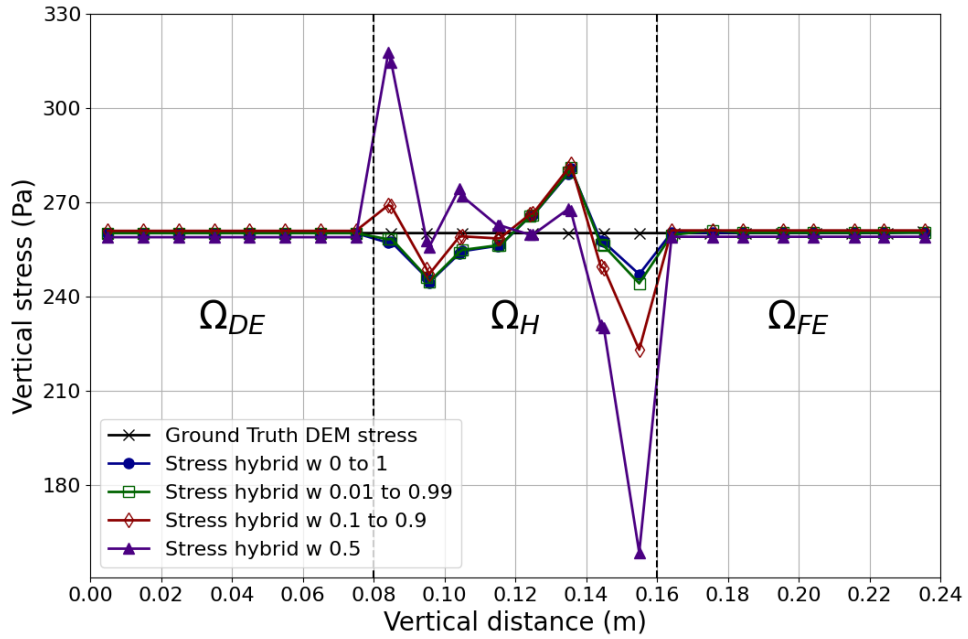


Figure 5.3: Vertical stress(σ_{yy}) profile in the hybrid model showing variation of stress profile for different weightage distribution. $\varepsilon = 1 \times 10^{10} \text{ Pa/m}^2$ was used.

the expected stress levels. In this verification problem, a penalty value of $\varepsilon = 10^{10}$ was found to provide the best agreement between the predicted stresses in the hybrid model and both the DEM reference stresses and the FEM-continuum stresses, as if the DEM and FEM models were independently representing the entire domain.

Effect of the weighting function

The effect of the weighting function on stress and strain distribution was also explored as shown in Figs. 5.3 and 5.4, respectively. Various weighting functions were tested, including a constant weighting of 0.5 throughout the hybrid zone, similar to the approach adopted by Chen et al. [4] in their formulation. It was observed that the stress became increasingly non-uniform in the hybrid zone when the weighting was applied more abruptly. This also influenced strain distribution as shown in Fig. 5.4; the hybrid strain depicted in the figure is determined by weighting and adding the FEM and DEM strains at each location. Abrupt changes in weights, for example from 1 to 0.5, resulted in non-uniform strains in the hybrid zone. Using a constant weighting function, such as $w = 0.5$, forces an equal contribution from both models, without considering the local behaviour of the system at the boundaries of the overlap region. This simplistic approach results

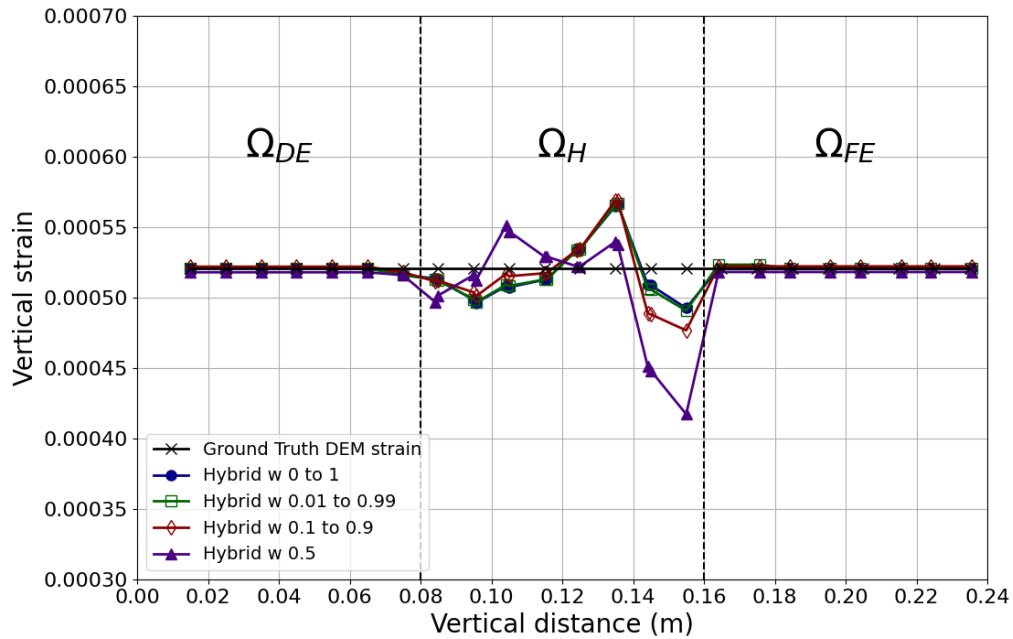


Figure 5.4: Vertical strain (ϵ_{yy}) profile in the hybrid model showing variation of strain profile for different weightage distribution. $\epsilon = 1 \times 10^{10} \text{ Pa/m}^2$ was used.

in a lack of smoothness in the transition between the two models, which manifests as discontinuities or non-physical behaviour in the computed solutions. Bauman et al. [2] have demonstrated that this approach can result in inaccurate representations of displacement and stress fields in discrete-continuum coupled systems, and that more adaptive weighting functions (like linear or cubic) are needed to ensure smooth transitions and better physical accuracy. A gradual weighting factor range between 0 and 1 was found to be a good approximation for maintaining consistency in stress and strain distribution within the domain. The corresponding stress-strain profile is presented in Fig. 5.5.

The combined stress and strain has a good agreement with the results from the reference DEM simulations and the strain distribution in the hybrid model remained relatively constant across the domain, indicating consistency in the hybrid formulation. Overall, the findings indicate that the choice of penalty parameter and weighting functions plays a critical role in determining the accuracy and uniformity of stress and strain profiles within the hybrid model.

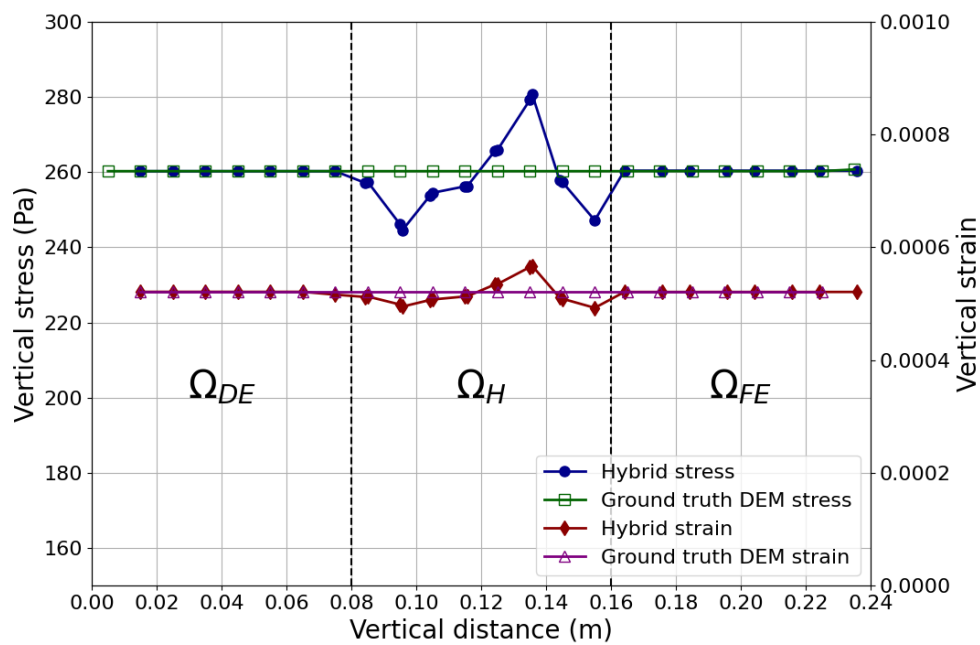


Figure 5.5: Vertical stress (σ_{yy}) and strain (ϵ_{yy}) profile in the hybrid model. The x-axis shows the vertical distance of the column starting from 0 m at the bottom. A penalty $\epsilon = 1 \times 10^{10} \text{ Pa/m}^2$ and weighting of 0 to 1 is used.

5.3 Polydisperse particle compaction test

5.3.1 Model setup and parameters

The configuration of the polydisperse particle system is depicted in Figure 5.6, where the simulation is set up within a rigid, frictionless container filled with particles of varying sizes. In this system, the particles are subjected to a high coefficient of contact friction, with a value of 2, ensuring that no slip occurs during compression. Under these conditions, the particles can be effectively modelled using an elastic continuum framework. The densely packed assembly of the polydisperse spheres is modelled as an isotropic elastic continuum whereby the Poisson's ratio, the Young's modulus, and the density are required for accurate representation. The two elastic parameters must be calibrated from the loading response of the reference DEM model. For an isotropic elastic medium confined within a rigid cylinder and subjected to controlled compression from the top, the two elastic parameters can be calibrated from the loading response using the following method:

Assuming vertical loading, the isotropic linear elastic equations can be written as:

$$\varepsilon_{xx} = \frac{1}{E} [\sigma_{xx} - \nu(\sigma_{yy} + \sigma_{zz})] \quad (5.1)$$

$$\varepsilon_{yy} = \frac{1}{E} [\sigma_{yy} - \nu(\sigma_{xx} + \sigma_{zz})] \quad (5.2)$$

For a rigid container, we assume $\sigma_{xx} = \sigma_{zz}$ and $\varepsilon_{xx} = \varepsilon_{zz} = 0$.

From Eq.(5.1),

$$0 = -\nu\sigma_{yy} + (1 - \nu)\sigma_{xx} \quad (5.3)$$

$$\frac{\sigma_{xx}}{\sigma_{yy}} = \frac{\nu}{1 - \nu} \quad (5.4)$$

Defining lateral stress ratio $\frac{\sigma_{xx}}{\sigma_{yy}} = k$, we get poisson ratio:

$$\nu = \frac{k}{1 + k} \quad (5.5)$$

From Eq.(5.2) , using $\sigma_{xx} = k\sigma_{yy}$ and $\sigma_{xx} = \sigma_{zz}$,

$$E\varepsilon_{yy} = \sigma_{yy} - 2\nu k\sigma_{yy} \quad (5.6)$$

$$E = (1 - 2\nu k) \frac{\sigma_{yy}}{\varepsilon_{yy}} \quad (5.7)$$

This calibration is achieved by compressing the particles within the rigid, frictionless container and analysing the lateral stress ratio during loading, as illustrated in Figure 5.7. Poisson's ratio is derived from this lateral stress ratio. Similarly, Young's modulus is determined from the stress-strain relationship obtained during compression tests, as illustrated in Figure 5.8. Additionally, the coarse-grained density is computed to match the mass distribution of the DEM particles in the FEM material model.

A complete list of the calibrated material properties is provided in Tables 5.3 and 5.4. The hybrid model is designed to closely replicate the mechanical behaviour of the reference DEM domain. The simulation domain is divided into three equal zones: the top third modelled using DEM, the middle third using a hybrid DEM-FEM approach, and the bottom third using FEM. In the following results, loading is applied by moving the top plate at a velocity of 0.0015 m/s for 0.3 seconds, after which it is held stationary for 0.2 seconds. The results are plotted for the final timestep. We use a penalty parameter of $\varepsilon = 10^{11} \text{ Pa/m}^2$ and linear weighting of 0 to 1 to obtain a reasonable agreement between the pure DEM and hybrid simulation results.

Table 5.3: DEM specifications for polydisperse particle compaction test.

Specification	Value
Avg. diameter of particle	0.0043 m
Range - particle diameters	0.0020 m to 0.0073 m
Particle density	2500 kg/m ³
K_{normal} Particle-Particle	500 N/m
$K_{\text{tangential}}$ Particle-Particle	500 N/m
Coefficient of Restitution	0.1
Coefficient of Friction	2
Timestep	1×10^{-5} s
No. of particles (hybrid)	3995
No. of particles (pure DEM)	5909
Wall parameters	COF = 0, COR = 0.1
K_{normal} Particle-Wall	1000 N/m




Figure 5.6: Reference simulation domain containing polydisperse particles and the equivalent hybrid domain.

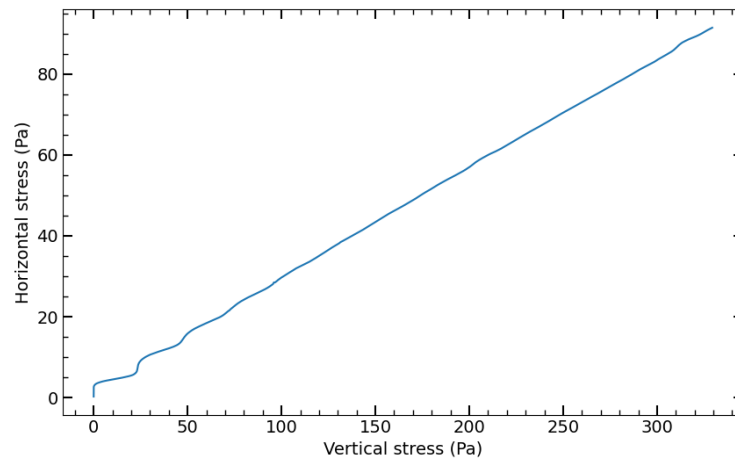


Figure 5.7: The lateral stress ratio (k) of 0.27 is derived from the horizontal and vertical stress data obtained via reference DEM simulations, using the last 60% of the data points from the plot for calculation.

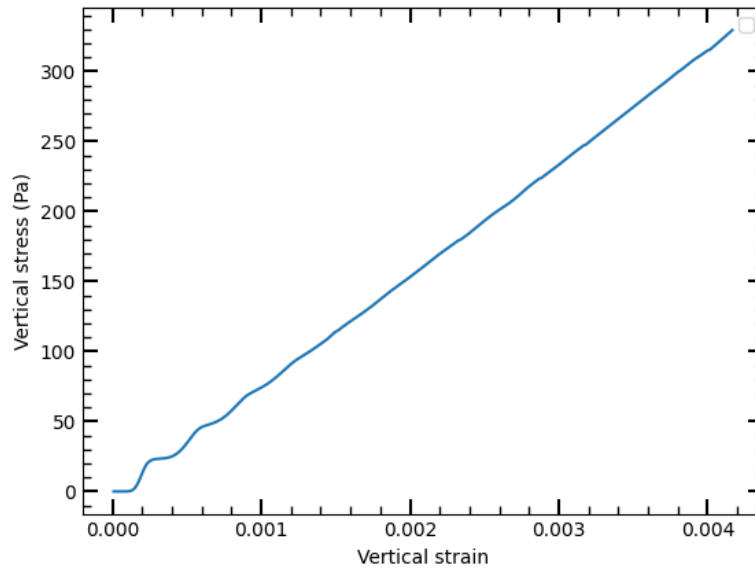


Figure 5.8: The stress-strain ratio, calculated as 81,233.16 Pa, was determined from reference DEM simulations by analysing the last 60 % of the data points from the plot.

5.3.2 Results

Stress profiles

Figures 5.9 and 5.10 present the vertical and horizontal stress profiles respectively, in the hybrid model. The predicted stresses in the hybrid model (shown as combined stress) are compared with the reference DEM stress and full-domain FEM stress (simulation done purely with FEM representing the whole domain). The vertical and horizontal stresses in the hybrid model, formed by adding the stress contributions from both the DEM and FEM parts, is a close match with the reference DEM and full-domain FEM stresses. Averaging across the available data points, the hybrid model predicts the vertical stress within 3% of the reference DEM values; mesh-refinement studies presented below demonstrate that this error can be reduced even further.

Figure 5.11 illustrates the lateral stress ratio $k = \sigma_h / \sigma_v$ obtained from the hybrid simulation in comparison to the reference results from DEM simulations, demonstrating an excellent agreement between the hybrid simulation and the reference DEM results. The lateral stress ratio $k = 0.27$ is generally uniform across the FEM, the hybrid and the DEM domains. This again confirms that the coupled DEM-FEM system indeed behaves in close agreement with the pure DEM particles under uniaxial compression.

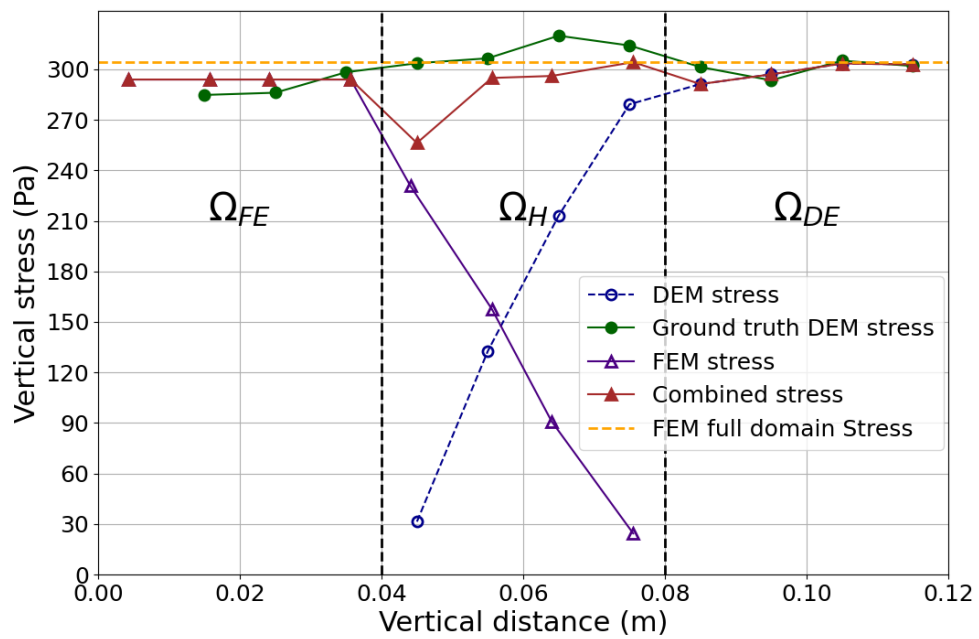


Figure 5.9: Vertical stress (σ_{yy}) profile in the hybrid model. The x-axis shows the vertical distance of the column starting from 0 m at the bottom. A penalty $\varepsilon = 10^{11} \text{ Pa/m}^2$ and a weighting of 0 to 1.0 are used.




Figure 5.10: Horizontal stress (σ_{xx}) profile in the hybrid model. The x-axis shows the vertical distance of the column starting from 0 m at the bottom. A penalty $\varepsilon = 10^{11}$ Pa/m² and a weighting of 0 to 1.0 are used.

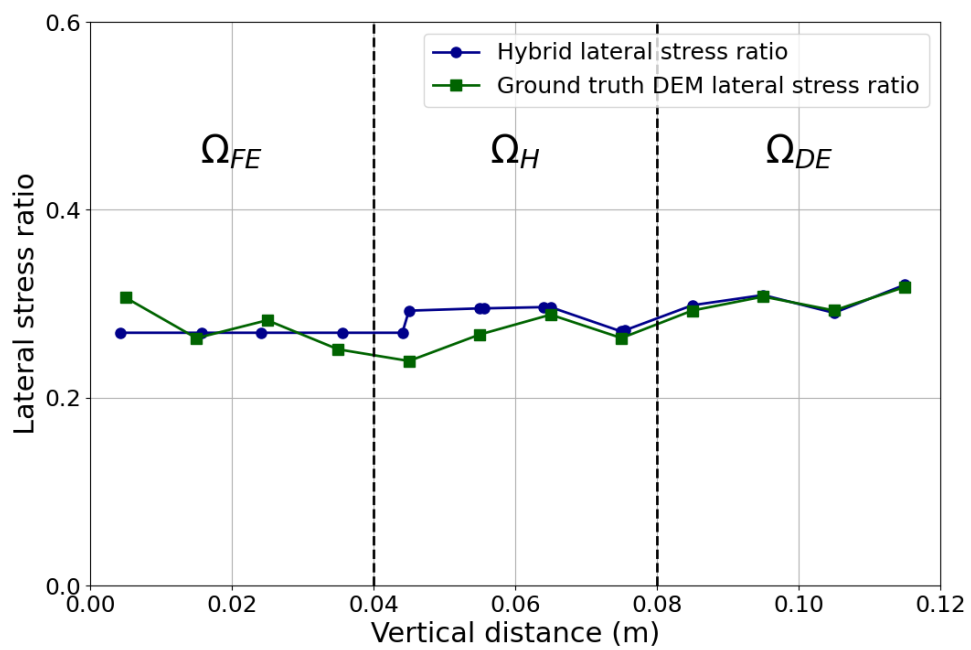


Figure 5.11: Lateral stress ratio for the hybrid simulation compared with reference DEM simulation results. A penalty $\varepsilon = 10^{11} \text{ Pa/m}^2$ and a weighting of 0 to 1.0 are used.

Table 5.4: FEM specifications for polydisperse particle compaction test.

