6th European Conference on Computational Mechanics (ECCM 6)
7th European Conference on Computational Fluid Dynamics (ECFD 7)
1115 June 2018, Glasgow, UK

A PARALLEL MULTISCALE DEM-VOF METHOD FOR LARGE-SCALE SIMULATIONS OF THREE-PHASE FLOWS

GABRIELE POZZETTI[1], XAVIER BESSERON[1] ALBAN ROUSSET[1] and BERNHARD PETERS[1]

[1] University of Luxembourg
6 Avenue de la Fonte Esch-Sur-Alzette Luxmbourg
pozzetti.gabriele@gmail.com www.xdem.de

Key words: Parallel Computing, Multiscale , DEM-VOF Method, Computing Methods

Abstract. A parallel dual-grid multiscale DEM-VOF coupling is here investigated. Dual-grid multiscale couplings have been recently used to address different engineering problems involving the interaction between granular phases and complex fluid flows. Nevertheless, previous studies did not focus on the parallel performance of such a coupling and were, therefore, limited to relatively small applications. In this contribution, we propose an insight into the performance of the dual-grid multiscale DEM-VOF method for three-phase flows when operated in parallel. In particular,we focus on a famous benchmark case for three-phase flows and assess the influence of the partitioning algorithm on the scalability of the dual-grid algorithm.

1 INTRODUCTION

In this work, a parallelization strategy for a multiscale DEM-VOF method is investigated. This CFD-DEM coupling was recently used to approach several different engineering problems, ranging from process industry [1, 3] to additive manufacturing [4]. Nevertheless, achieving good parallel performance is mandatory in order to address a wider range of engineering applications within a reasonable computational time [2].

The dual-grid multiscale DEM-VOF method introduced in [1] holds significant advantages over standard DEM-VOF method as it allows obtaining grid-convergent results and a better interface tracking. Despite the fact that its computational cost is comparable with the mono-scale DEM-VOF method, its enhanced complexity makes its parallel execution non trivial, as the inter-scale communication introduced in with the dual-grid approach can represent a possible issue for the algorithm performance.

In [12], we proposed a parallelization strategy for generic CFD-DEM couplings that allows overcoming and important inter-physics communication bottleneck. This strategy consists of imposing a co-location constraint between the partitions of the CFD and DEM domains and perform the inter-physics exchanges locally. As shown in [12] this allows improving the performance of the inter-physic exchange, but it reduces the flexibility in the partitioning of the domain, not allowing partitioning configurations that violates the

co-location constraint. In [12], we pointed out how this limitation would be more severe in case of non-uniformly distributed loads.

In this contribution, we investigate the parallel performance of the multiscale DEM-VOF method and compare it with a standard DEM-VOF coupling parallelized with an overlapping domain approach [5]. We show how the enhanced complexity of the multiscale method can be translated in an higher flexibility in the domain partitioning, and how this advantage allows it to address large-scale problems and scale over hundreds of processes.

2 METHODOLOGY

In the multiscale DEM-VOF method, two different CFD grids are adopted, one for the coupling between the CFD and the DEM code and one for the solution of the fluid equations [1]. As shown in figure 1, a coarse grid is chosen to perform the coupling between CFD and DEM code at a bulk scale, while a finer and non-uniform grid is adopted to discretize the CFD equations. An interpolation strategy between the grids ensures the correct exchange of information between the bulk scale at which the inter-physics coupling is performed and the fluid fine scale at which the fluid equations are solved. In [1] the method was shown to produce grid-convergent results and to provide a higher accuracy if compared to a standard DEM-VOF method. Nevertheless, its enhanced complexity can represent an issue for its parallelization, that is needed to approach large-scale problems [2].