Specification	Value
Solver type	Dynamic
Young's modulus (E)	71.933 kPa
Rayleigh alpha damping	200 s^{-1}
Poisson's ratio	0.212
Density	1539.51 kg/m^3
Analysis type	Non-linear
Timestep	$1 \times 10^{-5} \text{ s}$
Default No. of elements	36
Default mesh size	0.02 m

Effect of the penalty parameter and the weighting function

Similar to the 1D granular column example in Sec. 5.2, the effect of the penalty parameter and the weight function on the hybrid model's accuracy is investigated here. The ratio between the FEM mesh and particle sizes is kept the same as before. However, the penalty parameter is varied from 10^8 to 10^{11} . It can be observed in Fig. 5.12 and Fig. 5.13 that the effect of the penalty parameter becomes negligible after 10^9 . A low penalty parameter (e.g. 10^8) underestimates both the stress predictions, making the hybrid system "softer" compared to the pure DEM system. In both the 1D granular column and the polydisperse granular system, the overall stress response is affected by the penalty parameter. However, in the 1D case, the effect of the penalty parameter is also prominent near the coupled elements close to the boundaries of the hybrid zone. In contrast, this localised effect is not observed in the polydisperse system, likely because the polydispersity and randomness in the microstructure appear to spatially smooth the impact of the kinematic constraint on the hybrid zone. Note that for polydisperse systems, the penalty parameter was chosen to ensure convergence of the coupled solver. However, in the one-dimensional column case, localised stress spikes can occur in the hybrid region, surpassing the reference stress profile. In subsequent parametric sweeps varying both penalty and weighting, cases with a reduced penalty of 10^9 and weights 0.01–0.99 were evaluated. Benchmark comparisons showed no discernible change in trends, indicating insensitivity to penalty over 10^9 – 10^{11} and to weighting across 0–1 and 0.01–0.99.

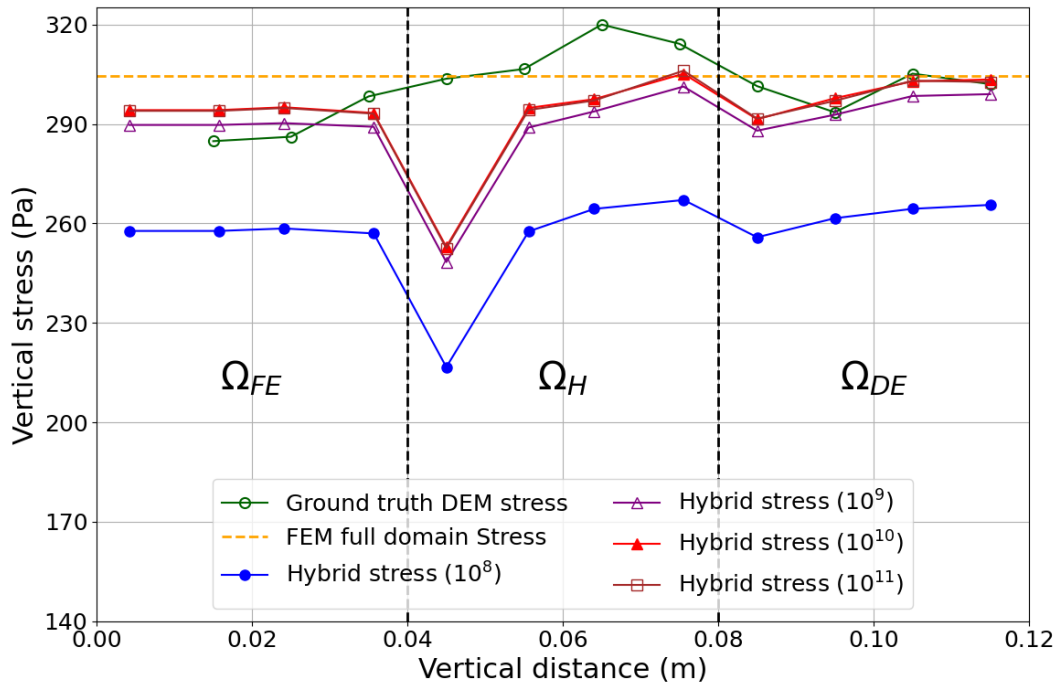


Figure 5.12: Vertical stresses (σ_{yy}) for different penalties. A weighting of 0.01 to 0.99 are used.

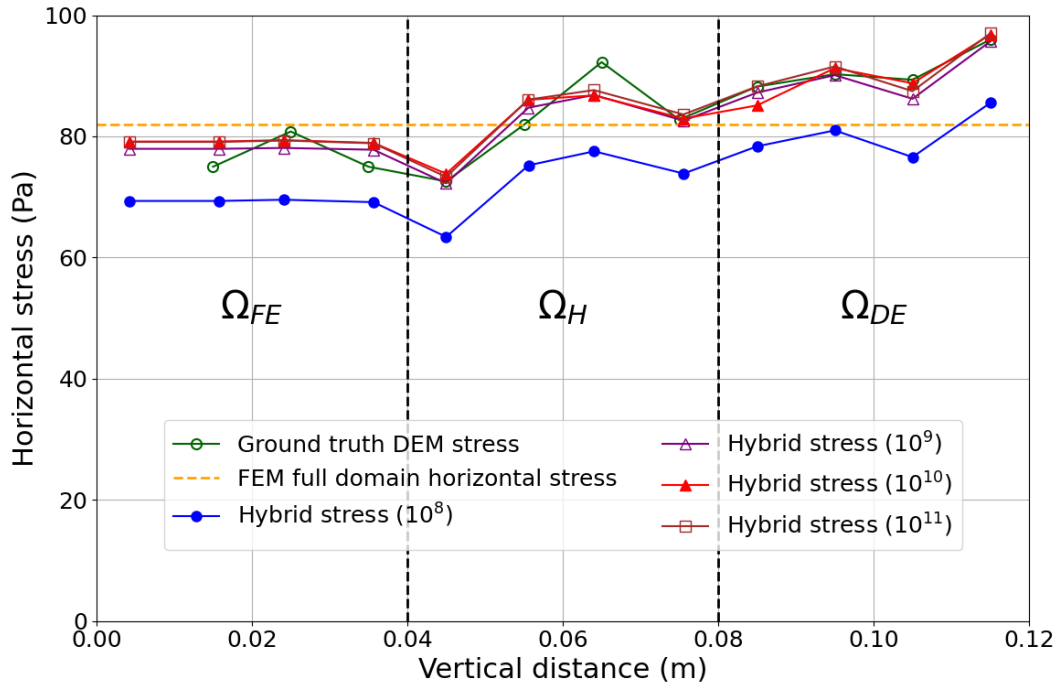


Figure 5.13: Horizontal stresses (σ_{xx}) for different penalties. A weighting of 0.01 to 0.99 are used.

Figure 5.14 illustrates the variation in vertical stress profiles corresponding to different weighting functions. The results exhibit trends consistent with those observed in the 1D granular column verification example. Notably, when a constant weighting factor of 0.5 is applied across the hybrid zone, the hybrid model tends to slightly underpredict the stress magnitudes outside the hybrid region. This discrepancy highlights the sensitivity of the stress distribution to the choice of weighting function and underscores the importance of adopting an appropriate weighting scheme to ensure accurate stress transfer between the DEM and FEM domains. The weighting function that smoothly transitions from 0 to 1 across the hybrid zone provides the closest stress predictions to the reference results, while effectively minimising stress fluctuations.




Figure 5.14: Vertical stresses (σ_{yy}) for different weighting distribution. A penalty of $\varepsilon = 10^9 \text{ Pa/m}^2$ is used.

Effect of the mesh-to-particle size ratio

The advantage of using a hybrid approach is that parts of a detailed many-particle (DEM) simulations can be confidently replaced by a continuum (FEM) model that represents the particle system's bulk response. The previous sections demonstrate that this is indeed possible using an equivalent continuum material model. The

computations were performed using an element-to-particle size ratio of 4.66 (i.e. approximately 100 particles per finite element). In the following, we vary the element-to-particle size ratio from 1.17 to 4.66 to investigate how the element-to-particle size ratio affects the accuracy of the hybrid model.

Fig. 5.15 illustrates the vertical stress profile for various element-to-particle size ratios, representing approximately 1 to 100 particles per finite element within the hybrid zone. The profiles reveal that: (1) the discrepancy between the reference DEM solution and the hybrid solution diminishes as the element-to-particle size ratio decreases, and (2) the accuracy is particularly affected by the element-to-particle size ratio near the boundaries of the hybrid zone. Additionally, as shown in Fig. 5.15, the initial dip in combined stress at a vertical position of 0.44 m reduces with decreasing mesh size.

We observe a larger stress spike at the FEM–hybrid interface than at the DEM–hybrid interface. Given the spatial variation of the FEM weight $(1 - w)$ within the hybrid zone, the spike increases where the FEM weight is higher or in closer proximity to the FEM, indicating that the artefact is primarily associated with FEM weighting rather than the DEM contribution. Although the stress spikes have negligible impact on stresses outside the hybrid zone, they may nevertheless induce instabilities in dynamic simulations. This numerical artefact could be further mitigated by (i) employing higher-order finite elements to increase Gauss-point density in the FEM-weighted region, (ii) refining the mesh across the hybrid zone to raise Gauss-point sampling, (iii) increasing the length of the hybrid zone by introducing more elements within it, or (iv) using enriched homogenisation methods such as those presented in [6].

5.4 Conclusions

Through two progressively complex benchmarks—a monodisperse 1D column and a polydisperse uniaxial compaction—we verified that the hybrid DEM–FEM formulation transfers stress and strain reliably across the overlap. The penalty-based coupling approach proved effective, with the choice of penalty parameter and weighting function playing a critical role in minimising artifacts in the overlap region. Stress and strain distributions in the hybrid model closely matched reference DEM simulations when the penalty parameter and weighting were properly calibrated and an equivalent continuum model is used in the FEM to represent the discrete particle system.

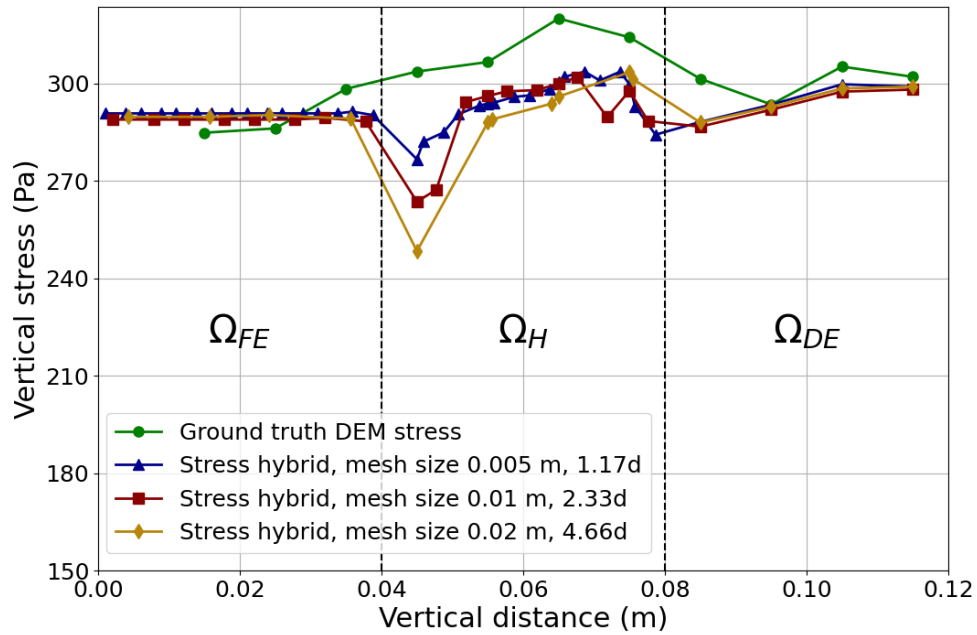


Figure 5.15: Vertical stresses (σ_{yy}) for different mesh sizes. A penalty $\varepsilon = 10^9 \text{ Pa/m}^2$ and a weighting of 0.01 to 0.99 are used.

For the polydisperse granular assembly, the hybrid model is shown to effectively capture the stress distributions which are consistent with the reference DEM results. The calibration of the continuum material properties, including the Poisson's ratio and the Young's modulus, is critically important for ensuring a seamless transition between the DEM and FEM subdomains and achieving a satisfactory prediction. FEM mesh refinement which allows a finer interpolation of the variables in the computation provides further improvement. The findings validate the hybrid framework as a reliable approach for modelling complex granular systems. These results provide practical guidance for selecting ε , weighting, and mesh resolution, and demonstrate that—with suitable parameter choices—the overlap region functions as an accurate transfer layer with negligible discontinuities.

Chapter 6

Dimensionless scaling of the penalty parameter

6.1 Introduction

Coupling the Discrete Element Method (DEM) and the Finite Element Method (FEM) in an overlapping domain has enabled tractable simulations of granular systems that mix quasi-static bulk response with localised micro-mechanics [2–4, 6, 18, 19]. Within this “bridging” or Arlequin setting, kinematic compatibility can be enforced either by a *penalty* that adds a coupling body-force density proportional to the displacement (or velocity) mismatch—widely used in granular hybrids and coarse-grained couplings [3, 5, 6]—or by *Lagrange multipliers* that impose the constraint in the weak form—prevalent in adaptive switching, high-stress tunnel/rock problems, and foundational Arlequin/bridging formulations [2, 4, 18, 26, 35, 44]. The comparative advantages of these strategies are reviewed elsewhere in the thesis; here we focus on the practical question of how to choose the magnitude of the penalty.

A penalty that is too small leaves a noticeable kinematic mismatch and under-transmits tractions through the overlap, while an excessively large penalty can create stiffness imbalance, local stress spikes near the overlap boundaries (especially with abrupt weighting), and adverse conditioning for the solver. To make this choice systematic rather than ad hoc, we propose a compact, discretisation-agnostic scaling

$$\varepsilon = C \frac{E}{d_{avg} L_c}, \quad (6.1)$$

where E denotes the Young's modulus of the material, d is a characteristic (or mean) particle diameter, L_c is the hybrid-zone thickness, and C is a dimensionless constant that absorbs nonidealities from weighting, mapping and engagement. The analysis underlying this scaling is quasi-static in nature and neglects inertial effects, which are assumed to be negligible in the regimes considered here.

When confinement effects are relevant, the effective stiffness governing normal traction transfer across the overlap is increased due to lateral stress coupling. This can be approximately accounted for by introducing a multiplicative factor $1/(1 - 2\nu K)$, where ν is an effective Poisson's ratio and K is the lateral stress ratio. This correction reflects the fact that, under confined conditions, normal deformation is partially constrained by transverse stresses, thereby amplifying the apparent stiffness that the penalty must counteract. The E -based scaling in Eq. (6.1) remains the central result; the corresponding discrete interpretation is deferred to the derivation.

As illustrated in Fig. 6.1, if the penalty parameter ε is held fixed while the material stiffness E is increased, the enforced kinematic compatibility becomes progressively weaker relative to the elastic response of the material. As a consequence, a larger displacement mismatch develops within the overlap, leading to an under-transmission of stresses and a growing deviation of the Hybrid prediction from the DEM reference solution, as measured by the mean stress deviation PD (see Eq. (6.3)). This trend demonstrates that a penalty chosen independently of material stiffness cannot provide consistent coupling performance across materials. Also, Cheng et al. (2023) reported a practical heuristic for confined granular hybrids that $\varepsilon d^2 \approx \text{const}$ across resolutions. Taken together, these observations motivate choosing a penalty that scales with the elastic response and a characteristic particle size, rather than prescribing a single absolute value.

The remainder of the chapter proceeds as follows. A short derivation based on a balance-of-resultants argument is first presented, linking the overlap coupling action to the confined continuum response. The proposed scaling is then verified across different configurations and particle stiffnesses. Finally, a convenient operating range for the dimensionless constant C is identified, and practical guidance is provided for selecting ε and diagnosing deviations, for example those arising from nonuniform engagement across the overlap.




Figure 6.1: Observation: with a fixed penalty $\varepsilon = 1 \times 10^8 \text{ Pa/m}^2$ and linear weights (0 \rightarrow 1) across the overlap, increasing the particle normal stiffness K_{pp}^N progressively increases the mean-stress deviation (PD) between the hybrid and DEM stress profiles. This motivates scaling ε with particle stiffness.

6.2 Methods in brief

6.2.1 Coupling summary

Within the overlap Ω^H , the hybrid virtual work is the weighted sum of DEM and FEM contributions plus a penalty term that enforces kinematic compatibility of homogenised DEM and FEM displacements, see Eq.(3.31).

$$\delta W_C = \int_{\Omega^H} \varepsilon(u^{\text{DE}} - u^{\text{FE}}) \cdot (\delta x - \delta X) \forall, \quad (6.2)$$

with linear (0→1) weights across Ω^H unless stated otherwise. Here the quantity

$$b^C = \varepsilon(u^{\text{DE}} - u^{\text{FE}})$$

is a *coupling force density* (body-force per unit volume) acting in the overlap coupling zone.

6.2.2 Error metric and setups

Accuracy is measured by the mean-stress percentage difference,

$$\text{PD}(\%) = 100 \times \frac{\sigma_m^H - \sigma_m^{\text{DEM}}}{\sigma_m^{\text{DEM}}}. \quad (6.3)$$

We verify the scaling on two canonical tests: (i) a mono-disperse, frictionless 1D column permitting an exact continuum analog with $\nu \simeq 0$; and (ii) a confined polydisperse compaction where large contact friction suppresses slip and an isotropic elastic calibration (from lateral-stress ratio and stress–strain slope) is meaningful. We vary particle stiffness, particle diameter, mesh sizes, particles-per-element, and overlap thickness L_c , and sweep ε .

6.3 Coupling penalty: scaling law and empirical behaviour

6.3.1 Derivation

Because b^C is a body-force density, the *resultant* coupling force transferred through the overlap is obtained by volume integration:

$$\mathbf{R}_{\text{cpl}} = \int_{\Omega^H} b^C \mathcal{V} = \int_{\Omega^H} \varepsilon (u^{\text{DE}} - u^{\text{FE}}) \mathcal{V} \quad (6.4)$$

Assuming near-uniform engagement across Ω^H at the level of scaling, the integral may be approximated by the hybrid volume times a representative mismatch ΔL_h :

$$\|\mathbf{R}_{\text{cpl}}\| \sim \varepsilon V_{\Omega^H} \Delta L_h, \quad V_{\Omega^H} = A L_c \quad (6.5)$$

We take y as the axis of action (loading/constraint direction) in a confined system, so axial resultants are carried by σ_{yy} . We equate \mathbf{R}_{cpl} with the axial stress resultants:

$$\varepsilon A L_c \Delta L_h \sim \sigma_{yy} A \implies \varepsilon L_c \Delta L_h \sim \sigma_{yy} \quad (6.6)$$

For a confined configuration with equal lateral stresses (denoting the lateral-stress ratio by $k = \sigma_{xx}/\sigma_{yy} = \sigma_{zz}/\sigma_{yy}$), an isotropic elastic estimate gives

$$\sigma_{yy} = \frac{E}{1 - 2\nu k} \varepsilon_{yy} \quad (6.7)$$

Identifying $\varepsilon_{yy} \approx \Delta \ell / \ell$ and writing $\ell = n d$ and $\Delta \ell = n \Delta \ell_d$, with $\Delta L_h \propto \Delta \ell_d$ under linear response, unknown proportionality factors cancel at the scaling level. Introducing a dimensionless constant C that absorbs mapping, weighting, and mild nonuniform engagement, we obtain the general scaling for a confined configuration

$$\varepsilon = C \frac{E}{(1 - 2\nu k) d L_c} \quad (6.8)$$

which reduces, for a monodisperse column confined by frictionless boundaries where $\nu = 0$, to

$$\boxed{\varepsilon = C \frac{E}{d L_c}} \quad (6.9)$$

Using $K_{pp}^N = E d$ for a cubic packing then gives the discrete counterpart (for $\nu = 0$)

$$\varepsilon = C \frac{K_{pp}^N}{d^2 L_c} \quad (6.10)$$

6.3.2 Verification across configurations

For each case in Table 6.1, reference DEM is run and ε is calibrated such that $PD = -5.00\% \pm 0.03\%$. For clarity, the axes and particle–mesh layout used in Case 3 are illustrated in Fig. 6.2. The inferred $C = \varepsilon d L_c / E = \varepsilon d^2 L_c / K_{pp}^N$ versus K_{pp}^N is shown in Fig. 6.3. Across Cases 1–3 and 5, $C \approx 0.77$ is nearly case-invariant; Case 4 is slightly lower (~ 0.69), consistent with mild nonuniform coupling–traction engagement across L_c .

Table 6.1: Configurations used for scaling verification: particle diameter d , hybrid-zone thickness L_c , element size, and particles per element n .

Case	d [m]	L_c [m]	Element size (x, y, z) [m]	n
Case 1	0.01	0.08	(0.01, 0.01, 0.01)	1
Case 2	0.01	0.08	(0.01, 0.02, 0.01)	2
Case 3	0.01	0.08	(0.02, 0.02, 0.02)	8
Case 4	0.02	0.08	(0.02, 0.02, 0.02)	1
Case 5	0.01	0.04	(0.01, 0.01, 0.01)	1




Figure 6.2: Axes and layout for Case 3 ($d = 0.01$ m, element 0.02^3 m³, $n = 8$).




Figure 6.3: Penalty scaling across Cases 1–5 for target PD = −5% : C clusters near 0.77 when overlap engagement is uniform; deviations reflect small nonuniformities.

To reduce PD further, we instead target essentially zero deviation on Case 3, the calibrated penalties collapse to $C \approx 9\text{--}10$. Fig. 6.4 illustrates this local calibration, and the global trend is made explicit in Fig. 6.5, which plots PD versus the dimensionless C with a secondary (top) axis for $\varepsilon = CE/(dL_c)$. The figure shows a clear *saturation plateau*: a dashed vertical line at $C=10$ marks the onset of the high-fidelity regime, and moving to the right of this line ($C > 10$) the curves remain essentially flat with $\text{PD} \approx 0\%$. Consistent with this, $\varepsilon \approx 10E/(dL_c)$ (or $\varepsilon \approx 10K_{pp}^N/(d^2L_c)$) drives $\text{PD} \rightarrow 0\%$, and further increases in ε produce no discernible improvement.




Figure 6.4: Case 3 penalties tuned for target $\text{PD} = -0.01\%$: the inferred constant is $C \approx 9.23$, i.e. a high-fidelity saturation regime with $C \sim 10$.

6.3.3 Polydisperse particle compaction in a rigid container: penalty sweep

Setup and calibration. We test the proposed penalty scaling in the polydisperse granular assembly problem with confined rigid frictionless rectangular container. The domain $(L_x, L_y, L_z) = (0.06\text{ m}, 0.12\text{ m}, 0.06\text{ m})$ is partitioned along y into three layers—a top DEM region (Ω_{DE}), a middle overlap (Ω^{H}), and a bottom FEM region (Ω_{FE})—as




Figure 6.5: Mean stress deviation PD (%) versus the dimensionless penalty C (bottom axis), with the corresponding $\varepsilon = CE/(dL_c)$ shown on the top secondary axis for Case 3. The dashed vertical line at $C=10$ marks the onset of the high-fidelity regime; for $C > 10$ the curve plateaus with $PD \approx 0\%$.

shown in Fig. 5.6. With mean particle diameter $d = 0.0043$ m, penalties from the simple scaling (6.9) at $C = 10$ give $\varepsilon = 4.19 \times 10^9$ Pa/m² for the lower-stiffness case and $\varepsilon = 1.8 \times 10^{10}$ Pa/m² for the higher-stiffness case. Including the confinement factor $1/(1 - 2\nu K)$ would only slightly increase these targets; choosing $C = 10$ under the simple form already yields high-fidelity responses. A high coefficient of friction (COF) is used to promote an approximately linear load–displacement response.




Figure 6.6: Polydisperse system, lower contact stiffness: penalty sweep around the reference $C = 10$ level. Parameters: $K_{pp}^N = K_{pp}^T = 500$ N/m; COF = 2; imposed displacement $\Delta y = 0.0045$ m.

Observed deviations and saturation. The resulting mean-stress deviations are PD = -0.150% for the lower-stiffness system and PD = 2.324% for the higher-stiffness system. Parametric sweeps indicate a *plateau* at $C \approx 10$: increasing ε beyond these values yields negligible changes in the stress response. The residual PD therefore reflects DEM–continuum calibration (e.g. consistency of E with the contact law and fabric, choice of representative d , and measured L_c) rather than inappropriate coupling strength. Within calibration-consistent settings, $C \approx 10$ provides the best attainable prediction relative to the DEM reference solution at both stiffness levels.




Figure 6.7: Polydisperse system, higher contact stiffness: penalty sweep around the reference $C = 10$ level. Parameters: $K_{pp}^N = K_{pp}^T = 2000\text{N/m}$; COF = 10; imposed displacement $\Delta y = 0.015\text{ m}$.

6.4 Scope and caveats

In practice, choose $\varepsilon \approx 10E/(dL_c)$ and use smooth (e.g., linear or cubic) weights across Ω^H [2]. Because the DEM and continuum descriptions are calibrated from different observables, small mismatches in E , K_{pp}^N , representative d , or overlap normalisation can introduce a bias floor in PD that persists even when ε is very large. Our sweeps consistently show a saturation plateau around $C \approx 10$: increasing ε beyond this range brings no further accuracy gains and may only stiffen the algebraic problem. The scaling remains predictive provided the coupling weights are well normalised and the continuum calibration is consistent with the DEM configuration. This scaling may also be used in elasto-plastic cases when E and ν are interpreted as pseudo-elastic parameters that include the bulk plastic response along the relevant load path.

We have also assumed a hybrid zone with uniform cross-sectional area. When the area varies along the coupling direction (e.g., radially changing sections), the balance of coupling resultants suggests a local multiplicative correction to the simple scaling. In practice, however, the simple form with $C \approx 10$ already places the solution near the high-fidelity plateau, so explicit cross-sectional factors are often unnecessary.

If the hybrid-zone thickness L_c varies within a given problem, using a single, fixed penalty ε will generally lead to spatial variations in PD within Ω^H . Prescribing a fixed dimensionless constant C instead, and updating the penalty locally via $\varepsilon(x) = CE/(dL_c(x))$, allows the system to adapt the coupling strength and maintain a more uniform PD across the hybrid zone.

6.5 Conclusions

A single-parameter scaling, $\varepsilon = CE/(dL_c)$, organises the choice of the coupling penalty in DEM–FEM volume coupling to achieve near-kinematic compatibility between DEM and FEM displacements. Empirically, $C \approx 10$ yields a high-fidelity plateau beyond which larger penalties provide no further benefit. For the polydisperse assembly with mean diameter $d = 0.0043$ m and the two stiffness levels tested, the penalties implied by $C = 10$ produce PD values of -0.150% and 2.324% ; the remaining discrepancy is governed by DEM–continuum calibration rather than insufficient coupling strength. The rule is simple to apply and robust to mesh and material variations.

Elasto-plastic hybrid modelling of granular media

7.1 Uniaxial compaction of an elasto-plastic polydisperse granular column

Building on the previous polydisperse compaction study and our sensitivity analysis of the coupling parameters, we extend the hybrid approach to elasto-plastic behaviour by coupling a DEM domain with a continuum FEM domain governed by elasto-plastic Drucker–Prager (DP) model. The DP model is adopted here as a deliberately simple elasto-plastic continuum representation, rather than as an attempt to reproduce the full complexity of granular constitutive behaviour. The primary objective of this study is to assess the performance of the hybrid DEM–continuum coupling under elasto-plastic deformation, and to examine how stress transmission, localisation and redistribution evolve within the overlap region, particularly under cyclic loading. By employing a basic pressure dependent yield criterion, the DP model enables systematic observation of stress variations within the hybrid framework while minimising additional modelling complexity. We first revisit the same polydisperse configuration to establish DP benchmarks by varying the (interparticle) friction coefficient, and then proceed to an industrial-scale silo simulation within the elasto-plastic framework.

7.2 Problem definition and numerical setup

The configuration of the polydisperse particle system is the same as in Figure 5.6, where a rigid, frictionless container is used to constrain the lateral deformation. Nevertheless, three sets of DEM simulations are run with particle coefficients of friction (COF) = 0.4, 0.8, and 10.0. The complete list of DEM properties are shown in Table 7.1. The FEM domain employs a rate-independent DP constitutive law, whose cohesion is set to a small value of 10 Pa for numerical stability. Table 7.2 collects the FEM parameters that remain unchanged across all cases—that is, those not affected by the COF-dependent elasto-plastic parameter calibration.

7.3 Determination of elastic- and elasto-plastic material parameters

As in the polydisperse compaction study, Young's modulus E and Poisson's ratio ν were first determined directly from the DEM compression response using Eqs. (5.5) to (5.7) in Sec. 5.3. Because these values are inferred from a loading path involving plasticity, the resulting ν can be regarded as a *pseudo-elastic* property calibration that implicitly reflects plastic deformation present in the simulated response. With (E, ν) thus fixed for each coefficient of friction (COF), the elasto-plastic parameters of the Drucker–Prager (DP) model were identified by Bayesian calibration against the DEM stress histories. The DEM configuration and the FEM/coupling setup used throughout this section are summarised in Tables 7.1 and 7.2, respectively, while the final calibrated parameter sets are reported in Table 7.3.

The Bayesian inference was carried out with the open-source `GrainLearning` toolkit [61, 62], which formulates parameter estimation for history-dependent models as an *iterative Bayesian filtering* problem. In the present context, the internal friction angle ϕ and dilation angle χ were treated as uncertain, with independent uniform bounds ϕ in $[3^\circ, 60^\circ]$ and χ in $[5^\circ, 90^\circ]$. The observation vector comprised the vertical and lateral stresses, $(\sigma_{yy}, \sigma_{xx})$, sampled along the imposed strain history from the continuum (Kratos) compression test. A diagonal multivariate normal error model with normalized tolerance `sigma_tol = 0.01` was used for the likelihood.

The calibration procedure follows a sequential Monte Carlo (SMC) procedure, in which the unknown material parameters are represented by an evolving ensemble of weighted samples rather than by a single deterministic estimate. At the start of each iteration, a finite set of parameter realisations is drawn from a proposal density, and each realisation is propagated through the forward FEM model to generate a corresponding stress–strain response. This ensemble-based strategy motivates the term Monte Carlo, as the parameter space is explored through repeated stochastic sampling rather than gradient-based optimisation.

The sequential aspect arises from the incremental assimilation of the stress–strain history along the loading path. As loading progresses, each parameter sample is assigned an importance weight reflecting its consistency with the observed response, such that poorly performing realisations are progressively downweighted while plausible ones are retained. This gradual incorporation of information improves numerical robustness when calibrating nonlinear elasto-plastic models, compared to single-shot or fully deterministic calibration strategies.

At the end of each iteration, the weighted ensemble is used to construct a probabilistic surrogate of the joint parameter–observation density in the form of a Gaussian mixture, implemented here via a Dirichlet-process mixture model. This surrogate serves as an adaptive proposal distribution for subsequent iterations, focusing sampling effort in regions of high posterior probability while preserving sufficient diversity. Iterations are continued until convergence is achieved, as indicated by the coefficient of variation of posterior expectations falling below the prescribed threshold.

The maximum-a-posteriori (MAP) estimates of (ϕ, χ) from the final iteration were then paired with the fixed (E, ν) to yield the continuum property sets listed in Table 7.3. These calibrated sets are used without further adjustment in the pure-FEM and hybrid DEM–FEM simulations reported in the remainder of this chapter.

7.4 Micro-slip statistics

The difference in the constitutive parameter values shown in Table 7.3 is fundamentally caused by the amount of inter-particle slip (or mobilised friction) although the initial particle configuration is identical. To quantify inter-particle slip we post-process the mobilised friction, defined at every contact as the tangential-to-normal force ratio

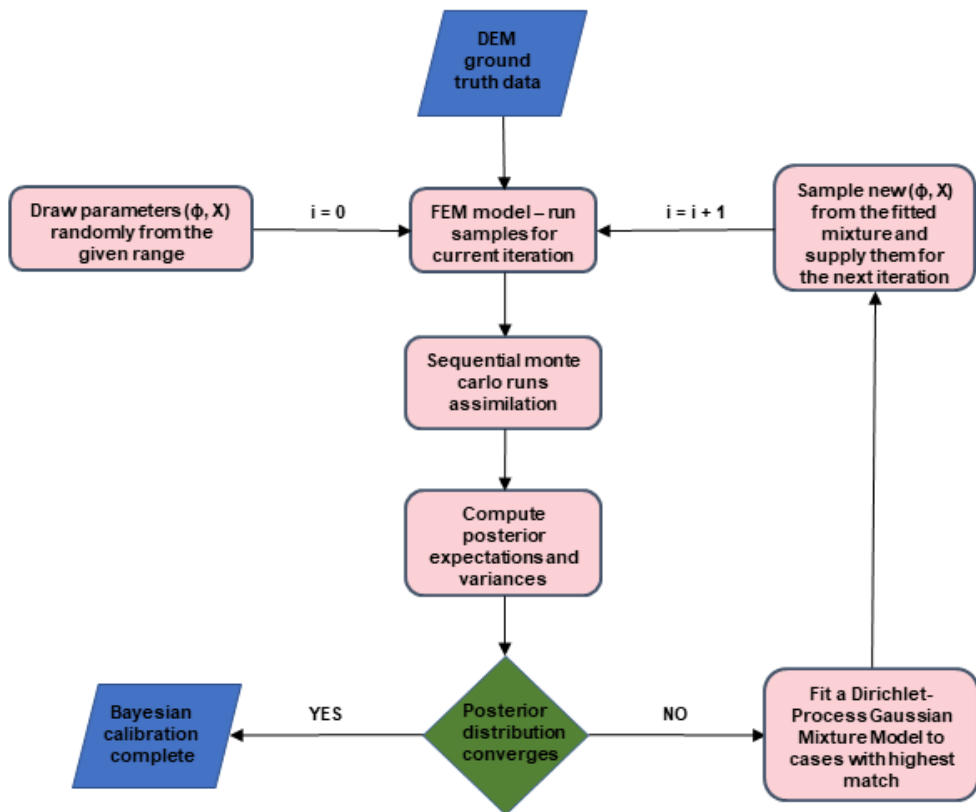


Figure 7.1: Schematic of the iterative Bayesian calibration procedure in GrainLearning used to identify the Drucker–Prager plasticity angles (ϕ, χ) . DEM target curves define the prior; batches of Kratos FEM simulations are assimilated via sequential Monte Carlo (SMC) with DPGMM-adapted proposals; parameters are resampled and iterated until the posterior coefficient of variation falls below 2%.

Table 7.1: DEM specifications for uniaxial compaction of an elasto-plastic polydisperse granular column

Specification	Value
Avg. diameter of particle	0.0043 m
Range - particle diameters	0.0020 m to 0.0073 m
Particle density	2500 kg/m ³
K_{normal} Particle-Particle	2000 N/m
$K_{\text{tangential}}$ Particle-Particle	2000 N/m
Coefficient of Restitution	0.1
Coefficient of Friction	0.4, 0.8, and 10
Coefficient of Rolling Friction	0.02
Timestep	1×10^{-5} s
No. of particles (hybrid)	3995
No. of particles (pure DEM)	5909
Wall parameters	COF = 0, COR = 0.1
K_{normal} Particle-Wall	4000 N/m

Table 7.2: FEM and coupling specifications for uniaxial compaction of an elasto-plastic polydisperse granular column

Specification	Value
Solver type	Dynamic
Rayleigh alpha damping	20 s^{-1}
Rayleigh beta damping	0.01 s
Density	1539.51 kg/m ³
Analysis type	Non-linear
Timestep	1×10^{-5} s
No. of elements	36
Mesh size	0.02 m
Penalty (ϵ)	10^{11} Pa/m^2
Weighting	Linear (0 to 1)

Table 7.3: Calibrated Drucker–Prager parameters for each DEM friction case.

Parameter	DEM particle COF		
	0.4	0.8	10
Young's modulus E (kPa)	250	315.02	402.50
Poisson's ratio ν	0.25	0.195	0.1262
Friction angle ϕ ($^{\circ}$)	24.24	32.20	44.74
Dilation angle χ ($^{\circ}$)	8.32	7.50	7.07

$\eta = \|\mathbf{f}_t\|/\|\mathbf{f}_n\|$, using the pure-DEM results at the peak compressed state. Figure 7.2 plots the distribution of the normalised mobilised friction η/COF for each friction case, expressed as the percentage of contacts in each histogram bin; the contacts at $\eta/\text{COF} = 1$ mark the transition from sticking to sliding behaviour. These distributions clearly show the tendency of increasing friction mobilisation as COF decreases from 10 to 0.4.




Figure 7.2: Distribution of contact-level normalised mobilised friction η/COF for the three friction coefficients.

7.5 Loading–unloading response for $\text{COF} = 0.8$

Figures 7.3 and 7.4 compare the cyclic response obtained from three simulations: (i) a pure DEM model, (ii) a pure FEM model calibrated with the Drucker–Prager parameters of Table 7.3, and (iii) the DEM–FEM hybrid model. Figure 7.3 shows the vertical stress–strain response σ_{yy} – ϵ_{yy} of the granular assembly. The loading branch lies above the unloading branch, producing the expected hysteresis loop. The pure DEM model captures the material’s nonlinear elasto-plastic behaviour during the unloading path providing the reference for comparison. The pure FEM model, however, with a simple DP constitutive law, is unable to capture the subtle hardening behaviour during compaction, but approximates the stress-strain curves during loading and unloading with two different slopes. This key difference between the pure DEM and FEM models can be observed from the evolution of the lateral-to-vertical stress path σ_{xx} – σ_{yy} , shown in Fig. 7.4.

In both figures the stress–strain curves given by the hybrid model lie closer to the DEM reference during both loading and unloading, especially when the stress level is at the low and high end. In the lateral-versus-vertical stress plot the loading branch lies below the unloading branch, forming a counter-clockwise loop. It seems that σ_{xx} is consistently lower than the DEM reference solution during unloading, possibly due to the irreversible change of the granular assembly’s microstructure. The good agreement between the hybrid and the pure DEM solution confirms the consistency of the elasto-plastic coupling.

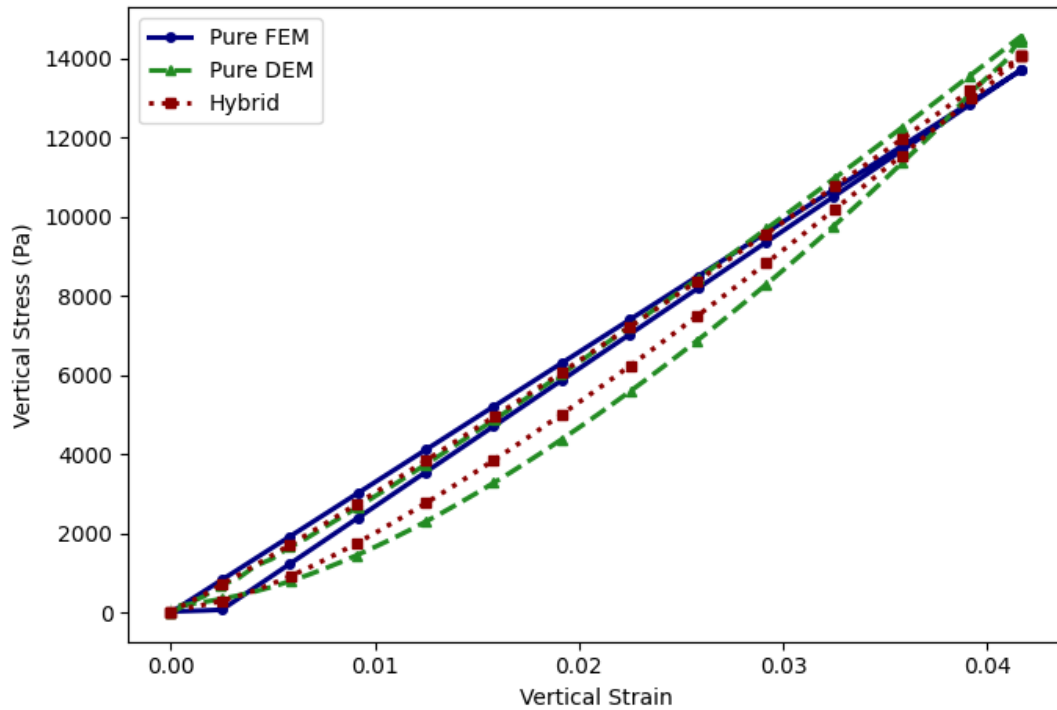


Figure 7.3: Axial response for COF = 0.8: vertical stress σ_{yy} versus vertical strain ϵ_{yy} during a complete loading–unloading cycle. Comparison between Pure DEM, Pure FEM and Hybrid models.

7.6 Influence of particle friction on macroscopic behaviour

Figures 7.5 and 7.6 compares the hybrid results with respective reference DEM results for the three investigated coefficients of friction (COF = 0.4, 0.8, 10) , all subjected to an identical imposed vertical-strain load. With increasing inter-particle friction the axial response becomes progressively stiffer (Fig. 7.5), and while the lateral-to-vertical




Figure 7.4: Lateral response for COF = 0.8: lateral stress σ_{xx} versus vertical stress σ_{yy} over a complete loading–unloading cycle. Comparison between Pure DEM, Pure FEM and Hybrid models.

stress ratio decreases (Fig. 7.6), trends that agree with classical granular-plasticity theory. To quantify the match between the hybrid and pure-DEM curves we use the *coefficient of determination*,

$$R^2 = 1 - \frac{\sum_i (\sigma_i^{\text{DEM}} - \sigma_i^{\text{hyb}})^2}{\sum_i (\sigma_i^{\text{DEM}} - \bar{\sigma}^{\text{DEM}})^2},$$

where σ_i^{DEM} denotes the stress from the pure-DEM simulation at sample point i , σ_i^{hyb} is the corresponding stress from the hybrid FEM–DEM model, $\bar{\sigma}^{\text{DEM}}$ is the arithmetic mean of all DEM stresses. A value of $R^2 = 1$ indicates perfect agreement.

Across all COFs, good agreement is obtained between the pure DEM and the hybrid solution, with a R^2 value larger than 0.95. For σ_{yy} in particular, a trend of increasing R^2 with COF can be observed, which is consistent with micro-slip statistics, as discussed in Sec. 7.4. The results of this benchmark test confirm that the volume-coupling strategy leads to accurate hybrid FEM-DEM models (with very high R^2 values in comparison to DEM reference solution) although the FEM model is limited in capturing the material's nonlinear in a loading-unloading cycle.