In this contribution, the dynamic module of the XDEM Platform [1, 6, 4], is used to resolve the dynamic of the discrete entities moving in the presence of a multiphase flow. We here repropose the equations resolved by this module, for a deeper insight into those equation we refer the reader to [1], and citation within. Positions and orientations of the particles are updated at every time-step according to

$$m_i \frac{d^2}{dt^2}\mathbf{x}_i = \mathbf{F}_{coll} + \mathbf{F}_{drag} + \mathbf{F}_g, \tag{1}$$

$$\mathbf{I}_i \frac{d^2}{dt^2}\boldsymbol{\phi}_i = \mathbf{M}_{coll} + \mathbf{M}_{ext}, \tag{2}$$

where $\mathbf{x}_i$ are the positions, m_i the masses, and $\boldsymbol{\phi}$ the orientations of the entities. The term $\mathbf{F}_c$ accounts for the force arising from collisions

$$\mathbf{F}_{coll} = \sum_{i \neq j} \mathbf{F}_{ij}(\mathbf{x}_j, \mathbf{u}_j, \boldsymbol{\phi}_j, \boldsymbol{\omega}_j), \tag{3}$$

with $\mathbf{u}_j$ the velocity of particle j, and $\boldsymbol{\omega}$ the angular velocity. The term $\mathbf{M}_{coll}$ is the torque acting on the particle due to collisions

$$\mathbf{M}_{coll} = \sum_{i \neq j} \mathbf{M}_{ij}(\mathbf{x}_j, \mathbf{u}_j, \boldsymbol{\phi}_j, \boldsymbol{\omega}_j), \tag{4}$$

with $\mathbf{M}_{ij}$ the torque acted to particle i from particle j . The term $\mathbf{F}_{drag}$ takes into account the force rising from the interaction with the fluid. Finally, $\mathbf{F}_g$ accounts for the

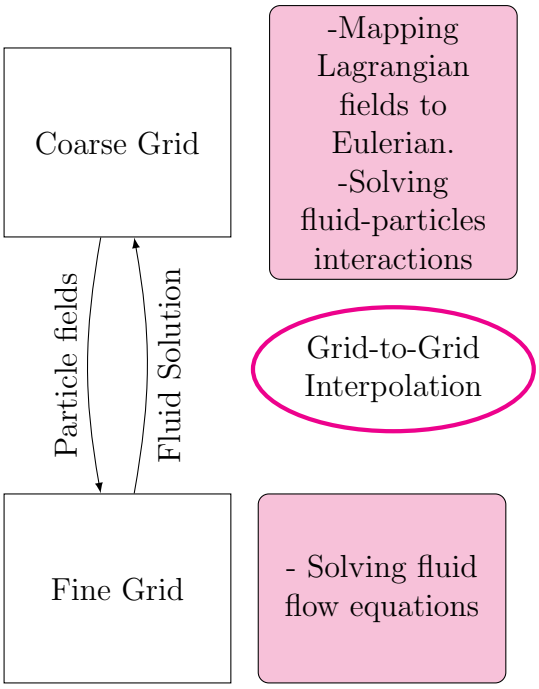

Figure 1: Schematic of the solution procedure for the bulk and fine length-scale in the simulation. The two boxes represent the different models adopted, while the arrows show schematically the communication between the scales. A coarse grid (top) is used to perform the volume averaging and to solve the fluid-particle interaction. Particle-related fields are mapped to the supporting domain (bottom) then, a finer grid is used to solve the fluid equations.

gravitational force. For the estimation of $\mathbf{F}_{drag}$ we adopt a semi-empirical

$$\mathbf{F}_{drag} = \beta(\mathbf{u}_f - \mathbf{u}_p), \tag{5}$$

$$\beta = \beta(\mathbf{u}_f - \mathbf{u}_p, \rho_f, \rho_p, d_p, A_p, \mu_f, \epsilon), \tag{6}$$

where $\mathbf{u}_f$, $\mathbf{u}_p$ the fluid and particle velocity respectively, ρ_f, ρ_p the respective densities, d_p, A_p the particle characteristic length and area, μ_f the fluid viscosity, and ϵ the porosity, defined as the ratio between the volume occupied by the fluid and the total volume of the CFD cell. For the sake of generality, β is modeled according to [8].

During its parallel execution, XDEM geometrically decomposes the simulation domain in regularly fixed-size cells that are used to distribute the load between the processes. As better described in [12], every process is only responsible for a set of cells that will define its sub-domain. In this way, every process only performs the calculation and holds knowledge of the particles that belong to its sub-domain.

The XDEM platform offers different possibilities for the domain partitioning [7]. in particular, a dedicated partitioner that is able to force a co-location constraint between XDEM and OpenFoam partitions as proposed in [12]. In this contribution, we compare the performance obtained with the above mentioned partitioner, a Scotch partitioner, Zoltan-RIB and and one based on Orthogonal Recursive Bisection (ORB).

2.1 Equations solved in the CFD Domain

The fluid flow equations are instead resolved with the OpenFOAM-extend libraries. The solver proposed in [1, 9] is adopted for tackling the general case of an unsteady incompressible multiphase flow through an evolving porous media. The set of equations governing such incompressible particle-laden flow takes the form:

$$\frac{\partial \epsilon \rho_f \mathbf{u}_f}{\partial t} + \nabla \cdot (\epsilon \rho_f \mathbf{u}_f \mathbf{u}_f) = -\epsilon \nabla p + \nabla \cdot \left(\epsilon \mu_f \left(\nabla \mathbf{u}_f + \nabla^T \mathbf{u}_f \right) \right) + \mathbf{T_\Gamma} + \mathbf{F_b} + \mathbf{F_{fpi}}, \quad (7)$$

$$\nabla \cdot \epsilon \mathbf{u}_f = \frac{\partial \epsilon}{\partial \tau}. \quad (8)$$

with $\mathbf{u}_f$ the fluid velocity, p the fluid pressure, $\mathbf{F_b}$ a generic body force. $\mathbf{F_{fpi}}$ is the fluid-particle interaction force, that is the exact counterpart of $\mathbf{F}_{drag}$. The definition of $\mathbf{F_{fpi}}$ is treated with the semi-implicit algorithm proposed in [10]. Viscosity and density are functions on the fluid phase and can be written in the form:

$$\rho_f(\mathbf{x}) = \rho_1 \alpha(\mathbf{x}) + \rho_2 (1 - \alpha(\mathbf{x})), \quad (9)$$

$$\mu_f(\mathbf{x}) = \mu_1 \alpha(\mathbf{x}) + \mu_2 (1 - \alpha(\mathbf{x})), \quad (10)$$

where α is the volume fraction defined as

$$\alpha = \frac{1}{V} \int_V \chi(\mathbf{x}) d\mathbf{x}, \quad (11)$$

$$\chi = \begin{cases} 1 & \text{if first fluid,} \\ 0 & \text{if second fluid.} \end{cases} \quad (12)$$

α is modeled as a scalar transported by the fluid flow for which:

$$\frac{\partial \epsilon \alpha}{\partial t} + \nabla \cdot (\epsilon \alpha \mathbf{u}_f) + \nabla \cdot (\epsilon \alpha (1 - \alpha) \mathbf{u}_c) = 0, \quad (13)$$

must hold. In equation 13, $\mathbf{u}_c$ is the relative velocity between the two-phases referred to as compression velocity. The third term is introduced in order to avoid an excessive numerical dissipation.

The OpenFOAM libraries [11] are as well parallelized using domain decomposition. The computational domain, where the CFD equations are defined, is split into sub-domains assigned to each process available at run time. A separate copy of the code is run on each MPI process. The single processes exchange information through a dedicated boundary class as described in [11]. The OpenFOAM libraries offer different partitioning algorithm. In this contribution, we partition the CFD grid of the monoscale DEM-VOF coupling and the coarse grid of the multiscale DEM-VOF coupling enforcing a co-location with the DEM domain as described in [12] For the fine grid we instead use the native Scotch partitioner of OpenFOAM.

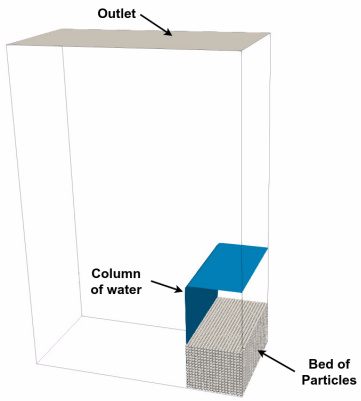

Figure 2: Three-phase dam-break. Simulation setup.