7.6.1 Stress and strain distribution in the hybrid domain

Figures 7.8, 7.9, and 7.10 compare, for three friction levels (COF = 0.4, 0.8, 10), the vertical stress field σ_{yy} and the lateral stress ratio $K = \sigma_{xx}/\sigma_{yy}$ sampled at three representative stages of the load–unload cycle: midway loading ($\epsilon_{yy} = 0.021$), full load ($\epsilon_{yy} = 0.042$), and midway unloading ($\epsilon_{yy} = 0.021$). To interpret these continuum-scale fields, the hybrid region is partitioned into two overlap layers and an upper DEM layer (Fig. 7.7), and the contact-level slip statistics in those subregions are contrasted with the reference DEM using the consolidated histograms in Fig. 7.11. Unless stated otherwise, “pure DEM” denotes the DEM reference solution.

For each curve in Figs. 7.8–7.10 we report in the legend a mean percent error (MPE) computed on the set of hybrid sampling points N_H used for plotting. If y_i are those abscissae, H_i the hybrid ordinates, and G_i the DEM ordinates interpolated at y_i , then

$$\text{MPE} = \frac{1}{N_H} \sum_{i=1}^{N_H} \left(\frac{H_i - G_i}{G_i} \right) \times 100\%. \quad (7.1)$$




Figure 7.5: Hybrid DEM–FEM simulations: vertical stress σ_{yy} versus vertical strain ϵ_{yy} for different coefficients of friction. Comparison between Pure DEM and Hybrid models.




Figure 7.6: Hybrid DEM–FEM simulations: lateral stress σ_{xx} versus vertical stress σ_{yy} for different coefficients of friction. Comparison between Pure DEM and Hybrid models.

The MPE is shown in the legend for each plot and is reported separately for σ_{yy} and for K .

Starting with $\text{COF} = 0.4$ (Fig. 7.8), the hybrid model reproduces $\sigma_{yy}(y)$ closely at midway loading (Fig. 7.8a), while $K(y)$ is modestly higher than in DEM (Fig. 7.8b), consistent with the steeper stress-path slopes seen earlier. This behaviour is expected: the elasto-plastic continuum parameters were calibrated to the full hysteretic loop rather than to a single stage, so small stage-wise deviations are acceptable. At full load the hybrid slightly underestimates σ_{yy} relative to DEM, while K remains higher; during midway unloading, σ_{yy} largely recovers but the trend in K reverses and the hybrid lateral ratio falls below the DEM profile through out the vertical distance. This is consistent with irreversible microslip accumulated in DEM that is not captured exactly by the continuum plasticity; the corresponding histograms for $\text{COF} = 0.4$ (Fig. 7.11 a) show a larger fraction of slipping contacts in DEM than in the hybrid model, particularly in the high-mobilisation tail $\eta/\text{COF} \in [0.9, 1]$. Within the hybrid model, the fraction of slipping contacts is relatively uniform across Hybrid zone 1, Hybrid zone 2, and the DEM zone, yet K increases towards the top part of the RVE during unloading (Fig. 7.8b); this vertical gradient aligns with the upward increase of coupling weights—and hence the DEM contribution—from the FEM side to the DEM side, although the K in the DEM zone remains below the Pure DEM value. Across unloading, a larger fraction of contacts enters the no-slip bin; this trend recurs for all COFs (see panels a, b, and c in Fig. 7.11).

Proceeding to $\text{COF} = 0.8$ (Fig. 7.9), the overall picture mirrors the $\text{COF} = 0.4$ case but with a notable difference in the high-slip tail during unloading. For normalised mobilisation $\eta/\text{COF} \in [0.9, 1.0]$, the slip fraction ranks *Hybrid zone 1* > *Hybrid zone 2* > *DEM layer* in Fig. 7.11 b. This elevated slip activity within the overlap brings the midway-unloading $K(y)$ of the hybrid (Fig. 7.9b) closer to the DEM profile than in the $\text{COF} = 0.4$ scenario. A plausible mechanism is enhanced recovery of sliding with increasing elevation as the top platen retracts, which promotes larger local slip activity in the upper hybrid layers for higher friction.

For very high friction, $\text{COF} = 10$ (Fig. 7.10), the hybrid and DEM agree closely at all stages, reflecting weak plasticity and a minimal hysteresis loop. The microslip distributions (Fig. 7.11 c) likewise match well between hybrid model and Pure DEM simulation; for clarity the abscissa is truncated at $\eta/\text{COF} \leq 0.2$ because only a small fraction of contacts mobilises beyond this level.

Two ingredients govern these trends. First, contact-level slip controls lateral accommodation: increases in the fraction of highly mobilised contacts—especially in the right-hand tail of η/COF —enhance lateral compliance, and, on unloading, preserve residual lateral stress produced during compression. Configurations with more slip therefore tend to exhibit higher K during unloading, as seen when contrasting the hybrid and DEM at $\text{COF} = 0.4$ and, to a lesser degree, at $\text{COF} = 0.8$. Second, the vertical gradient of coupling weights in the overlap modulates both vertical stress transfer and lateral confinement, producing a systematic increase of K with height during unloading. However, unless the slip activity within the hybrid model matches that of DEM, the hybrid model K remains below the Pure DEM curve during unloading for $\text{COF} = 0.4$ and 0.8 .

Synthesising across the three frictions, differences in K between the hybrid model and Pure DEM are traced to the differences in slip activity: at $\text{COF} = 0.4$ the larger irreversible microslip in DEM yields a higher unloading K than in the hybrid; at $\text{COF} = 0.8$ the reduction in slip activity narrows this gap; at $\text{COF} = 10$ limited slip and hysteresis lead to near-coincident responses. These observations are consistent with the calibration strategy (fit to the full cycle) and with the intended role of the coupling weights: to transition the response smoothly from continuum-dominated to DEM-dominated region while preserving stress continuity across the overlap. Finally, in all cases the fraction of contacts in the no-slip partition increases during unloading, indicating progressive elastic re-engagement after peak compression; the magnitude and vertical distribution of this recovery depend on COF and on the local DEM/FEM weighting, explaining the gentle K increase toward the DEM side in the hybrid plots and the uniformly higher K in Pure DEM when irreversible slip is substantial. These discrepancies underscore the need for a continuum constitutive description that faithfully captures contact-scale plasticity and path dependence; otherwise, the unloading evolution of K and its vertical gradient cannot be reproduced quantitatively.






Figure 7.7: Domain of different zones in the hybrid model which are split into Hybrid zone 1 , Hybrid zone 2, and DEM zone based on the vertical y-coordinates.




(a) Vertical stress vs. vertical distance

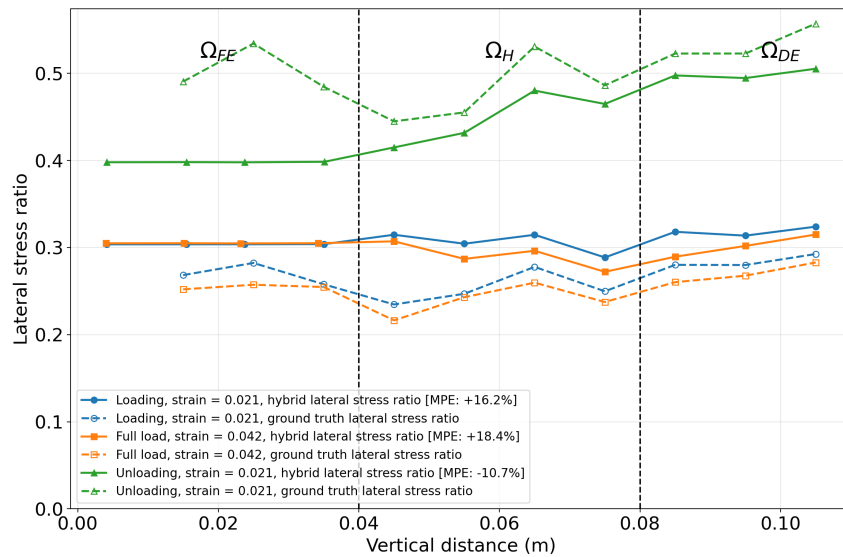


(b) Lateral stress ratio vs. vertical distance

Figure 7.8: Stress distributions in the hybrid model compared to ground truth DEM at three different loading levels Midway loading (strain = 0.021) , Full load (strain = 0.042), and Midway unloading (strain = 0.021) for COF = 0.4.




(a) Vertical stress vs. vertical distance




(b) Lateral stress ratio vs. vertical distance

Figure 7.9: Stress distributions in the hybrid model compared to ground truth DEM at three different loading levels Midway loading (strain = 0.021) , Full load (strain = 0.042), and Midway unloading (strain = 0.021) for COF = 0.8.




(a) Vertical stress vs. vertical distance




(b) Lateral stress ratio vs. vertical distance


Figure 7.10: Stress distributions in the hybrid model compared to ground truth DEM at three different loading levels Midway loading (strain = 0.021) , Full load (strain = 0.042), and Midway unloading (strain = 0.021) for COF = 10.



(a) COF = 0.4: microslip statistics across Hybrid zone 1, Hybrid zone 2, DEM zone, and pure DEM at midway loading, full load, and midway unloading.



(b) COF = 0.8: microslip statistics across Hybrid zone 1, Hybrid zone 2, DEM zone, and pure DEM at midway loading, full load, and midway unloading.



(c) COF = 10: microslip statistics across Hybrid zone 1, Hybrid zone 2, DEM zone, and pure DEM at midway loading, full load, and midway unloading (abscissa truncated to $\eta/COF \leq 0.2$ to capture the majority of particle contacts and to enhance visual clarity).

Figure 7.11: Consolidated micro slip statistics (η/COF) for the three friction cases. Each panel compares the hybrid subzones (Hybrid zone 1, Hybrid zone 2, and the upper DEM layer) against pure DEM at the three stages of the loading cycle.

7.7 Conclusions

This chapter presented an elasto–plastic hybrid DEM–FEM framework in which the continuum region is governed by a Drucker–Prager (DP) law calibrated to DEM responses. Elastic constants (E, ν) were inferred from DEM compression, and the plastic angles (ϕ, χ) were identified by iterative Bayesian filtering, yielding COF-dependent parameter sets that were used unchanged in pure-FEM and hybrid simulations. Across three friction levels (COF = 0.4, 0.8, 10), the hybrid model reproduced the DEM stress paths under cyclic loading with high fidelity (typical $R^2 > 0.95$ against DEM), captured the expected stiffening of axial response with COF, and tracked the lateral–to–vertical stress ratio trends during loading and unloading.

Micro–slip statistics clarified the mechanisms behind residual differences in the unloading branches: departures in the high- η /COF slip tail correlate with modest gaps in the lateral stress ratio K , while very high friction (COF = 10) shows close agreement owing to limited hysteresis. Within the overlap, smooth weights and a sufficiently large penalty produced stable transfers of stress and strain, and the DEM \leftrightarrow FEM mappings preserved action–reaction without spurious energy. Taken together, the results indicate that the volume–coupled hybrid with a calibrated DP continuum provides an accurate and robust surrogate for frictional granular assemblies across a range of particle frictions—substantially improving over pure FEM on cyclic paths.

The principal limitations stem from the simplicity of the local DP model in reproducing subtle hardening and path dependence. Future improvements may include enriched (e.g., nonlocal or dilatancy-evolving) constitutive laws to further reduce the discrepancies of the model with pure DEM.

Chapter 8

Silo Discharge of a polydisperse granular system

8.1 Introduction

Granular discharge from silos is a canonical benchmark for validating discrete continuum hybrids because it exercises the full spectrum of granular behaviour: a rapidly sheared funnel-flow core, quasi-static wall zones with stress arching, and transient stress reorganisation as the free surface descends [1, 9, 11, 63–67]. Particle polydispersity and wall friction amplify these effects and challenge continuum closures, making silo flow a stringent test of whether a hybrid DEM–FEM framework can deliver DEM-level fidelity where it matters while retaining continuum-level efficiency elsewhere.

This chapter deploys the hybrid coupling developed in the preceding chapters on an industrial-type, gravity-driven discharge of a polydisperse assembly and benchmarks it against a high-fidelity pure DEM reference under identical geometry, material parameters, and orifice conditions. The goals are fourfold: (i) to assess how well the hybrid reproduces the evolution of vertical and lateral Cauchy stresses, σ_{yy} and σ_{xx} , across the silo; (ii) to evaluate the lateral stress ratio $k = \sigma_{xx}/\sigma_{yy}$ as a compact diagnostic of continuum response; (iii) to compare wall normal and shear stresses along the confining silo walls during discharge; and (iv) to quantify the mismatch between the two models using consistent space–time error measures on σ_{yy} , σ_{xx} , k , and wall stresses.

8.2 Problem setup and domain decomposition

8.2.1 Geometry and modelling context

An industrial-type, gravity-driven discharge of a polydisperse granular assembly is considered in a cuboidal silo with dimensions $L_x = 1$ m, $L_y = 1.04$ m, and $L_z = 0.015$ m, as illustrated in Fig. 8.2. Two models are run under *identical* geometry, wall conditions, and time stepping: (i) a *pure DEM* reference and (ii) a *hybrid* DEM–FEM model in which a quasi-static bulk is replaced by a continuum region coupled through an overlap. The particle size distribution used in the silo matches the RVE (COF = 0.8) employed in the elasto–plastic study (see Chapter 7); it exhibits a close-to-Gaussian spread, and the minimum, maximum, and mean diameters are in tight agreement with the RVE statistics (Fig. 8.1). The complete list of silo particle properties is shown in Table 8.1.




Figure 8.1: Particle size distributions: silo vs RVE. Bin width = 1.22×10^{-4} m.

The continuum (FEM) material in the silo adopts the elasto–plastic Drucker–Prager parameters calibrated for the COF = 0.8 RVE case (Table 7.3). A small density offset exists between the RVE study and the silo setup: $\rho_{\text{RVE}} = 1539.51 \text{ kg m}^{-3}$ (Table 7.2) versus $\rho_{\text{silo}} = 1478.7 \text{ kg m}^{-3}$ (Table 8.2). Taking the silo pure-DEM bulk as reference solution, the relative difference is 4.11%. This small deviation is retained to remain faithful to the RVE calibration while preserving the target silo mass.

Table 8.1: DEM specifications for silo particles

Specification	Value
Avg. diameter of particle	0.0043 m
Range - particle diameters	0.0019 m to 0.0078 m
Particle density	2500 kg/m ³
K _{normal} Particle-Particle	2000 N/m
K _{tangential} Particle-Particle	2000 N/m
Coef. of Restitution	0.1
Coef. of Friction	0.8
Coef. of Rolling Friction	0.02
Timestep	1 × 10 ⁻⁵ s
No. of particles (pure DEM)	194040
COF Particle-Wall	COF _z = 0, COF _x = 0.8, COF _y = 0.8
COR Particle-Wall	0.1
Coef. of Rolling Friction Particle-Wall	0.02
K _{normal} Particle-Wall	4000 N/m
K _{tangential} Particle-Wall	4000 N/m

Table 8.2: FEM and coupling specifications for silo discharge problem

Specification	Value
Solver type	Dynamic
Rayleigh alpha damping	20 s ⁻¹
Rayleigh beta damping	0.05 s
Density	1478.7 kg/m ³
Analysis type	Non-linear
Timestep	1 × 10 ⁻⁵ s
Mesh size	0.02 m
Penalty (ϵ)	10 ⁷ Pa/m ²
Weighting	Linear (0 to 1)

8.2.2 Particle generation, stress-free initialisation, and settling protocol

Particles are generated to fill the silo up to a height of 1.04 m with no pre-stresses (see Fig. 8.2a). From this stress-free configuration, gravity is applied to the DEM particles and ramped over many computational timesteps to minimize any transient dynamics. The assembly settles by approximately 3% settlement (from 1.04 m to ≈ 1.01 m), establishing the *settled silo* baseline. The *same* two-stage procedure is used for the hybrid model: (a) the DEM region is initialised stress-free and (b) the FEM region is also started without pre-stresses; gravity is then applied to DEM and FEM *at the same pace and with the same time step*, ensuring compatible development of self-weight stresses across the coupled system. Following settlement, the outlet is opened at $t=0$ and discharge proceeds to $t=0.50$ s, with diagnostics reported at

$$t \in \{0, 0.05, 0.10, 0.20, 0.40, 0.50\} \text{ s.}$$


8.2.3 Hybrid overlap topology and weighting

In Fig. 8.2b, the hybrid overlap (between DEM and FEM) is four elements thick and *involves a 90° bend*, positioned to capture the quasi-static core as faithfully as possible within the given geometry. Blending is achieved by assigning node-layer weights across five consecutive layers spanning the overlap, from the FEM side to the DEM side. Nodes are grouped into layers that sandwich the intervening elements, and each layer is given a distinct weight as shown in Table 8.3.

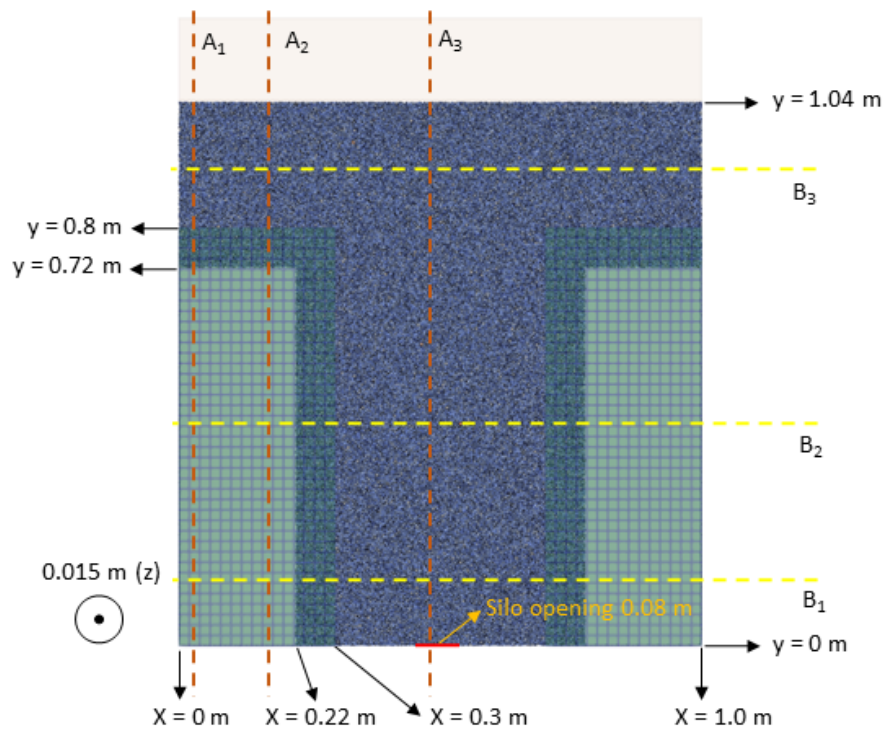
Table 8.3: Node-layer weighting scheme across the hybrid FEM–DEM overlap region

Layer index	1	2	3	4	5
Weight w	0	0.25	0.50	0.75	1.00

These nodal values generate a *continuous weighting field* $w(\mathbf{x})$ *inside* each element via standard finite-element interpolation; for particles located within the overlap, the nodal weights are likewise *interpolated to particle positions using shape functions*, so that DEM–FEM blending is smooth for both field equations and particle kinematics. Kinematic compatibility across the overlap is enforced weakly through a displacement-penalty term of magnitude ε (Table 8.2). DEM and FEM advance with the same Δt to avoid subcycling artefacts. The coupling is realised in KRATOS (Chapter 4) and adheres to the DEM–FEM framework set out in Chapter 3.



(a) pure DEM silo domain.



(b) Hybrid silo domain. A symmetric finite-element region replaces the quasi-static DEM bulk; the hybrid-zone boundaries are indicated. For subsequent analysis, vertical cuts A_1 , A_2 , and A_3 are taken at $x = 0.025$, 0.175 , and 0.475 m, and lateral cuts B_1 , B_2 , and B_3 at $y = 0.125$, 0.425 , and 0.925 m.

Figure 8.2: Problem setup and domain decomposition. (a) pure DEM configuration. (b) Hybrid configuration obtained by embedding a finite-element region.

8.2.4 Sampling, coarse graining, and wall statistics

To ensure pointwise comparability between the DEM and FEM fields, stresses are evaluated at identical spatial locations in both formulations. Specifically, DEM Cauchy stresses are coarse-grained over a cuboidal kernel of

$$\Delta x = 0.05 \text{ m}, \quad \Delta y = 0.05 \text{ m}, \quad \Delta z = 0.015 \text{ m},$$

and FEM Gauss-point stresses are interpolated onto the *same* query locations. On the confining walls, normal and shear stresses are reported as running averages, sampled at 0.02 m intervals along the vertical coordinate y . Each running average is computed as the wall force divided by the area of a strip 0.02 m tall (along y) and 0.015 m thick in the out-of-plane z direction. To reduce symbol clutter in wall-stress plots, every other marker is suppressed.

8.2.5 Sections and reported fields

Spatial diagnostics are extracted along three vertical cuts,

$$A_1 : x = 0.025 \text{ m}, \quad A_2 : x = 0.175 \text{ m}, \quad A_3 : x = 0.475 \text{ m},$$

and three lateral cuts,

$$B_1 : y = 0.125 \text{ m}, \quad B_2 : y = 0.425 \text{ m}, \quad B_3 : y = 0.925 \text{ m},$$

as indicated in Fig. 8.2. We report the vertical and lateral Cauchy stresses (σ_{yy}, σ_{xx}), the lateral stress ratio $k = \sigma_{xx}/\sigma_{yy}$, and wall normal/shear stresses through the discharge window.