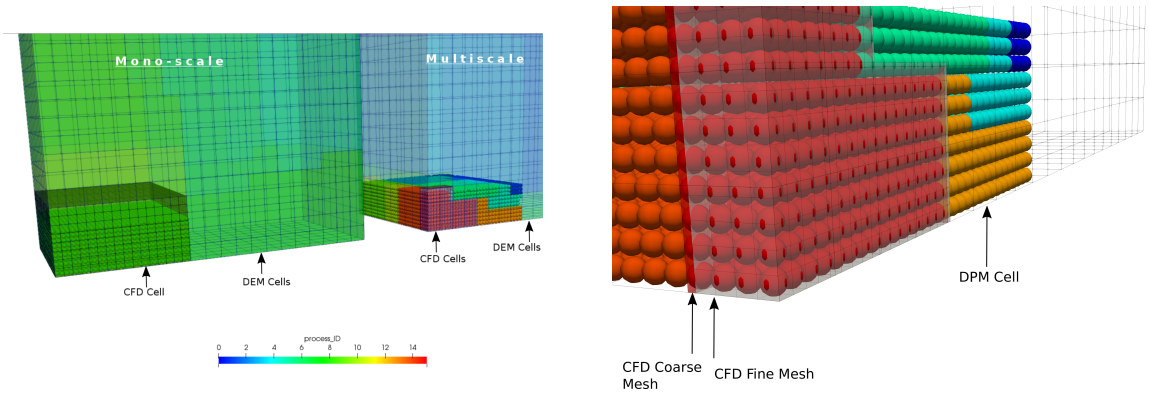

(a) (b)

Figure 3: Three-phase dam-break. Different Partition strategies with mono-scale and multiscale approach 3a, and corner zoom for the Multiscale approach 3b.

3 TEST-CASES

The dam break is a very famous benchmark for two-phase flows simulations. In [1], the benchmark was chosen to underline the advantages of the multiscale DEM-VOF method over the classical DEM-VOF method in terms of accuracy and computational traceability. Here, we propose the parallel execution of this benchmark to show how the enhanced complexity of the multiscale approach can translate into a higher flexibility for the domain partitioning. This offers significant advantages in terms of parallel performance and, in particular, allows to maintain the advantages of the strategy proposed in [12], and at the same time to use advanced partitioning strategies. With this purpose, we compare the parallel execution of a mono-scale DEM-VOF coupling using a co-located partitions strategy, with the parallel execution of a multiscale DEM-VOF coupling where the coarse grid is partitioned using a co-located partitions strategy, while the fine grid is partitioned

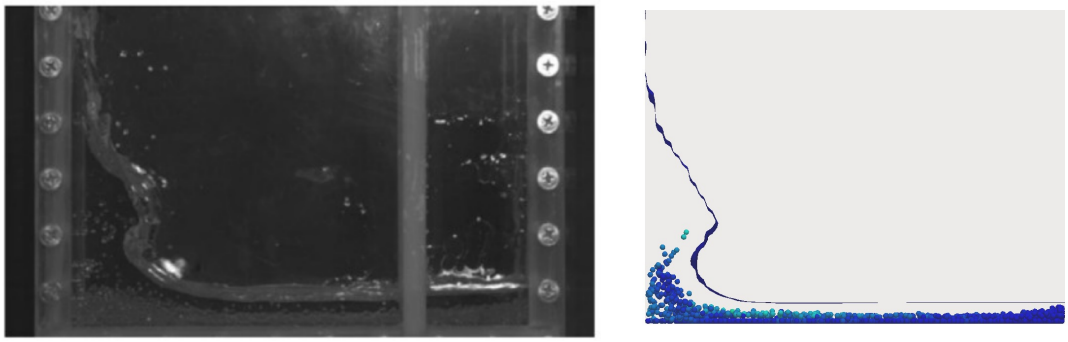

Figure 4: Three-phase dam-break. Comparison between Experimental data as in [1] and numerical simulation with the parallel multiscale DEM-VOF Method. Good agreement can be observed.

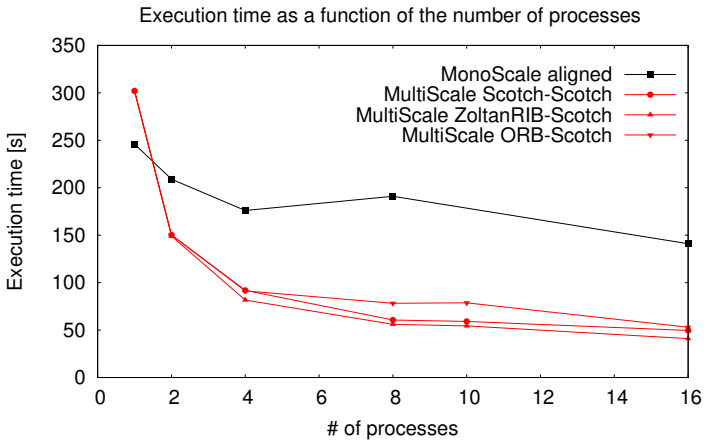

Figure 5: Three-phase dam-break. Execution time as a function of the number of processes. better parallel performance is observed for the multiscale approach

with native OpenFOAM algorithms.