8.2.6 Error measures and presentation

Vertical profiles (along A_1 – A_3) are summarised with a *weighted* mean percent error (WMPE) to emphasise higher-magnitude regions, while lateral profiles (along B_1 – B_3) use the *mean percent error* (MPE) as defined in Eq. (7.1). These measures are applied to σ_{yy} , σ_{xx} , k , and wall stresses at the designated time instants.

Let $\{D_i\}_{i=1}^N$ denote the reference solution (pure DEM) values and $\{H_i\}_{i=1}^N$ the hybrid values sampled at the same locations along a given profile. We define the WMPE in a “ratio of sums” form that compares the *magnitudes* of the fields:

$$\text{WMPE} = \frac{\sum_{i=1}^N (|H_i| - |D_i|)}{\sum_{i=1}^N |D_i|} \times 100\%. \quad (8.1)$$

This sign-preserving definition reports positive percentages when the hybrid *overpredicts* the stress magnitude relative to the Pure DEM (on average over the section) and negative values when it *underpredicts*.

8.3 Settled-state silo comparison

8.3.1 Comparison of stresses between pure FEM and pure DEM models

The simulation is advanced with gravity increased linearly to its target value and then held while numerical damping further dissipates any remaining oscillations, yielding a final settled state. To verify the FE continuum model, an additional pure FEM simulation is performed through the settling phase. Figure 8.3a compares vertical stress σ_{yy} profiles at several longitudinal cross sections between the pure FEM and pure DEM solutions.

The weighted mean percent error (WMPE) improves progressively from A_1 to A_3 . This trend is consistent with the friction law used on the FEM side, $\min(\Delta u K_t, \mu N)$, being largely fully mobilised in the settled state, whereas in the pure DEM model the wall friction is not fully mobilised. A more detailed discussion of wall stresses in the settled state is provided later.

Figure 8.3b shows the lateral stress σ_{xx} at several vertical cross sections. In general, the pure FEM underpredicts the lateral stresses, and both models exhibit similar slopes over 0–0.8 m. In the range 0.8–1.0 m, the pure DEM indicates a high apparent Poisson’s ratio, which is expected near a free surface where the mean stress is small and incipient wall arching and contact reorganisation can elevate σ_h/σ_v . The pure FEM additionally shows an increase in lateral stress σ_{xx} when moving from A_1 to A_3


at the free surface; although not typical for silo statics to start the free surface with a non zero lateral stress, a pseudo-elastic continuum calibration combined with high wall friction can generate bending moment stress at the mid-free surface of the silo without material yielding and produce this trend.

Figure 8.3c shows the lateral stress σ_{zz} at several vertical cross sections. The σ_{zz} profiles show reasonable agreement between models, though pure FEM still tends to underpredict the pure DEM results overall. At section A_1 , close to the wall, σ_{yy} approaches a saturation level, consistent with the Janssen effect [63, 68]. The mobilisation of wall friction diverts a portion of the vertical load into wall shear sustained by the wall-normal pressure, reducing the growth of vertical stress with depth and producing the observed arching. Due to Poisson effects, the arching observed in σ_{yy} manifests in σ_{zz} as a similar non-linear variation with depth. Notably, non-linear variation is not pronounced in σ_{xx} for either model, σ_{xx} keeps increasing with depth, with only a weak near-wall saturation over the last few sampling points. This behaviour is consistent with near-full mobilisation of wall friction at settlement, for which $\tau_w \approx \mu \sigma_{xx}$ constrains the wall-normal stress to the level needed to sustain the wall shear that carries part of the vertical load. The resulting boundary condition limits lateral gradients and suppresses a strong non-linear signature in σ_{xx} , even as fabric anisotropy develops.


8.3.2 Comparison of stresses between Hybrid and pure DEM models

Figures 8.4a, 8.4b, and 8.4c present the vertical and lateral stress distributions along the same vertical cross sections considered previously. The plots indicate the distinct computational regions with labels Ω_{FE} , Ω_H , and Ω_{DE} . Note that the A_3 line lies entirely within Ω_{DE} . Comparing Hybrid to the pure FEM results at A_3 , the Hybrid model attains a lower WMPE, which is consistent with the fact that this cross section samples only the DEM region in the Hybrid simulation.


For the A_1 and A_2 cross sections, the Hybrid profiles between approximately 0.80 and 1.04m closely follow the pure DEM curves because the local domain is Ω_{DE} . As these sections enter the overlap/hybrid region Ω_H , the Hybrid solution transitions smoothly and trends toward the behaviour observed in the pure FEM solution. Across all three figures the coupling produces a continuous stress field without spurious jumps at the domain interfaces.



(a) Comparing vertical stress σ_{yy} .



(b) Comparing lateral stress σ_{xx} .



(c) Comparing lateral stress σ_{zz} .


Figure 8.3: Settled silo comparison: pure FEM vs pure DEM stresses along various vertical cross sections at A1,A2, and A3.

In contrast to the pure FEM vs pure DEM comparisons, Figure 8.4b demonstrates that along the vertical cross-section A_3 , the lateral stress aligns closely with the reference DEM solution. This alignment occurs because, along A_3 , the hybrid domain consists of pure DEM particles rather than a pseudo-elastic continuum, thereby eliminating bending-moment effects near the free surface. Notably, the DEM lateral stress (σ_{xx}) exhibits fluctuations near the walls along the A_1 vertical cross-section, which are absent in both the pure FEM and hybrid models. In a settled state, particles form force chains that result in local variations in lateral stress near boundaries, rather than a monotonic increase with depth. This phenomenon is a direct consequence of the discrete nature of particle-particle interactions, which cannot be captured by the simplified continuum model employed here.


The wall effects observed in the hybrid model follow a similar pattern to those of the pure FEM model, with both differing from the pure DEM results in a comparable manner. This suggests that the coupling framework functions effectively, and the remaining discrepancies between the hybrid and reference DEM results are primarily attributable to the inherent limitations of the continuum representation rather than the coupling strategy itself. Specifically, the FEM sub-domain employs a calibrated pseudo-elastic response combined with an idealised frictional wall law; while this reproduces global stress trends, it cannot fully replicate the local stress redistribution and contact reorganisation observed in the discrete system, particularly in regions of low mean stress near the free surface. The smooth, continuous stress transition across the overlap region indicates that the coupling procedure does not introduce spurious numerical artefacts. Consequently, the residual differences reflect the fidelity of the chosen constitutive description in approximating underlying granular behaviour, highlighting constitutive calibration as a critical factor in the accuracy of hybrid simulations.

8.3.3 Comparison of wall stresses between Pure DEM, Pure FEM, and Hybrid models


In both the hybrid-region Ω_{FE} and the pure FEM model, wall friction is represented by a regularised Coulomb law, $\tau_{xy} = \min(K_t \Delta u_t, \mu \sigma_{xx})$, with K_t taken equal to the particle tangential contact stiffness. Figure 8.5a compares the wall-normal stress σ_{xx} and wall shear τ_{xy} at the $-x$ wall ($x = 0$ m) for the Pure FEM and Pure DEM simulations. In the Pure FEM case, τ_{xy} closely follows the Coulomb limit with friction coefficient $\mu = 0.8$, i.e., $\tau_{xy} \approx \mu \sigma_{xx}$, indicating a near-full mobilisation of wall friction at settlement. In contrast, the Pure DEM wall shear remains well below $\mu \sigma_{xx}$, suggesting that friction



(a) Comparing vertical stress σ_{yy} .



(b) Comparing lateral stress σ_{xx} .




(c) Comparing lateral stress σ_{zz} .


Figure 8.4: Settled silo comparison: Hybrid vs pure DEM stresses along various vertical cross sections at A1,A2, and A3.

is not fully mobilised at the wall. This shortfall is consistent with the filling and settling procedure: residual oscillations and contact rearrangements can intermittently affect the tangential force development during filling, limiting sustained shear mobilisation at the boundary in the DEM settling stage.

Figure 8.5b compares the Hybrid and Pure DEM results. Within Ω_{DE} , the Hybrid wall stresses follow the Pure DEM response; upon entering Ω_{FE} , the Hybrid solution transitions smoothly toward the Pure FEM trend, consistent with continuum behaviour in that region. The principal divergence between Hybrid and Pure DEM appears in the wall shear τ_{xy} . In the FE zone, τ_{xy} is more strongly mobilised, so a larger fraction of the weight is carried by wall shear sustained by the wall-normal pressure, which lowers the vertical stress in the regions near to the wall (Janssen-type effect) [68]. By contrast, in the Pure DEM case the initial mobilisation of τ_{xy} is weaker, leading to higher vertical stresses near the walls in the same quasi-static zones. This disparity is a leading contributor to the calibration mismatch at settlement and it carries through into the subsequent dynamic response. Overall, the deviations discussed here are dominated by the initial shear boundary condition—i.e., the degree of wall-friction mobilisation—rather than by shortcomings of the coupling strategy itself.



(a) Wall stress comparing pure DEM vs pure FEM.



(b) Wall stress comparing Hybrid vs pure DEM.

Figure 8.5: Settled silo: wall normal and wall shear stress comparisons.

8.4 Silo discharge analysis

8.4.1 Analysis along lateral cross sections


Time evolution of vertical stress σ_{yy}

To examine the time evolution of stresses during discharge, lateral cross sections B_1 , B_2 , and B_3 , located in the lower, middle, and upper portions of the silo, respectively, are analysed to characterise the full-domain response. The B_3 cross section lies entirely within the Ω_{DE} region of the Hybrid model.


Figure 8.6a shows vertical stress σ_{yy} along B_1 . At $t = 0$, the profile exhibits a maximum at mid-span and a reduction near the walls, consistent with wall-friction effects discussed earlier. As discharge progresses, the mid-span vertical stress decreases rapidly and reaches a minimum of approximately 1000 Pa around $t = 0.20$ s. The zone of reduced σ_{yy} expands laterally, indicating a widening flowing zone. In the quasi-static Ω_{FE} region, σ_{yy} increases relative to the initial state: as the orifice opens and the core σ_{yy} drops, enhanced shear at the interface with the quasi-static zone must be equilibrated by that region, thereby elevating the local vertical stress and causing a stress redistribution [66]. Although the initial settled-state mismatch in wall-shear mobilisation offsets absolute levels, the Hybrid model captures the observed trends and mechanisms in pure DEM simulation [65].

Figure 8.6b shows a similar evolution along B_2 with a temporal delay relative to B_1 . Here σ_{yy} approaches a saturated level near 1000 Pa by $t \approx 0.40$ s. As discharge initiates, the core unloads while a quasi-static band adjacent to the stagnant peripheral zone takes a larger share of the load, so σ_{yy} in that band can rise relative to the core. This transient compression can momentarily mobilise wall friction and produce brief, mass-flow-like near-wall signatures (i.e., boundary-layer-type stress redistribution). Similar near-wall shear layers are documented under mass-flow conditions [69]. Meanwhile, the principal shear remains concentrated on the internal interface between the flowing core and the stagnant annulus.


Figure 8.6c presents σ_{yy} along the upper cross section B_3 . At $t = 0.05$ s, the Hybrid model's mean percent error (MPE) changes sign—transitioning from underprediction to overprediction and then back—yet the magnitude remains modest, with $|\text{MPE}| \lesssim 10\%$ throughout. Since B_3 lies entirely within Ω_{DE} , the Hybrid prediction follows the DEM response closely, with discrepancies dominated by the initial differences in wall-friction mobilisation rather than by the domain coupling.



(a) σ_{yy} vs x along B_1 .



(b) σ_{yy} vs x along B_2 .




(c) σ_{yy} vs x along B_3 .

Figure 8.6: Time evolution of vertical stress σ_{yy} vs X at three different lateral cross sections - comparison between Hybrid and pure DEM model.


Time evolution of lateral stress σ_{xx}

Figures 8.7a, 8.7b, and 8.7c trace the evolution of the lateral stress σ_{xx} along the three lateral cross sections. In Fig. 8.7a (along lateral cross section B_1), σ_{xx} remains essentially unchanged at $t = 0.05$ s even though σ_{yy} drops markedly at that instant: particle motion has not yet fully developed, so the lateral stress continues to confine the assembly. With continued discharge, σ_{xx} relaxes toward a steady level of ≈ 2000 Pa. The stress evolution as discharge commences follow similar pattern to the vertical stress, where over time, there is a reduction in the lateral stress in the core flow zone with a corresponding increase in the lateral stress in the surrounding quasi-static zones against the walls. There is, however, a distinct difference in the slight time delay in the lateral stress evolution. As discharge initiates, the core unloads while a quasi-static band adjacent to the wall carries a larger share of the load. In this band, mobilisation of wall friction and passive confinement redirect stresses laterally, so σ_{xx} adjacent to the wall can increase even as σ_{yy} in near the wall decreases. This transient near-wall amplification of σ_{xx} produces brief, mass-flow-like boundary-layer signatures; similar near-wall shear layers under mass-flow conditions are documented in [69].


In Fig. 8.7b (B_2), σ_{xx} exhibits modest variation along the section except near the silo edges where the stresses drop due to arching. The stress magnitude shows an overall increase over time. This trend reflects a combination of quasi-static confinement (Poisson-type coupling in regions of higher σ_{yy}) and wall-induced effects. By contrast, Fig. 8.7c (B_3) shows a substantial increase of σ_{xx} at the centre as discharge proceeds: funnel-flow entrains material from above, drawing it toward the core and elevating σ_{xx} centrally, while σ_{xx} decreases in the flanks where material is being pulled inward and rises again at the boundaries due to the wall effects noted above. The hybrid model reproduces the stress *trends* well; although its accuracy is sensitive to the initial mobilisation of wall friction, the overall discrepancy is moderate, with MPE $\approx 20\%$.



(a) σ_{xx} vs x along B_1 .



(b) σ_{xx} vs x along B_2 .

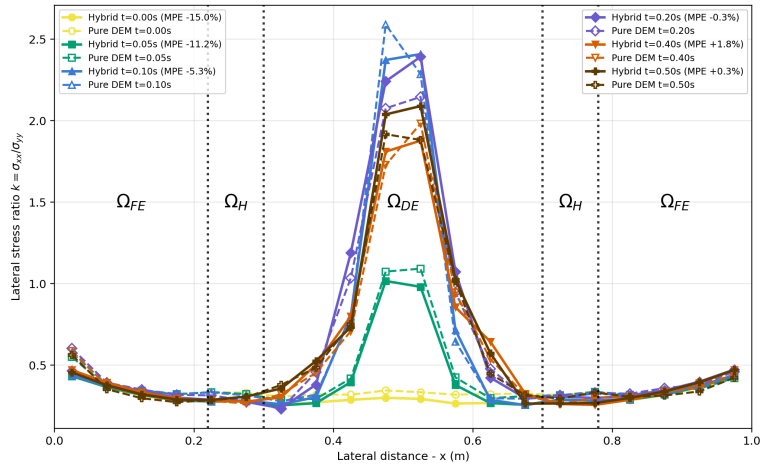


(c) σ_{xx} vs x along B_3 .


Figure 8.7: Time evolution of lateral stress σ_{xx} vs X at three different lateral cross sections - comparison between Hybrid and pure DEM model.

Time evolution of the lateral stress ratio k


Figures 8.8a–8.8c report the time evolution of the lateral stress ratio k along the three lateral cross sections B_1 – B_3 . Near the outlet (Fig. 8.8a, B_1), k rises sharply as discharge begins: σ_{yy} unloads while σ_{xx} remains approximately constant or decreases only slightly. Consequently, k increases from a pre-discharge reference of ≈ 0.25 to peaks around ≈ 2.0 , indicating a local rotation of principal stresses within the developing flow region. At mid-height (Fig. 8.8b, B_2), the growth of k is more gradual as the flow develops, but reaches larger maxima (≈ 3.5) than at B_1 . This amplification is consistent with stronger confinement in the central portion of the silo, where funnel-flow draws material inward and sustains higher lateral support relative to the vertically unloading core. In the upper section (Fig. 8.8c, B_3), the pattern mirrors the behaviour of σ_{xx} : changes in σ_{yy} are comparatively modest, so the evolution of k is dominated by the increase of σ_{xx} as material is pulled toward the core, with lower values on the flanks and a recovery near the walls due to boundary confinement effects.



(a) k vs x along B_1 .



(b) k vs x along B_2 .



(c) k vs x along B_3 .

Figure 8.8: Time evolution of lateral stress ratio k vs X at three different lateral cross sections - comparison between Hybrid and pure DEM model.

8.4.2 Analysis along vertical cross sections


Time evolution of vertical stress σ_{yy}

Figures 8.9a, 8.9b, and 8.9c present the time evolution of σ_{yy} along the vertical sections A_1 – A_3 .


Near the wall (Fig. 8.9a, A_1), the profiles exhibit the expected Janssen-type arching: σ_{yy} tends toward a saturation level with depth. Because the FEM boundary condition has fully mobilised wall friction before the onset of discharge, the arching is more pronounced in the FEM response, whereas the DEM curves approach saturation more gradually since contact friction is only partially mobilised.

A short distance from the wall (Fig. 8.9b, A_2), σ_{yy} varies approximately linearly with depth and its magnitude increases systematically over time. This trend reflects progressive load transfer to the quasi-static regions surrounding the developing flow channel, which support a growing fraction of the overburden as discharge proceeds.


Close to the core (Fig. 8.9c, A_3 ; for the Hybrid case this line lies within Ω_{DE}), σ_{yy} decreases with time and approaches $\sim 10^3$ Pa at the bottom. The associated low-stress front propagates upward through the core, indicating the formation and upward migration of the developing flow region. Consistently, the upper half of the silo shows a concurrent reduction of σ_{yy} with time as the core unloads.



(a) σ_{yy} vs y along A_1 .



(b) σ_{yy} vs y along A_2 .



(c) σ_{yy} vs y along A_3 .

Figure 8.9: Time evolution of vertical stress σ_{yy} vs y at three different vertical cross sections - comparison between Hybrid and pure DEM model.


Time evolution of lateral stress σ_{xx}

Figures 8.10a–8.10c summarise the evolution of σ_{xx} along the vertical sections A_1 – A_3 .


Near the wall (Fig. 8.10a, A_1), the Hybrid model underpredicts σ_{xx} relative to the pure DEM at early times. This bias originates from the settled state: wall friction in FEM is fully mobilised from the outset, whereas the DEM wall friction is only partially mobilised. The resulting offset persists into the discharge phase. A local stress spike appears close to the Ω_{DE} – Ω_H interface; given the proximity to the boundary, enhanced wall adherence combined with the displacement-penalty coupling (which reacts to DEM–FEM kinematic mismatch) plausibly amplifies σ_{xx} locally. Despite these offsets, the Hybrid solution reproduces the *temporal trends* in σ_{xx} as discharge progresses.

At A_2 (Fig. 8.10b), σ_{xx} increases approximately linearly with depth, with a mild tendency toward saturation near the base as the flow develops. The overall magnitude of σ_{xx} grows with time, reflecting progressive load transfer and lateral redirection as the core unloads. In the upper portion—coincident with Ω_{DE} —the Hybrid follows the pure DEM; upon entering Ω_H and subsequently Ω_{FE} , it transitions smoothly toward the continuum response.


Along A_3 in the core flow zone (Fig. 8.10c; Ω_{DE} for the Hybrid case), σ_{xx} is nearly linear with depth at settlement. As discharge proceeds, σ_{xx} increases in the upper zone and decreases in the lower zone, consistent with increasing lateral confinement in the feeding region of the channel. The bottom portion simultaneously relaxes to lower σ_{xx} once the outflow becomes fully developed. By $t = 0.5$ s, the upper half exhibits an approximately uniform σ_{xx} from mid-height to just below the free surface, while the peak magnitude diminishes as it propagates upward. The Hybrid model captures this spatio-temporal redistribution with a WMPE of about -10% .



(a) σ_{xx} vs y along A_1 .



(b) σ_{xx} vs y along A_2 .




(c) σ_{xx} vs y along A_3 .


Figure 8.10: Time evolution of vertical stress σ_{xx} vs y at three different vertical cross sections - comparison between Hybrid and pure DEM model.

Time evolution of the lateral stress ratio k


Figures 8.11a–8.11c show the evolution of k along the vertical sections A_1 – A_3 . For A_1 (Fig. 8.11a), the Hybrid response within Ω_{FE} shows a slight decrease of k with height before entering Ω_H ; once in Ω_{DE} , the curve converges toward the pure DEM trend. Near the free surface and close to the wall, k is initially elevated—consistent with the lateral stress acting as the major principal stress—but it diminishes near the wall as discharge draws material into the core. The cross-stream behaviour of k near the free surface discussed earlier is consistent with these vertical profiles (see Sec. 8.4.1). At A_2 (Fig. 8.11b), the Hybrid model closely matches the pure DEM distribution, with WMPE typically within $\pm 10\%$ over the reported times. For A_3 (Fig. 8.11c; Ω_{DE} in the Hybrid case), a central peak in k develops and migrates upward as discharge advances, reaching its largest value at $t = 0.5$ s. The Hybrid solution reproduces both the upward migration and peak formation, with a slight overall overprediction in magnitude.



(a) k vs X along A_1 .



(b) k vs X along A_2 .




(c) k vs X along A_3 .


Figure 8.11: Time evolution of lateral stress ratio k vs y at three different vertical cross sections - comparison between Hybrid and pure DEM model.

8.4.3 Analysis of wall stresses

Figure 8.12a presents the time evolution of the wall-normal stress and mirrors the behaviour of σ_{xx} along A_1 (see Sec. 8.4.2). Figure 8.12b shows the corresponding wall shear τ_{xy} . In the pure DEM, wall friction is not fully mobilised at the onset of discharge and mobilises progressively thereafter; this initial mismatch introduces an offset in the quasi-static Ω_{FE} region, leading to underprediction relative to the hybrid and FEM trends. Despite this offset, the Hybrid model reproduces the temporal increase in wall-normal stress and the concomitant rise in wall shear, capturing the overall evolution of the wall pressure development on the silo walls during discharge.



(a) Wall normal stress on the $-x$ wall.



(b) Wall shear stress on the $-x$ wall.

Figure 8.12: Time evolution of wall stress (normal and shear) during discharge - comparison between Hybrid and pure DEM model.

8.5 Conclusions

The hybrid DEM–FEM framework reproduced the key features of industrial silo discharge with compact, quantitative fidelity. Across lateral sections (B_1 – B_3) and vertical sections (A_1 – A_3), the Hybrid tracked the unloading of the core in σ_{yy} , the near–wall build-up and redistribution of σ_{xx} , the migration of high- k regions, and the coupled evolution of wall-normal and wall-shear stresses, while maintaining interface continuity without spurious jumps. Despite an initial bias from unequal wall-friction mobilisation at settlement (fully mobilised in FEM, only partial in DEM), errors remained modest—MPE for lateral profiles and WMPE for vertical profiles were both typically within $\sim 10\%$.

These results highlight the value of placing DEM where discrete physics dominate and using a calibrated continuum elsewhere: the Hybrid attains DEM-level fidelity in the flow core and near the free surface, yet preserves continuum-level efficiency in quasi-static zones. By contrast, a standalone FEM would struggle with core shear localisation, rapid stress reorganisation, and variations of $k = \sigma_{xx}/\sigma_{yy}$ tied to frictional redirection. Minor refinements—harmonising wall-friction mobilisation during settlement is expected to trim the residual bias further, Likewise, adopting a better constitutive model that permits spatial variations in plasticity could capture the complex trends with higher fidelity, but the present study already provides a concise, defensible demonstration that DEM–FEM coupling can capture discharge physics beyond the reach of a pure continuum model at practical cost.

Chapter 9

Conclusions

9.1 Conclusions

This thesis has developed and rigorously assessed a *volume-coupled* DEM–FEM framework that allocates discrete resolution where particle physics dominate and deploys a calibrated continuum elsewhere, aiming to achieve DEM–level fidelity at tractable cost. While a full wall–clock accounting of computational savings was beyond scope, the coupling itself adds no appreciable per–step overhead beyond the intrinsic costs of each solver operating in its own subdomain; in successful configurations, large quasi–static regions can be modelled by FEM with DEM reserved for dynamically active zones, offering substantial potential for speedups in industrial settings. A central contribution of this work is to underpin that promise with *quantitative* validation on large–scale problems, reporting error metrics that give users defensible confidence—an aspect under–emphasised in much prior literature.

The formulation is energy–consistent: smooth spatial weights partition mass and internal work across the overlap, an *unweighted* penalty enforces kinematic compatibility, and conservative DEM \leftrightarrow FEM mappings preserve action–reaction without spurious energy. Implemented in KRATOS MULTIPHYSICS as a dedicated application integrated via *CoSimulation*, the method is minimally intrusive to existing solvers and is *constitutive–model agnostic*: beyond linear elasticity and Drucker–Prager (DP) plasticity demonstrated here, it can extend to bonded–particle models and to stiff materials, where localised physics require high accuracy in DEM while surrounding zones respond at macroscopic scales that a continuum can represent effectively.

Verification on controlled elastic benchmarks—first a monodisperse 1D column, then a polydisperse compaction—demonstrated reliable stress/strain transfer across the overlap when three design choices are respected: (i) a penalty stiffness set by a compact, dimensionally consistent rule,

$$\varepsilon = C \frac{E}{dL_c},$$

(ii) smooth, monotone overlap weights spanning a small number of node layers, with *increased Gauss–point sampling* in FEM elements used to evaluate weight–scaled contributions robustly, and (iii) well–posed weighting functions—explicitly avoiding constant weights (e.g. $\omega \equiv 0.5$) that induce spurious jumps. Recasting results through the scaling collapses disparate test families (mono/polydisperse, mesh/particle ratios, element sizes, particle diameters, and overlap thicknesses) onto narrow trends and reveals a high–fidelity plateau near $C \simeq 10$. Beyond this value, larger ε yields negligible benefit; residual discrepancies are governed primarily by DEM \leftrightarrow continuum calibration (e.g. matching E, ν) rather than insufficient coupling strength.

Extending the continuum subdomain to elasto–plasticity with a calibrated DP law preserved these benefits under cyclic loading. Elastic constants were inferred from DEM compression, while frictional/plastic angles were identified against DEM stress paths using a Bayesian–inference workflow. The hybrid tracked DEM responses with high fidelity—capturing friction–dependent stiffening, the evolution of the lateral stress ratio $K = \sigma_{xx}/\sigma_{yy}$, and the shape of loading–unloading branches—while maintaining stable overlap behaviour. Micro–slip statistics linked small unloading gaps to evolving friction mobilisation across the domain at different load–unload stages, indicating where constitutive enrichment (e.g. dilatancy–evolving or nonlocal models) would most effectively reduce residuals.

Finally, in an industrially relevant silo–discharge benchmark, the hybrid reproduced the key features of a high–resolution reference DEM under identical geometry, wall friction, material parameters, and orifice size. It captured core unloading in σ_{yy} , redistribution of vertical and lateral stresses, migration of high– k regions, and the coupled evolution of wall–normal and wall–shear stresses, while preserving interface continuity. Errors—quantified by MPE for lateral profiles and WMPE for vertical profiles—were typically within $\sim 10\%$. A small early–time bias was traced to unequal wall–friction mobilisation during settlement (fully mobilised in FEM, partial in DEM); harmonising this preparation stage is expected to reduce the bias further. Likewise, adopting constitutive

models that permit stress path dependency would help capture complex trends with higher fidelity. Overall, a carefully tuned volume coupling can deliver DEM-like accuracy where it matters most, while retaining continuum efficiency across quasi-static zones that commonly dominate industrial applications.

9.2 Limitations

Despite its robustness, the framework inherits limitations from the local continuum model and from the transfer strategy. The DP law, while effective and computationally light, cannot fully represent path dependence associated with evolving dilatancy, non-local effects, or fabric anisotropy; these omissions are most visible in subtle unloading asymmetries and shear-band details. DEM-to-FEM and FEM-to-DEM mappings are local and projection-based; performance can degrade under extreme polydispersity, very sharp localisation, or strongly rate-dependent flows. The demonstrations, though representative, focus on silo-type configurations; broader classes—impact problems, rapidly varying free surfaces, and strong fluid-grain coupling—require additional validation. That said, the present model can be used *confidently* in quasi-static zones with minimal local variation. Finally, while the scaling $\varepsilon = CE/(dL_c)$ is simple and transferable, the optimal C may drift modestly with algorithmic choices (e.g. alternative weight shapes or mapping kernels) and should be monitored with standard diagnostics; in practice, choosing $C \approx 10$ consistently placed the penalty in a safe, high-fidelity regime. When the hybrid zone undergoes rapid motion or intense shear, DEM-FEM coupling can challenge continuum convergence and slow the co-simulation; algorithmic improvements and *dynamic calibration* of DEM and FEM parameters during such transients will be valuable to mitigate this effect.

9.3 Challenges faced

Implementing the volume-coupled DEM-FEM framework in KRATOS MULTIPHYSICS presented several practical challenges. Documentation gaps and evolving interfaces in an expert-oriented, community-maintained codebase made onboarding and maintenance difficult; sporadic regressions and elusive runtime errors caused setbacks and extended debugging cycles. Limited restart/checkpoint functionality complicated long, large-scale runs by raising the cost of failure and reducing experimentation throughput.

Unlike conventional co-simulation that merely exchanges forces, the present volume coupling required weighted partitions of mass and internal work that affected element assembly, mapping, contact forces, and reactions on both DEM and FEM sides. This tight entanglement with solver internals introduced solver-specific intricacies that had to be engineered and tested carefully. These experiences informed the final design choices (smooth weights, conservative mappings, and a simple penalty scaling) and highlight the value of stronger documentation, regression tests, and robust restart support for future users.

9.4 Future work

The results suggest clear, tractable extensions. On the constitutive side, nonlocal or gradient-enhanced plasticity and dilatancy-evolving laws could reduce unloading gaps and improve shear-band representation. On coupling, adaptive model switching—expanding or contracting the overlap and relocating DEM/FEM subdomains based on error indicators—would automate resolution placement as flow reorganises. Momentum- and energy-consistent coarse-graining for DEM \rightarrow FEM, paired with conservative back-projection, promises more accurate transfers in highly transient regimes. Uncertainty quantification for C and constitutive parameters, for example via Bayesian calibration from small DEM patches, would provide credible intervals alongside mean predictions. Multi-physics couplings (thermo-mechanical heating in shear bands, gas-grain interactions in aerated hoppers) and full-scale 3D industrial systems on heterogeneous hardware are natural proving grounds for scalability. In addition, adopting *different time steps* for FEM and DEM within a synchronised multi-rate co-simulation could further enhance performance without compromising stability.

9.5 Closing statement

In sum, this work delivers an energy-consistent, open-source DEM-FEM volume coupling with a practical penalty scaling law, verified benchmarks, an elasto-plastic extension, and an industrial-type problem validation. The framework consolidates best practices into actionable defaults—smooth weights over a handful of element-layers in the coupling zone, $C \simeq 10$ in $\varepsilon = CE/(dL_c)$, calibrated continuum properties,

increased Gauss–point sampling where weights are applied, and modest FEM refinement where gradients are steep—yielding reproducible, high–fidelity predictions at tractable cost. By enabling granular analysts to put DEM where physics demand it and FEM where efficiency allows it, the thesis provides a coherent path from particle–scale understanding to industrial–scale decision making.

Appendix A

Derivation of volume coupling terms equation.

The additional virtual work δC done by the kinematic constraint in the volume coupling domain Ω^H is given by

$$\delta C = \varepsilon \int_{\Omega^H} (\mathbf{u}^{\text{DE}} - \mathbf{u}^{\text{FE}}) \cdot (\delta \mathbf{x} - \delta \mathbf{X}) \forall \quad (\text{A.1})$$

$$= \underbrace{\varepsilon \int_{\Omega^H} (\mathbf{u}^{\text{DE}} - \mathbf{u}^{\text{FE}}) \cdot \delta \mathbf{x} \forall}_{:=\delta C^{\text{DE}}} + \underbrace{\varepsilon \int_{\Omega^H} (\mathbf{u}^{\text{FE}} - \mathbf{u}^{\text{DE}}) \cdot \delta \mathbf{X} \forall}_{:=\delta C^{\text{FE}}}, \quad (\text{A.2})$$

where \mathbf{u}^{DE} and \mathbf{u}^{FE} are the displacement fields coarse-grained from the particles and interpolated via the finite element basis functions ψ_i ($i = 1, \dots, N$), respectively, $\delta \mathbf{x}$ and $\delta \mathbf{X}$ denote the variation of position vector at the microscopic and macroscopic scales, and ε is the penalty parameter.

Substituting $\mathbf{u}^{\text{FE}} = \sum_{j=1}^N \mathbf{u}_j^{\text{FE}} \psi_j$, $\mathbf{u}^{\text{DE}} = \sum_{j=1}^N \mathbf{u}_j^{\text{DE}} \psi_j$, $\delta \mathbf{X} = \sum_{i=1}^N \delta \mathbf{X}_i \psi_i$, and $\delta \mathbf{x} = \sum_{i=1}^N \delta \mathbf{x}_i \psi_i$, the discretised forms of δC become:

$$\delta C^{\text{FE}} = \int_{\Omega^H} \underbrace{\varepsilon (\mathbf{u}^{\text{FE}} - \mathbf{u}^{\text{DE}})}_{:= -\mathbf{b}^{\text{C}}} \cdot \delta \mathbf{X} \forall \quad (\text{A.3})$$

$$= \sum_{i=1}^N \varepsilon \underbrace{\sum_{j=1}^N \int_{\Omega^H} \psi_i \psi_j \forall (\mathbf{u}_j^{\text{FE}} - \mathbf{u}_j^{\text{DE}}) \cdot \delta \mathbf{X}_i}_{:= -\mathbf{f}_i^{\text{C}}} \quad (\text{A.4})$$

$$\begin{aligned} \delta C^{\text{DE}} &= \varepsilon \int_{\Omega^H} (\mathbf{u}^{\text{DE}} - \mathbf{u}^{\text{FE}}) \cdot \delta \mathbf{x} \forall \\ &= \varepsilon \sum_{i=1}^N \sum_{j=1}^N \int_{\Omega^H} \psi_i \psi_j \forall (\mathbf{u}_j^{\text{DE}} - \mathbf{u}_j^{\text{FE}}) \cdot \delta \mathbf{x}_i \end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^N \mathbf{f}_i^{\mathbf{C}} \cdot \delta \mathbf{x}_i \\
&= \sum_{\alpha=1}^{N_p} \underbrace{\sum_{i=1}^N \Pi_{i\alpha} \mathbf{f}_i^{\mathbf{C}}}_{:= -\mathbf{f}_\alpha^{\mathbf{C}}} \cdot \delta \mathbf{x}_\alpha.
\end{aligned} \tag{A.5}$$

where the homogenisation operator Π (see Eq. (3.35)) is used to relate the virtual displacements from the finite elements to the particles via $\delta \mathbf{x}_i = \sum_{\alpha=1}^{N_p} \Pi_{i\alpha} \delta \mathbf{x}_\alpha$.

Substituting Eqs. (A.3) and (A.5) into the virtual work equation yields

$$\begin{aligned}
\delta W &= \delta W^{\text{FE}} + \delta C^{\text{FE}} + \delta W^{\text{DE}} + \delta C^{\text{DE}} \\
&= \int_{\Omega^{\text{FE}}} (1-w) \left\{ \boldsymbol{\sigma} : \delta \boldsymbol{\varepsilon} - \left(\mathbf{b} + \frac{1}{1-w} \mathbf{b}^{\mathbf{C}} - \rho \frac{\partial^2 \mathbf{X}}{\partial t^2} \right) \cdot \delta \mathbf{X} \right\} \mathcal{V} + \int_{\Gamma_X^{\text{FE}}} \mathbf{t} \cdot \delta \mathbf{X} \mathcal{A} \\
&\quad + \sum_{\alpha=1}^{N_p} \left[w_\alpha \left(m_\alpha \frac{2\mathbf{x}_\alpha}{t^2} - \mathbf{f}_\alpha^{\mathbf{b}} - \frac{1}{w_\alpha} \mathbf{f}_\alpha^{\mathbf{C}} \right) \cdot \delta \mathbf{x}_\alpha - \sum_{\beta}^{N_\alpha} w_{\alpha\beta} \mathbf{f}_{\alpha\beta} \cdot \delta \mathbf{x}_\alpha \right],
\end{aligned} \tag{A.6}$$

where the macroscopic and the microscopic variables are defined for the governing equations of FEM and DEM. Note that the additional coupling terms due to the kinematic constraint closely resemble that of a body force (density). Further, since the homogenisation operators involve a normalisation $\sum_{\alpha=1}^{N_p} \Pi_{i\alpha} = 1$, the coupling forces applied to the FEM nodes is equal and opposite to the coupling forces applied to the particles,

$$-\sum_{\alpha=1}^{N_p} \mathbf{f}_\alpha^{\mathbf{C}} = \sum_{\alpha=1}^{N_p} \sum_{i=1}^N \Pi_{i\alpha} \mathbf{f}_i^{\mathbf{C}} = \sum_{i=1}^N \mathbf{f}_i^{\mathbf{C}}. \tag{A.7}$$

Appendix B

Algorithms for implementation

Algorithms in this Appendix

- i) Algorithm 1: Initialize
- ii) Algorithm 2: Volume Coupling (SolveSolutionStep)
- iii) Algorithm 3: CalculateLinearMomentum()
- iv) Algorithm 4: DEM to FEM Mapping
- v) Algorithm 5: Calculation of nodal coupling forces
- vi) Algorithm 6: Particle coupling forces

Algorithm 1 Initialize

Import Applications

- i) Co-simulation application as `CoSimulation`: To couple FEM and DEM solver.
- ii) DEM application as `dem_analysis`: For simulating discrete particles.
- iii) Structural mechanics application as `structural_analysis`: For continuum granular flow.
- iv) Mapping application as `python_mapper_factory`: To allow mapping between DEM and FEM solvers.
- v) Import `AnalysisStage`: Class that acts as a parent class for all the solvers; contains the main `Run()` function.

Problem Setup

- i) `AnalysisStage.Run()` : Starts the simulation by initializing `CoSimulation` application and both DEM and FEM solver applications.
 - ii) `parameters = KratosMultiphysics.Parameters(parameter_file.read())` : Reads the parameters required for the co-simulation application; this includes mapper configurations and the types of solvers used; for the coupling process we use a *volume_coupling* solver.
 - iii) `CoSimulation._GetSolver()` : Creates class `DemFemVolumeCoupledSolver` (subclass of `CoSimulationCoupledSolver`), which controls the coupling via `SolveSolutionStep()`.
 - iv) `CoSimulation._GetAnalysisStage()` : Returns DEM and FEM `AnalysisStage` objects. DEM and FEM solver parameters are also read when these objects are created.
 - v) `dem_analysis.CreateModelParts()` : Prepares model parts; model parts contain information about the initial configuration of particles/nodes.
 - vi) `dem_analysis.Initialize()` : Initialize all variables.
 - vii) `structural_analysis.Initialize()` : Initialize all variables and the model part.
-

Algorithm 2 Volume Coupling (SolveSolutionStep)

Solution Loop (per time step)

- i) **DEM** → **FEM momentum/mass map and FEM nodal coupling computation**
 - (a) **For** each coupling_op in coupling_operations_dict.values():
 - i. **If** isinstance(coupling_op, ComputeDemMomentum):
 - A. coupling_op.InitializeCouplingIteration()
(compute DEM particle momentum).
 - B. **For** each solver_name, solver in solver_wrappers:
 - C. **If** solver_name == "structure":
_SynchronizeInputData("structure")
(send DEM momentum and mass to FEM; Alg. 4).
 - ii. **Else:**
 - A. coupling_op.InitializeCouplingIteration()
(compute FEM nodal coupling forces; Alg. 5).
 - ii) **FEM** → **DEM inverse map of coupling forces**
 - (a) **For** each solver_name, solver in solver_wrappers:
 - i. _SynchronizeOutputData(solver_name)
(for FEM, map nodal coupling forces back to DEM particles; Alg. 6).
 - iii) **Finalize coupling operations**
 - (a) **For** each coupling_op in coupling_operations_dict.values():
coupling_op.FinalizeCouplingIteration()
(e.g., scale external loads, add loads to selected/bottom particles).
 - iv) **Advance both physics solvers**
 - (a) **For** each solver_name, solver in solver_wrappers:
solver.SolveSolutionStep()
(run DEM and FEM solvers sequentially for this step).
 - v) **Advance in time**
 - (a) CoSimulation._GetSolver().AdvanceInTime() (proceed to the next time step).
-

Algorithm 3 CalculateLinearMomentum()

▷ All variables are defined in DEMApplication; each node represents a particle.

```
for all node ∈ spheres_model_part.Nodes do  
    node.LINEAR_MOMENTUM ← node.NODAL_MASS × node.VELOCITY  
end for
```

Algorithm 4 DEM to FEM Mapping

This describes the mapping process of linear momentum from the DEM solver to the FEM solver. The nearest element to a particle is found according to search-radius settings; once an element is found, the particle data are mapped to the nodes in a conservative way. Here, $\psi_{\alpha_j}^i$ represents the shape function associated with node i evaluated at the position of particle α_j ; the value of the shape function at the particle's position is the weight by which the momentum gets multiplied before assigning it to the i^{th} node. m_{α_j} and v_{α_j} are the mass and velocity of the particles respectively. P_i^{DE} is the homogenised particle momentum at each node i . The mass of the discrete particles also gets mapped in a similar way to m_i . Then the homogenised velocity V_i^{DE} can be calculated as $V_i^{DE} = P_i^{DE} / m_i$.

$$\begin{aligned} P_1^{DE} &\leftarrow \psi_{\alpha_1}^1 \times (m_{\alpha_1} v_{\alpha_1}) + \psi_{\alpha_2}^1 \times (m_{\alpha_2} v_{\alpha_2}) + \dots \\ P_2^{DE} &\leftarrow \psi_{\alpha_1}^2 \times (m_{\alpha_1} v_{\alpha_1}) + \psi_{\alpha_2}^2 \times (m_{\alpha_2} v_{\alpha_2}) + \dots \\ P_3^{DE} &\leftarrow \psi_{\alpha_1}^3 \times (m_{\alpha_1} v_{\alpha_1}) + \psi_{\alpha_2}^3 \times (m_{\alpha_2} v_{\alpha_2}) + \dots \\ P_4^{DE} &\leftarrow \psi_{\alpha_1}^4 \times (m_{\alpha_1} v_{\alpha_1}) + \psi_{\alpha_2}^4 \times (m_{\alpha_2} v_{\alpha_2}) + \dots \end{aligned}$$

Algorithm 5 Calculation of nodal coupling forces

i) Initialize a DEMFEM_VOLUME_COUPLING_FORCE vector to store the nodal coupling forces for each node in the element.

$$\mathbf{F}_l \leftarrow \mathbf{0} \in \mathbb{R}^d, \quad l = 1, \dots, n_e.$$

ii) Calculate the difference between the nodal velocities from the finite element method (FEM) and the projected velocities from the discrete particle method (DEM), multiply with timestep and store cumulatively in the PENALIZE_DISPLACEMENT vector.

$$\Delta \mathbf{u}_l \leftarrow \Delta \mathbf{u}_l + \Delta t (\mathbf{v}_l^{\text{FEM}} - \mathbf{v}_l^{\text{DEM}}), \quad l = 1, \dots, n_e.$$

iii) Set up memory for the shape functions (ψ) and their derivatives ($\frac{\partial \psi}{\partial x_i}$), evaluated at integration points.

$$\psi_l^{(ipt)} := \psi_l(\xi_{ipt}), \quad \left(\frac{\partial \psi_l}{\partial x_i} \right)^{(ipt)} := \frac{\partial \psi_l}{\partial x_i} \Big|_{\xi = \xi_{ipt}}.$$

iv) **For** each integration point ipt of the element:

(a) Retrieve the integration weight w for the current integration point.

$$w \leftarrow w_{ipt}.$$

(b) Calculate the Jacobian J of the shape function derivatives at the current integration point.

$$J \leftarrow \det \left(\frac{\partial \mathbf{x}}{\partial \xi} \right)_{\xi = \xi_{ipt}}.$$

(c) **For** each test function (node) l of the element:

i. **For** each node ll of the element again:

A. Compute the volume vol using the shape functions (ψ_l, ψ_{ll}), integration weight, and Jacobian.

$$vol \leftarrow \gamma w J \psi_l^{(ipt)} \psi_{ll}^{(ipt)}.$$

B. **For** each dimension i of the problem:

• Compute the displacement $displ$ using the difference in nodal velocities and the current time step.

$$displ \leftarrow (\Delta \mathbf{u}_{ll})_i.$$

• Update the nodal coupling forces for the current node l and dimension i using the computed volume, penalty parameter and displacement.

$$F_{l,i} \leftarrow F_{l,i} + vol \cdot displ.$$

Algorithm 6 Particle coupling forces

This describes the inverse mapping process of the coupling forces from FEM nodes to the particles contained within that element. The shape functions used do not change in this case, although the mapping process is not conservative. Here, f_i^C represents the coupling force stored at node i .

$$f_{\alpha_1}^C \leftarrow \psi_{\alpha_1}^1 \times (-f_1^C) + \psi_{\alpha_1}^2 \times (-f_2^C) + \psi_{\alpha_1}^3 \times (-f_3^C) + \psi_{\alpha_1}^4 \times (-f_4^C)$$

$$f_{\alpha_2}^C \leftarrow \psi_{\alpha_2}^1 \times (-f_1^C) + \psi_{\alpha_2}^2 \times (-f_2^C) + \psi_{\alpha_2}^3 \times (-f_3^C) + \psi_{\alpha_2}^4 \times (-f_4^C)$$

...