The simulation parameters are chosen accordingly to what presented in [1]. As shown in figure 2, a configuration without intermediate obstacle is adopted, in a box of dimensions $0.2m$ x $0.1m$ x $0.3m$. A column of water of extension $0.05m$ x $0.1m$ x $0.1m$ is posed at the left corner containing a uniformly layered bed of spherical particles. All the spheres have an identical diameter of $2.7\,mm$. As already proposed in [1], we adopt a fine grid discretized with $500k$ identical cubic cells while the coarse grid discretized with $6k$ identical volumes. For the mono-scale approach, a domain discretization of $48k$ cells is adopted, and no coarse grid is used.

The liquid (heavy phase) density and viscosity are $1000kg/m^3$ and $10^{-3}\,Pa\,s$ respectively. The gas (light phase) features a density of $1\,kg/m^3$ and a viscosity of $10^{-5}\,Pa\,s$. The particle density is $2500\,kg/m^3$. For the particle inter-collisions and the particle-wall collisions, a linear dashpot impact model is chosen with a spring constant of $1000\,N/m$, a restitution coefficient of 0.9 and a friction coefficient of 0.3.

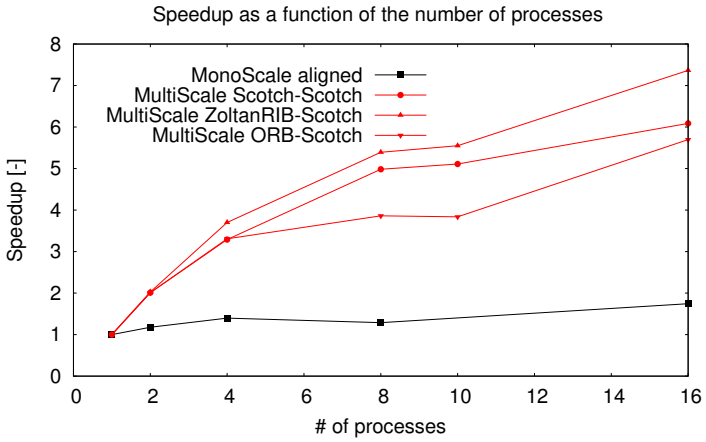

Figure 6: Three-phase dam-break. Speedup as a function of the number of processes. better parallel performance is observed for the multiscale approach

Figure 4 shows how the parallel execution with the strategy presented in this article, produces the same results as the sequential that have a very good agreement with experimental data. As extensively pointed out in [1], the accuracy obtained with this approach is significantly better than the one provided by the monoscale DEM-VOF.

In figure 3, the partitioning for the two cases is shown. One can observe how, in the mono-scale case, both the CFD and the DEM domains are partitioned uniformly in order to ensure perfect alignment between the partitions, according to a standard co-located partitions strategy. Due to the non-uniform distribution of the particles, this constraint results in all the particles to be assigned to a single process. On the other hand, in the multiscale approach, the coarse grid and the DPM grid are always perfectly aligned by construction. This is possible due to the fact that the coarse grid can be chosen independently from the fluid requirements. Therefore, the fine grid can be distributed in an independent way, trying to optimize the ratio between the load balance and the communication.

As preliminary discussed in [1], the usage of a finer CFD grid makes the multiscale approach more computationally expensive than a mono-scale DEM-VOF. Nevertheless, when operated in sequential, the overall computation time is comparable with the one involved in a mono-scale simulation. One can observe this in figure 5, where we report the computational time of the simulation time as a function of the number of MPI processes for both the mono-scale and the multiscale approaches. It is important to notice how, despite the multiscale simulation requires more time in sequential, it performs significantly better in the parallel execution. In particular, when using more than 2 processes, the multiscale approach is faster than the mono-scale despite the heavier CFD load. Furthermore, it can be observed how the partitioning strategy influences the performance of the coupling, that always remains better than the standard DEM-VOF. This can be explained by observing figure 6 that proposes the speedup of the two approaches as a function of the number of processes. It is clearly shown how the multiscale approach performs significantly better

than the mono-scale one.