Appendix C

Core utility functions

```
#ifndef KRATOS_DEMFEM_VOLUME_COUPLING_UTILITIES_H
#define KRATOS_DEMFEM_VOLUME_COUPLING_UTILITIES_H

//--- Kratos core ---
#include "includes/define.h"
#include "includes/model_part.h"
#include "includes/variables.h"

//--- Applications ---
#include "structural_mechanics_application_variables.h"
#include "DEMFEM_volume_coupling_application.h"

//--- STL ---
#include <unordered_map>
#include <functional>
#include <algorithm>
#include <cmath> // std::abs

namespace Kratos {

class DEMFEMVolumeCouplingUtilities
{
public:
KRATOS_CLASS_POINTER_DEFINITION(DEMFEMVolumeCouplingUtilities);

//-----
```

```

// Apply friction-capped tangential point loads on boundary nodes (eg:
//   at x=0/1 or y=0)
// Direction opposes displacement; magnitude = min(Kt*|u|, *|R_n|).
//-----

void AssignPointLoads(
    ModelPart& rFEMModelPart,
    const double cof, // Coulomb friction coefficient
    const double tangential_stiffness, // Kt
    const double coord_tol = 1e-8 // geometric tolerance for boundary
    check
) {
    constexpr double disp_tol = 1e-12;

    for (auto& rNode : rFEMModelPart.Nodes()) {
        const double x = rNode.X();
        const double y = rNode.Y();

        const bool on_bottom = std::abs(y) < coord_tol; // y = 0
        const bool on_left = std::abs(x) < coord_tol; // x = 0
        const bool on_right = std::abs(x - 1.0) < coord_tol; // x = 1

        if (!(on_bottom || on_left || on_right)) continue; // interior /
            top

        auto& u = rNode.FastGetSolutionStepValue(DISPLACEMENT);
        const auto& R = rNode.FastGetSolutionStepValue(REACTION);

        const double u_norm = norm_2(u);
        if (u_norm < disp_tol) { // no motion no load
            rNode.FastGetSolutionStepValue(POINT_LOAD).clear();
            continue;
        }

        const array_1d<double,3> dir_neg = -u / u_norm; // -

        const double f_trial = tangential_stiffness * u_norm; // Kt|u|
        double normal_reaction = 0.0;
    }
}

```

```

        if (on_left || on_right) normal_reaction = std::abs(R[0]); // |
            R_x|
        else if (on_bottom) normal_reaction = std::abs(R[1]); // |R_y|

        const double f_coulomb = cof * normal_reaction; // |R_n|
        const double magnitude = std::min(f_trial, f_coulomb);

        rNode.FastGetSolutionStepValue(POINT_LOAD) = dir_neg * magnitude
            ;
    }
}

//-----

// Assign nodal coupling weights to whole layers by SubModelPart name.
//-----

void SetNodalCouplingWeightsFromLayers(
    ModelPart& rFEMModelPart,
    const std::unordered_map<std::string,double>& rLayerWeights)
{
    // 0) Initialise all nodes to zero
    for (auto& rNode : rFEMModelPart.Nodes()) {
        rNode.FastGetSolutionStepValue(NODAL_COUPLING_WEIGHT) = 0.0;
    }

    // 1) DFS through SubModelPart tree and assign matches
    std::function<void(ModelPart&)> visit;
    visit = [&](ModelPart& rPart) {
        if (auto it = rLayerWeights.find(rPart.Name()); it !=
            rLayerWeights.end()) {
            const double weight = it->second;
            for (auto& rNode : rPart.Nodes()) {
                rNode.FastGetSolutionStepValue(NODAL_COUPLING_WEIGHT) =
                    weight;
            }
        }
    };
    for (auto& rChild : rPart.SubModelParts()) visit(rChild);
}

```

```

};
visit(rFEMModelPart);

// 2) Synchronise ghost nodes (MPI)
if (rFEMModelPart.IsDistributed()) {
    rFEMModelPart.GetCommunicator().SynchronizeVariable(
        NODAL_COUPLING_WEIGHT);
}
}

//-----

// Linear assignment of coupling weights along Y between two boundaries.

// Keeps ONE implementation only (nonlinear/cubic version removed).
//-----

void SetNodalCouplingWeightsOnFEMLinearly(
    ModelPart& rFEMModelPart,
    const double y_fem_boundary,
    const double y_dem_boundary,
    const double tolerance,
    const double weight_fem_boundary,
    const double weight_dem_boundary)
{
    const double denom = (y_dem_boundary - y_fem_boundary);
    if (std::abs(denom) < 1e-15) return; // degenerate interval

    const double y_min = std::min(y_fem_boundary, y_dem_boundary) -
        tolerance;
    const double y_max = std::max(y_fem_boundary, y_dem_boundary) +
        tolerance;

    for (auto& rNode : rFEMModelPart.Nodes()) {
        const double y = rNode.Y();
        if (y < y_min || y > y_max) {
            rNode.FastGetSolutionStepValue(NODAL_COUPLING_WEIGHT) = 0.0;
            continue;
        }
    }
}

```

```

    }
    const double t = (y - y_fem_boundary) / denom; // linear 01
        across the band
    const double w = weight_fem_boundary + (weight_dem_boundary -
        weight_fem_boundary) * t;
    rNode.FastGetSolutionStepValue(NODAL_COUPLING_WEIGHT) = w;
}
}

//-----

// Accumulate displacement difference via velocity mismatch (DEM vs FEM)
.
// PENALIZE_DISPLACEMENT (v_dem - v_fem) dt
//-----

void CalculateDisplacementDifference(ModelPart& rFEMModelPart, const
double& dt)
{
    for (auto& rNode : rFEMModelPart.Nodes()) {
        const double total_mass = rNode.FastGetSolutionStepValue(
            NODAL_MAU);
        if (total_mass == 0.0) continue; // outside hybrid region

        array_1d<double,3> zero = ZeroVector(3);
        rNode.FastGetSolutionStepValue(POINT_LOAD) = zero;
        rNode.FastGetSolutionStepValue(DEMFEM_VOLUME_COUPLING_FORCE) =
            zero;

        const array_1d<double,3> v_dem = rNode.FastGetSolutionStepValue(
            VELOCITY_MULTIPLIED_MASS) / total_mass;
        rNode.FastGetSolutionStepValue(PENALIZE_DISPLACEMENT) += (v_dem -
            rNode.FastGetSolutionStepValue(VELOCITY)) * dt;
    }
}

//-----

```

```

// Assemble nodal coupling forces from element contributions (penalty
    method).
//-----

void CalculateNodalCouplingForces(ModelPart& rFEMModelPart, const
    double& penalty_max)
{
    for (auto itElem = rFEMModelPart.ElementsBegin(); itElem !=
        rFEMModelPart.ElementsEnd(); ++itElem) {
        const auto& rGeom = itElem->GetGeometry();
        if (rGeom[0].FastGetSolutionStepValue(NODAL_MAUX) == 0.0)
            continue; // skip if not in hybrid region

        const auto integration_method = rGeom.
            GetDefaultIntegrationMethod();
        const auto& rGaussPts = rGeom.IntegrationPoints(
            integration_method);
        const Matrix N = rGeom.ShapeFunctionsValues();

        for (unsigned int g = 0; g < rGeom.IntegrationPointsNumber(); ++
            g) {
            const double J = rGeom.DeterminantOfJacobian(g);
            const double w = rGaussPts[g].Weight();

            for (unsigned int n = 0; n < rGeom.size(); ++n) {
                for (unsigned int m = 0; m < rGeom.size(); ++m) {
                    const double vol = penalty_max * w * J * N(g,n) * N(g,
                        m);
                    for (std::size_t d = 0; d < 3; ++d) {
                        const double contrib = vol * rGeom[n].
                            FastGetSolutionStepValue(PENALIZE_DISPLACEMENT
                                ) [d];
                        rGeom[n].FastGetSolutionStepValue(POINT_LOAD) [d]
                            += contrib;
                        rGeom[n].FastGetSolutionStepValue(
                            DEMFEM_VOLUME_COUPLING_FORCE) [d] += contrib;
                    }
                }
            }
        }
}

```

```

        }
    }
}

//-----

// Scale DEM coupling force per node by inverse mass (for mapping to
// DEM).
//-----

void CalculateNodalDEMCouplingForces(ModelPart& rFEMModelPart)
{
    for (auto& rNode : rFEMModelPart.Nodes()) {
        const double total_mass = rNode.FastGetSolutionStepValue(
            NODAL_MAU);
        if (total_mass == 0.0) continue;
        rNode.FastGetSolutionStepValue(DEMFEM_VOLUME_COUPLING_FORCE) =
            -rNode.FastGetSolutionStepValue(DEMFEM_VOLUME_COUPLING_FORCE)
            / total_mass;
    }
}

//----- DEM-side helpers
-----

// Multiply FEM-like kinematics by nodal DEM mass (momentum-style
// packing).
void CalculateMomentum(ModelPart& rDEMModelPart)
{
    for (auto& rNode : rDEMModelPart.Nodes()) {
        rNode.FastGetSolutionStepValue(DISPLACEMENT_MULTIPLIED_MASS) =
            rNode.FastGetSolutionStepValue(NODAL_MASS) * rNode.
                FastGetSolutionStepValue(DISPLACEMENT);
        rNode.FastGetSolutionStepValue(VELOCITY_MULTIPLIED_MASS) =
            rNode.FastGetSolutionStepValue(NODAL_MASS) * rNode.
                FastGetSolutionStepValue(VELOCITY);
    }
}

```

```

}

// Ensure DEM external force is properly scaled (fallback weight=1 if
// unset).
void CalculateDEMForces(ModelPart& rDEMModelPart)
{
    for (auto& rNode : rDEMModelPart.Nodes()) {
        double particle_weight = rNode.FastGetSolutionStepValue(
            PARTICLE_COUPLING_WEIGHT);
        if (particle_weight == 0.0) {
            rNode.FastGetSolutionStepValue(PARTICLE_COUPLING_WEIGHT) =
                1.0;
        } else {
            rNode.FastGetSolutionStepValue(EXTERNAL_APPLIED_FORCE) *=
                rNode.FastGetSolutionStepValue(NODAL_MASS);
        }
    }
};

}; // class DEMFEMVolumeCouplingUtilities

} // namespace Kratos

#endif // KRATOS_DEMFEM_VOLUME_COUPLING_UTILITIES_H

```

Appendix D

Sample co-simulation input scripts

D.1 CoSimulationParameters.json

This script orchestrates the DEM–FEM volume coupling: it selects a Kratos mapper (`nearest_element`) and defines the two key coupling operations—computing DEM momentum and assembling FEM coupling forces with a penalty—plus the full data exchange sequence (masses, packed kinematics to FEM; coupling forces and coupling weights back to DEM).

Listing D.1: Co-simulation driver for DEM–FEM volume coupling.

```
{
  "problem_data": {
    "start_time": 0.0,
    "end_time": 0.2,
    "echo_level": 0,
    "print_colors": true,
    "parallel_type": "OpenMP"
  },
  "solver_settings": {
    "type": "coupled_solvers.volume_coupling", // volume-coupled
      DEM-FEM run
    "echo_level": 0,

    "data_transfer_operators": {
      "mapper_1": {
        "type": "kratos_mapping",
        "mapper_settings": {
```

```

    "mapper_type": "nearest_element", // robust and simple
        mapper
    "search_radius": 0.1,
    "use_approximation": false,
    "search_settings": { }
}
}
},

"coupling_operations": {
    "compute_dem_momentum": {
        "type": "compute_dem_momentum", // packs DEM kinematics ( $m*
            u, m*v$ )
        "echo_level": 0,
        "solver": "dem",
        "model_part_name": "SpheresPart.
            DEMParts_Define_particle_groups_Auto1"
    },
    "compute_nodal_force_process": {
        "type": "compute_dem_fem_volume_coupling_force", // FEM
            penalty assembly
        "echo_level": 0,
        "solver": "structure",
        "model_part_name": "Structure.Parts_Solid_Solid_Auto1",
        "timestep": 1e-5,
        "penalty_max": 1e10, // stiffness of coupling
        "force_end_time": 0.025,
        "boundary_tolerance": 1e-3,
        "y_fem_boundary": 0.16, // linear weight band (FEM side)
        "y_dem_boundary": 0.08, // linear weight band (DEM side)
        "weight_fem_boundary": 0.0, //  $w=0$  at FEM boundary
        "weight_dem_boundary": 1.0 //  $w=1$  at DEM boundary
    }
},

"coupling_sequence": [

```

```

{ "name": "dem", "input_data_list": [], "output_data_list":
  [] },

{ "name": "structure",
  "input_data_list": [
    { "data": "nodalmass",
      "from_solver": "dem", "from_solver_data": "nodalmass_d",
      "data_transfer_operator": "mapper_1",
      "data_transfer_operator_options": ["use_transpose"] },

    { "data": "displacementmultipliedmass",
      "from_solver": "dem", "from_solver_data": "
        displacementmultipliedmass_d",
      "data_transfer_operator": "mapper_1",
      "data_transfer_operator_options": ["use_transpose"] },

    { "data": "velocitymultipliedmass",
      "from_solver": "dem", "from_solver_data": "
        velocitymultipliedmass_d",
      "data_transfer_operator": "mapper_1",
      "data_transfer_operator_options": ["use_transpose"] }
  ],
  "output_data_list": [
    { "data": "couplingforces", // FEM DEM: nodal coupling
      force
      "to_solver": "dem",
      "to_solver_data": "couplingforces_d",
      "data_transfer_operator": "mapper_1" },

    { "data": "nodalcouplingweight", // FEM DEM: weights
      used on DEM side
      "to_solver": "dem",
      "to_solver_data": "particlecouplingweight",
      "data_transfer_operator": "mapper_1" }
  ]
}

```

```

],

"solvers": {
  "dem": {
    "type": "solver_wrappers.kratos.dem_wrapper",
    "solver_wrapper_settings": {
      "input_file": "dem/ProjectParametersDEM",
      "move_mesh_model_part": [],
      "working_directory": "dem"
    },
    "data": {
      "particlecouplingweight": {
        "model_part_name": "SpheresPart.
          DEMParts_Define_particle_groups_Auto1",
        "variable_name": "PARTICLE_COUPLING_WEIGHT"
      },
      "couplingforces_d": {
        "model_part_name": "SpheresPart.
          DEMParts_Define_particle_groups_Auto1",
        "variable_name": "EXTERNAL_APPLIED_FORCE",
        "dimension": 3
      },
      "displacementmultipliedmass_d": {
        "model_part_name": "SpheresPart.
          DEMParts_Define_particle_groups_Auto1",
        "variable_name": "DISPLACEMENT_MULTIPLIED_MASS",
        "dimension": 3
      },
      "velocitymultipliedmass_d": {
        "model_part_name": "SpheresPart.
          DEMParts_Define_particle_groups_Auto1",
        "variable_name": "VELOCITY_MULTIPLIED_MASS",
        "dimension": 3
      },
      "nodalmass_d": {

```

```

        "model_part_name": "SpheresPart.
            DEMParts_Define_particle_groups_Auto1",
        "variable_name": "NODAL_MASS"
    }
}
},

"structure": {
    "type": "solver_wrappers.kratos.
        volume_coupling_structural_wrapper",
    "solver_wrapper_settings": {
        "input_file": "fem/ProjectParametersFEM"
    },
    "data": {
        "displacementmultipliedmass": {
            "model_part_name": "Structure.Parts_Solid_Solid_Auto1",
            "variable_name": "DISPLACEMENT_MULTIPLIED_MASS",
            "dimension": 3
        },
        "velocitymultipliedmass": {
            "model_part_name": "Structure.Parts_Solid_Solid_Auto1",
            "variable_name": "VELOCITY_MULTIPLIED_MASS",
            "dimension": 3
        },
        "nodalcouplingweight": {
            "model_part_name": "Structure.Parts_Solid_Solid_Auto1",
            "variable_name": "NODAL_COUPLING_WEIGHT"
        },
        "nodalmass": {
            "model_part_name": "Structure.Parts_Solid_Solid_Auto1",
            "variable_name": "NODAL_MAUX"
        },
        "couplingforces": {
            "model_part_name": "Structure.Parts_Solid_Solid_Auto1",
            "variable_name": "DEM_FEM_VOLUME_COUPLING_FORCE",
            "dimension": 3
        }
    }
}

```



```
    "strategy": "sphere_strategy",
    "material_import_settings": { "materials_filename": "
        MaterialsDEM.json" }
},
```

```
"VirtualMassCoefficient": 1.0,
"RollingFrictionOption": true,
"GlobalDamping": 0.0,
"ContactMeshOption": true,
```

```
"OutputFileType": "ascii",
"Multifile": "multiple_files",
"ElementType": "SphericPartDEMElement3D",
"TranslationalIntegrationScheme": "Symplectic_Euler",
"RotationalIntegrationScheme": "Direct_Integration",
```

```
"MaxTimeStep": 1e-5,
"FinalTime": 1,
"NeighbourSearchFrequency": 50,
"SearchTolerance": 0.001,
```

```
"GraphExportFreq": 0.001,
"VelTrapGraphExportFreq": 0.001,
"OutputTimeStep": 0.001,
```

```
"PostBoundingBox": true,
"PostLocalContactForce": true,
"PostDisplacement": true,
"PostRadius": true,
"PostVelocity": true,
"PostAngularVelocity": true,
"PostElasticForces": true,
"PostContactForces": true,
"PostRigidElementForces": true,
"PostStressStrainOption": true,
"PostTangentialElasticForces": true,
```

```

"PostTotalForces": true,
"PostPressure": true,
"PostShearStress": true,
"PostSkinSphere": true,
"PostNonDimensionalVolumeWear": false,
"PostParticleMoment": true,
"PostAppliedForces": true, // important: for coupled loads
    arriving as external forces
"PostEulerAngles": true,
"PostRollingResistanceMoment": true,
"post_vtk_option": true,

"problem_name": "particlesettling"
}

```

D.3 ProjectParametersFEM.json

FEM settings for the structural wrapper (implicit dynamic, backward_euler, small time step). ****Important:**** the *point-load entities* are created via `create_point_based_entities_process` with `PointLoadCondition3D1N`. This is crucial: it enables nodal POINT_LOAD conditions so the DEM–FEM coupling forces can be applied directly at nodes. Without these entities, the mapped coupling forces would have nowhere to go.

Listing D.3: FEM solver parameters; note the creation of point-load entities for coupling.

```

{
  "problem_data": {
    "problem_name": "fem_region_longer",
    "parallel_type": "OpenMP",
    "echo_level": 1,
    "start_time": 0.0,
    "end_time": 1
  },
  "solver_settings": {
    "solver_type": "Dynamic",

```

```

"model_part_name": "Structure",
"time_integration_method": "implicit",
"domain_size": 3,
"echo_level": 1,
"analysis_type": "non_linear",
"auxiliary_variables_list": ["NODAL_MAUX"], // used by
    coupling
"scheme_type": "backward_euler",
"rayleigh_alpha": 200,
"rayleigh_beta": 0.01,
"model_import_settings": {
    "input_type": "mdpa",
    "input_filename": "fem_region_longer"
},
"material_import_settings": {
    "materials_filename": "StructuralMaterials.json"
},
"time_stepping": { "time_step": 1e-4 },
"line_search": false,
"convergence_criterion": "residual_criterion",
"displacement_relative_tolerance": 0.0001,
"displacement_absolute_tolerance": 1e-9,
"residual_relative_tolerance": 0.0001,
"residual_absolute_tolerance": 1e-9,
"max_iteration": 10,
"rotation_dofs": false,
"volumetric_strain_dofs": false
},

"processes": {
    "constraints_process_list": [
        {
            "python_module": "fix_vector_variable_process",
            "kratos_module": "KratosMultiphysics",
            "help": "Fix selected components of a vector variable",
            "Parameters": {

```

```

    "model_part_name": "Structure.
      DISPLACEMENT_Displacement_Auto1",
    "variable_name": "DISPLACEMENT",
    "constrained": [true, false, true]
  }
},

{
  "python_module": "create_point_based_entites_process",
  "kratos_module": "KratosMultiphysics.
    CoSimulationApplication.processes",
  "Parameters": {
    "model_part_name": "Structure.
      SelfWeight3D_Self_weight_Auto1",
    "entity_name": "PointLoadCondition3D1N", // creates
      POINT_LOAD entities at nodes
    "entity_type": "condition",
    "properties_id": 0
  }
},

// Prescribed vertical motion on a moving set (loading)
{
  "python_module": "assign_vector_variable_process",
  "kratos_module": "KratosMultiphysics",
  "help": "Assign vector variable over an interval",
  "process_name": "AssignVectorVariableProcess",
  "Parameters": {
    "model_part_name": "Structure.DISPLACEMENT_Movingnodes",
    "variable_name": "DISPLACEMENT",
    "constrained": [true, true, true],
    "value": [0.0, "-0.005*t", 0.0],
    "interval": [0.0, 0.025]
  }
},
{

```

```

    "python_module": "assign_vector_variable_process",
    "kratos_module": "KratosMultiphysics",
    "help": "Hold final displacement after ramp",
    "process_name": "AssignVectorVariableProcess",
    "Parameters": {
        "model_part_name": "Structure.DISPLACEMENT_Movingnodes",
        "variable_name": "DISPLACEMENT",
        "constrained": [true, true, true],
        "value": [0.0, "-0.005*0.025", 0.0],
        "interval": [0.025, "End"]
    }
}
],
"loads_process_list": [],
"list_other_processes": []
},

"output_processes": {
    "gid_output": [
        {
            "python_module": "gid_output_process",
            "kratos_module": "KratosMultiphysics",
            "process_name": "GiDOutputProcess",
            "help": "GiD post-processing",
            "Parameters": {
                "model_part_name": "Structure",
                "output_name": "fem_region_longer",
                "postprocess_parameters": {
                    "result_file_configuration": {
                        "gidpost_flags": {
                            "GiDPostMode": "GiD_PostAscii",
                            "WriteDeformedMeshFlag": "WriteDeformed",
                            "WriteConditionsFlag": "WriteConditions",
                            "MultiFileFlag": "SingleFile"
                        },
                    },
                "file_label": "step",
            }
        }
    ]
}

```

```

        "output_control_type": "step",
        "output_interval": 100,
        "body_output": true,
        "node_output": false,
        "skin_output": false,
        "plane_output": [],
        "nodal_results": [
            "DISPLACEMENT", "VELOCITY", "REACTION", "
            VOLUME_ACCELERATION", "POINT_LOAD"
            // POINT_LOAD is included so coupled nodal loads
            can be visualised
        ],
        "gauss_point_results": [
            "VON_MISES_STRESS", "CAUCHY_STRESS_VECTOR", "
            GREEN_LAGRANGE_STRAIN_VECTOR"
        ],
        "nodal_nonhistorical_results": []
    },
    "point_data_configuration": []
}
}
}
],
"vtk_output": [ {
    "python_module": "vtk_output_process",
    "kratos_module": "KratosMultiphysics",
    "process_name": "VtkOutputProcess",
    "help": "ParaView output",
    "Parameters": {
        "model_part_name": "Structure",
        "output_control_type": "step",
        "output_interval": 100,
        "file_format": "ascii",
        "output_precision": 7,
        "output_sub_model_parts": false,
        "output_path": "vtk_output",

```

```
    "save_output_files_in_folder": true,  
    "nodal_solution_step_data_variables": ["DISPLACEMENT",  
      REACTION"],  
    "nodal_data_value_variables": [],  
    "element_data_value_variables": [],  
    "condition_data_value_variables": [],  
    "gauss_point_variables_extrapolated_to_nodes": ["  
      VON_MISES_STRESS"]  
  }  
}  
]  
}  
}
```

Bibliography

- [1] Thomas Weinhart, Carlos Labra, Stefan Luding, and Jin Y Ooi. Influence of coarse-graining parameters on the analysis of dem simulations of silo flow. *Powder technology*, 293:138–148, 2016.
- [2] Paul T Bauman, Hachmi Ben Dhia, Nadia Elkhodja, J Tinsley Oden, and Serge Prudhomme. On the application of the Arlequin method to the coupling of particle and continuum models. *Computational Mechanics*, 42:511–530, 2008. doi: 10.1007/s00466-008-0291-1.
- [3] Christian Wellmann and Peter Wriggers. A two-scale model of granular materials. *Computer Methods in Applied Mechanics and Engineering*, 205:46–58, 2012. doi: 10.1016/j.cma.2010.12.023.
- [4] Peter Yichen Chen, Maytee Chantharayukhonthorn, Yonghao Yue, Eitan Grinspun, and Ken Kamrin. Hybrid discrete-continuum modeling of shear localization in granular media. *Journal of the Mechanics and Physics of Solids*, 153:104404, 2021. doi: 10.1016/j.jmps.2021.104404.
- [5] Mohsin Ali Chaudry, Christian Woitzik, Alexander Düster, and Peter Wriggers. A multiscale dem–fem coupled approach for the investigation of granules as crash-absorber in ship building. *Computational Particle Mechanics*, 9(1):179–197, 2022.
- [6] Hongyang Cheng, Anthony R Thornton, Stefan Luding, Andrew L Hazel, and Thomas Weinhart. Concurrent multi-scale modeling of granular materials: Role of coarse-graining in FEM-DEM coupling. *Computer Methods in Applied Mechanics and Engineering*, 403:115651, 2023. doi: 10.1016/j.cma.2022.115651.
- [7] Jacques Duran. *Sables, poudres et grains: introduction à la physique des milieux granulaires*. Eyrolles, 1997.
- [8] Lyn Bates. Bulk solids characterisation-the need for industrial education in bulk technology. *Bulk Solids Handling*, 26(7):464–475, 2006.
- [9] Bruno Andreotti, Yoël Forterre, and Olivier Pouliquen. *Granular media: between fluid and solid*. Cambridge University Press, 2013.

- [10] Cun Zhang, Yanhong Chen, Yongle Wang, and Qingsheng Bai. Discrete element method simulation of granular materials considering particle breakage in geotechnical and mining engineering: A short review. *Green and Smart Mining Engineering*, 2024.
- [11] Renaud Delannay, Alexandre Valance, Anne Mangeney, Olivier Roche, and Patrick Richard. Granular and particle-laden flows: from laboratory experiments to field observations. *Journal of Physics D: Applied Physics*, 50(5):053001, 2017.
- [12] A Bajpai, A Bhateja, and R Kumar. Plume-surface interaction during lunar landing using a two-way coupled dsmc-dem approach. *Physical Review Fluids*, 9(2):024306, 2024.
- [13] P. A. Cundall and O. D. L. Strack. A discrete numerical model for granular assemblies. *Géotechnique*, 29(1):47–65, jan 1979. ISSN 0016-8505, 1751-7656. doi: 10.1680/geot.1979.29.1.47.
- [14] Jin Yeom Ooi and J Michael Rotter. Wall pressures in squat steel silos from simple finite element analysis. *Computers & Structures*, 37(4):361–374, 1990. doi: 10.1016/0045-7949(90)90026-X.
- [15] Jian Fei Chen, SK Yu, JY Ooi, and JM Rotter. Finite-element modeling of filling pressures in a full-scale silo. *Journal of Engineering Mechanics*, 127(10):1058–1066, 2001. doi: 10.1061/(ASCE)0733-9399(2001)127:10(1058).
- [16] Richard A Link and Alaa E Elwi. Incipient flow in silo-hopper configurations. *Journal of Engineering Mechanics*, 116(1):172–188, 1990. doi: 10.1061/(ASCE)0733-9399(1990)116:1(172).
- [17] Ulrich Häussler and Josef Eibl. Numerical investigations on discharging silos. *Journal of Engineering Mechanics*, 110(6):957–971, 1984. doi: 10.1061/(ASCE)0733-9399(1984)110:6(957).
- [18] P-A Guidault and Ted Belytschko. On the L2 and the H1 couplings for an overlapping domain decomposition method using Lagrange multipliers. *International Journal for Numerical Methods in Engineering*, 70(3):322–350, 2007. doi: 10.1002/nme.1882.
- [19] Hachmi Ben Dhia and Guillaume Rateau. The Arlequin method as a flexible engineering design tool. *International Journal for Numerical Methods in Engineering*, 62(11):1442–1462, 2005. doi: 10.1002/nme.1229.

- [20] Ted Belytschko, SP Xiao, George C Schatz, and RS Ruoff. Atomistic simulations of nanotube fracture. *Physical Review B*, 65(23):235430, 2002.
- [21] Veronika Singer, Klaus B Sautter, Antonia Larese, Roland Wüchner, and Kai-Uwe Bletzinger. A partitioned material point method and discrete element method coupling scheme. *Advanced Modeling and Simulation in Engineering Sciences*, 9(1):16, 2022.
- [22] Hachmi Ben Dhia. Multiscale mechanical problems: the Arlequin method. *Comptes Rendus de l'Académie des Sciences – Series IIB – Mechanics-Physics-Astronomy*, 12(326):899–904, 1998. doi: 10.1016/S1251-8069(99)80046-5.
- [23] Fangping Ye, Yue Zhang Qiang, Weijie Jiang, and Xiang Fu. Dem–fem coupling simulation of the transfer chute wear with the dynamic calibration dem parameters. *Processes*, 9(10):1847, 2021.
- [24] Klaus Bernd Sautter, Helene Hofmann, Corinna Wendeler, Peter Wilson, Philipp Bucher, Kai-Uwe Bletzinger, and Roland Wüchner. Advanced modeling and simulation of rockfall attenuator barriers via partitioned dem-fem coupling. *Frontiers in Built Environment*, 7:659382, 2021.
- [25] Guilherme Barros, Andre Pereira, Jerzy Rojek, John Carter, and Klaus Thoeni. Time domain coupling of the boundary and discrete element methods for 3d problems. *Computational Mechanics*, 74(4):779–797, 2024.
- [26] Manon Voisin-Leprince, Joaquin Garcia-Suarez, Guillaume Anciaux, and Jean-François Molinari. Finite element method–discrete element method bridging coupling for the modeling of gouge. *International Journal for Numerical Methods in Engineering*, 124(6):1474–1488, 2023.
- [27] Hachmi Ben Dhia. Global-local approaches: the Arlequin framework. *European Journal of Computational Mechanics*, 15(1-3):67–80, 2006. doi: 10.3166/remn.15.67-80.
- [28] Hachmi Ben Dhia. Mathematical analysis of the mixed Arlequin method. *Comptes Rendus de l'Académie des Sciences – Series I – Mathematics*, 332(7):649–654, 2001. doi: 10.1016/S0764-4442(01)01900-0.
- [29] Hachmi Ben Dhia and Guillaume Rateau. Application of the Arlequin method to some structures with defects. *Revue Européenne des Éléments Finis*, 11(2-4): 291–304, 2002. doi: 10.3166/reef.11.291-304.

- [30] Leon B Lucy. A numerical approach to the testing of the fission hypothesis. *The Astronomical Journal*, 82:1013–1024, 1977. doi: 10.1086/112164.
- [31] Joe J Monaghan. Smoothed particle hydrodynamics. *Annual Review of Astronomy and Astrophysics*, 30:543–574, 1992. doi: 10.1146/annurev.aa.30.090192.002551.
- [32] SG Bardenhagen, JU Brackbill, and Deborah Sulsky. Numerical study of stress distribution in sheared granular material in two dimensions. *Physical Review E*, 62(3):3882, 2000. doi: 10.1103/PhysRevE.62.3882.
- [33] Yonghao Yue, Breannan Smith, Peter Yichen Chen, Maytee Chantharayukhonthorn, Ken Kamrin, and Eitan Grinspun. Hybrid grains: Adaptive coupling of discrete and continuum simulations of granular media. *ACM Transactions on Graphics (TOG)*, 37(6):1–19, 2018.
- [34] Manon Voisin-Leprince, Joaquin Garcia-Suarez, Guillaume Anciaux, and Jean-François Molinari. Two-scale concurrent simulations for crack propagation using fem–dem bridging coupling. *Computational Particle Mechanics*, 11(5):2235–2243, 2024.
- [35] Yifan Bai, Xiaojing Li, Weimin Yang, Zhenhao Xu, and Mingying Lv. Multiscale analysis of tunnel surrounding rock disturbance: a pfc3d-flac3d coupling algorithm with the overlapping domain method. *Computers and Geotechnics*, 147:104752, 2022.
- [36] Romain Ruysen and Hachmi Ben Dhia. A finite addition of matter elements method for modeling and solution of an slm thermal problem by a multiscale method. *International Journal for Numerical Methods in Engineering*, 123(8): 1760–1790, 2022.
- [37] Wei Sun, Jian-Min Zhang, Jacob Fish, and Rui Wang. Superposition-based concurrent multiscale approaches for porodynamics. *International Journal for Numerical and Analytical Methods in Geomechanics*, 48(16):3909–3932, 2024.
- [38] Philip P Müller, Falk K Wittel, and David S Kammer. Damage preserving transformation for materials with microstructure. *European Journal of Mechanics-A/Solids*, 100:104995, 2023.
- [39] Xue Zhang, Ziping Zhu, and Baoyin Sun. A practical multiscale modeling strategy to model wet-type beam–column connections for seismic analysis of precast rc frame structures. *Soil Dynamics and Earthquake Engineering*, 168:107824, 2023.

- [40] Shuai Zhang, Wei Ge, and Chunjiang Liu. Spatial–temporal multiscale discrete–continuum simulation of granular flow. *Physics of Fluids*, 35(5), 2023.
- [41] Hachmi Ben Dhia and Chokri Zammali. Multiscale analysis of impacted thin structures. In *Computational of Shell and Spacial Structures*, pages 1–4, 2005.
- [42] Hachmi Ben Dhia. Further insights by theoretical investigations of the multiscale arlequin method. *International Journal for Multiscale Computational Engineering*, 6(3), 2008.
- [43] Ludovic Chamoin, Serge Prudhomme, H Ben Dhia, and Tinsley Oden. Ghost forces and spurious effects in atomic-to-continuum coupling methods by the arlequin approach. *International journal for numerical methods in engineering*, 83(8-9):1081–1113, 2010.
- [44] SP Xiao and Ted Belytschko. A bridging domain method for coupling continua with molecular dynamics. *Computer methods in applied mechanics and engineering*, 193(17-20):1645–1669, 2004.
- [45] P-A Guidault and Ted Belytschko. Bridging domain methods for coupled atomistic–continuum models with l2 or h1 couplings. *International journal for numerical methods in engineering*, 77(11):1566–1592, 2009.
- [46] Jerzy Rojek and Eugenio Oñate. Multiscale analysis using a coupled discrete/finite element model. *Interaction and Multiscale Mechanics*, 1(1):1–31, 2007.
- [47] Hua Qiao and Wei-qiu Chen. Analysis of the penalty version of the arlequin framework for the prediction of structural responses with large deformations. *Journal of Zhejiang University-SCIENCE A*, 12(7):552–560, 2011.
- [48] Fei Han and Gilles Lubineau. Coupling of nonlocal and local continuum models by the arlequin approach. *International Journal for Numerical Methods in Engineering*, 89(6):671–685, 2012.
- [49] Hossein Talebi, Mohammad Silani, Stéphane PA Bordas, Pierre Kerfriden, and Timon Rabczuk. A computational library for multiscale modeling of material failure. *Computational Mechanics*, 53:1047–1071, 2014.
- [50] Ilias N Giannakeas, Theodosios K Papathanasiou, Arash S Fallah, and Hamid Bahai. Coupling xfem and peridynamics for brittle fracture simulation—part i: feasibility and effectiveness. *Computational Mechanics*, 66:103–122, 2020.

- [51] Thomas Schwager and Thorsten Pöschel. Coefficient of restitution for viscoelastic spheres: The effect of delayed recovery. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 78(5):051304, 2008.
- [52] Jun Ai, Jian-Fei Chen, J Michael Rotter, and Jin Y Ooi. Assessment of rolling resistance models in discrete element simulations. *Powder Technology*, 206(3): 269–282, 2011.
- [53] Thomas J. R. Hughes. *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*. Prentice Hall, 1987.
- [54] O. C. Zienkiewicz, R. L. Taylor, and J. Z. Zhu. *The Finite Element Method: Its Basis and Fundamentals*. Butterworth–Heinemann, 7 edition, 2013.
- [55] Klaus-Jürgen Bathe. *Finite Element Procedures*. Prentice Hall, 1996.
- [56] Juan C. Simo and Thomas J. R. Hughes. *Computational Inelasticity*, volume 7 of *Interdisciplinary Applied Mathematics*. Springer, 1998.
- [57] Daniel C. Drucker and William Prager. Soil mechanics and plastic analysis or limit design. *Quarterly of Applied Mathematics*, 10(2):157–165, 1952.
- [58] Wai-Fah Chen and Da-Jian Han. *Plasticity for Structural Engineers*. J. Ross Publishing, reprint ed. edition, 2007.
- [59] Eduardo A. de Souza Neto, Djordje Perić, and David R. J. Owen. *Computational Methods for Plasticity: Theory and Applications*. Wiley, 2008.
- [60] M. Dosta, D. Andre, V. Angelidakis, R.A. Caulk, M.A. Celigueta, B. Chareyre, J.-F. Dietiker, J. Girardot, N. Govender, C. Hubert, R. Kobyłka, A.F. Moura, V. Skorych, D.K. Weatherley, and T. Weinhart. Comparing open-source DEM frameworks for simulations of common bulk processes. *Computer Physics Communications*, 296: 109066, 2024. ISSN 0010-4655. doi: 10.1016/j.cpc.2023.109066. URL <https://www.sciencedirect.com/science/article/pii/S0010465523004113>.
- [61] Hongyang Cheng, Takayuki Shuku, Klaus Thoeni, Pamela Tempone, Stefan Luding, and Vanessa Magnanimo. An iterative bayesian filtering framework for fast and automated calibration of dem models. *Computer methods in applied mechanics and engineering*, 350:268–294, 2019.

- [62] Hongyang Cheng, Luisa Orozco, Retief Lubbe, Aron Jansen, Philipp Hartmann, and Klaus Thoeni. Grainlearning: A bayesian uncertainty quantification toolbox for discrete and continuum numerical models of granular materials. *Journal of Open Source Software*, 9(97):6338, 2024. doi: 10.21105/joss.06338. URL <https://doi.org/10.21105/joss.06338>.
- [63] R. M. Nedderman. *Statics and Kinematics of Granular Materials*. Cambridge University Press, Cambridge, 1992. ISBN 9780521404442.
- [64] Songxiong Ding, H Li, JY Ooi, and JM Rotter. Prediction of flow patterns during silo discharges using a finite element approach and its preliminary experimental verification. *Particuology*, 18:42–49, 2015.
- [65] AM Sanad, JY Ooi, JMFG Holst, and JM Rotter. Computations of granular flow and pressures in a flat-bottomed silo. *Journal of engineering mechanics*, 127(10): 1033–1043, 2001.
- [66] Yin Wang, Yong Lu, and Jin Y Ooi. Numerical modelling of dynamic pressure and flow in hopper discharge using the arbitrary lagrangian–eulerian formulation. *Engineering Structures*, 56:1308–1320, 2013.
- [67] Yin Wang, Yong Lu, and Jin Y Ooi. A numerical study of wall pressure and granular flow in a flat-bottomed silo. *Powder Technology*, 282:43–54, 2015.
- [68] Matthias Sperl. Experiments on corn pressure in silo cells—translation and comment of Janssen’s paper from 1895. *Granular Matter*, 8(2):59–65, 2006. doi: 10.1007/s10035-006-0003-2.
- [69] Xin Wang, Yifu Shi, Bing Luo, Cai Liang, Daoyin Liu, Jiliang Ma, and Xiaoping Chen. Flow profile and wall normal stress characteristics in pattern-transformable flow silos. *Chemical Engineering Research and Design*, 182:381–394, 2022.