In conclusion, the multiscale approach to the DEM-VOF coupling has the important advantage over the standard DEM-VOF method of allowing adopting more complex partitioning strategies. This leads to improved parallel performance up to the counter-intuitive result of allowing a simulation that uses an heavier CFD grid to run faster than one that uses a coarser one. Therefore, if compared to a standard monoscale method, the multiscale approach is shown to provide not only a better accuracy but also a reduced computational time when parallel execution is involved.

3.1 CONCLUSIONS

We investigated the parallel performance of a dual-grid multiscale DEM-VOF coupling. Perfectly aligned co-located partitions between the DEM domain and the CFD domain associated with the bulk scale were chosen, while an independent discretization for the CFD domain associated with the fine scale was adopted. This allowed avoiding inter-process communication between the CFD and the DEM part and, at the same time, keeping flexibility on the domain partitioning. A benchmark case has been discussed to assess consistency and performances of the proposed strategy. Results showed how the dual-grid multiscale approach can achieve better parallel performance than a single-grid CFD-DEM coupling for inhomogeneous cases. Furthermore, different partitioning algorithms have been tested and results reported for offering guidance to future works. One of the main benefits of the current strategy consists of the possibility of adopting different partitioning algorithms for the DEM domain and the CFD fine grid, that can be completely independent from one another without triggering inter-physics communication. For future studies, an in-depth analysis of the parallel performance of the grid-to-grid interpolation would be of great interest.

Acknowledgments

This research is in the framework of the project DigitalTwin, supported by the programme Investissement pour la compétitivité et emploi - European Regional Development Fund (Grant agreement: 2016-01-002-06). The experiments presented in this paper were carried out using the HPC facilities of the University of Luxembourg.

REFERENCES

[1] Gabriele Pozetti and B. Peters, *A multiscale DEM-VOF method for the simulation of three-phase flows*. International Journal of Multiphase flows, 2018.

[2] D.K. Kafui, S.Johnson, C.Thornton, and J.P.K.Seville *Parallelization of a lagrangianeulerian dem/cfd code for application to fluidized beds* Powder Technology, 207(1):270-278,2011

[3] Gabriele Pozzetti, Bernhard Peters , A Numerical approach for the evaluation of particle-induced erosion in an Abrasive Waterjet focusing tube. doi:10.1016/j.powtec.2018.04.006 Gabriele Pozetti and B. Peters, *A multiscale DEM-*

VOF method for the simulation of three-phase flows. International Journal of Multiphase flows, 2018.

[4] Bernhard Peters, and Gabriele Pozzetti, *Flow characteristics of metallic powder grains for additive manufacturing*. EPJ WC, 2017.

[5] Gabriele Pozzetti, Xavier Besseron, Alban Rousset W.Mainassara, Bernhard Peters, *On the performance of an overlapping-domain parallelization strategy for Eulerian-Lagrangian multiphysics software AIP CP*, 2018.

[6] Mehdi Baniasadi, Maryam Baniasadi, and Bernhard Peters. *Coupled cfd-dem with heat and mass transfer to investigate the melting of a granular packed bed.* Chemical Engineering Science,

[7] Alban Rousset, Xavier Besseron, and Bernhard Peters. *Parallelizing xdem: Load-balancing policies and efficiency, a study.* International Conference on Particle-Based Methods (PARTICLES17), 2017.

[8] L. Schiller and Z. Naumann. *A Drag Coefficient Corre-lation,*. VDI Zeitung, Vol. 77

[9] Gabriele Pozzetti and Bernhard Peters. *On the choice of a phase interchange strategy for a multiscale DEM-VOF Method*. AIP Conference Proceedings

[10] Heng Xiao and Jin Sun. *Algorithms in a Robust Hybrid CFD-DEM Solver for Particle-Laden Flows*. Communication Computer Physics, 2010.

[11] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby. *A tensorial approach to computational continuum mechanics using object-oriented techniques*. Computers in Physics 1998.

[12] G. Pozzetti, X. Besseron, A. Rousset, and B. Peters. *A co-located partitions strategy for parallel CFD-DEM couplings*. ArXiv e-prints, February 2018.