

28 College of Marine Science and Technology, China University of Geosciences, China

29 xue.zhang2@liverpool.ac.uk

30

31 **Abstract**

32 As one of the arbitrary Lagrangian-Eulerian methods, the material point method (MPM) owns
33 intrinsic advantages in simulation of large deformation problems by combining the merits of the
34 Lagrangian and Eulerian approaches. Significant computational intensity is involved in the
35 calculations of the MPM due to its very fine mesh needed to achieve a sufficiently high accuracy. A
36 new multiple-GPU parallel strategy is developed based on a single-root complex architecture of the
37 computer purely within a CUDA environment. Peer-to-Peer (P2P) communication between the GPUs
38 is performed to exchange the information of the crossing particles and ghost element nodes, which is
39 faster than the heavy send/receive operations between different computers through the infiniband
40 network. Domain decomposition is performed to split the whole computational task over the GPUs
41 with a number of subdomains. The computations within each subdomain are allocated on a
42 corresponding GPU using an enhanced ‘Particle-List’ scheme to tackle the data race during the
43 interpolation from associated particles to common nodes. The acceleration effect of the parallelisation
44 is evaluated with two benchmark cases, mini-slump test after a dam break and cone penetration test
45 in clay, where the maximum speedups with 1 and 8 GPUs are 88 and 604, respectively.

46

47 **Keywords:** material point method, parallel computation, cone penetration test, mini slump test

48

49 **1. Introduction**

50 The material point method (MPM), one of the arbitrary Lagrangian Eulerian methods, owns intrinsic
51 advantages in simulation of large deformation problems by combining the merits of the Lagrangian
52 and Eulerian methods^[1-4]. The Lagrangian particles, inheriting all the history-dependent information
53 of material, are allowed to move through the background Eulerian mesh, while the mesh is always
54 fixed in space to avoid the potential mesh distortion. The MPM, introduced to solid mechanics^[5] from
55 computational fluid dynamics^[6], was used to simulate high explosive explosions^[7], propagation of
56 wood cracks^[8], impact between solid bodies^[9-12], fluid-structure interactions^[13] and computer
57 animations^[14-16]. In the recent decade the MPM was applied to geotechnical engineering to investigate
58 runout of submarine landslides^[17-20], penetration and pull-out of structures^[21-23] and flow of granular
59 materials^[24-26]. Coupling analysis of pore or free water and soil, mainly used in the analysis of slope
60 stability^[27-30], is a new trend of the MPM simulations.

61 One of the main obstacles to the widespread application of the MPM is its low computational
62 efficiency, especially for large-scale and long-period problems. As the particles mostly are not at the

63 optimum locations for integration in the elements^[31-33], the mesh adopted in the MPM should be much
64 finer than that in large deformation finite element analysis to obtain sufficient accuracies^[3, 34-35].
65 Structured elements, used as often as the unstructured elements^[36-37], bring extra computational loads
66 with identical mesh size from the concerning domain to the far field. Elements in singularity zone
67 around structures need to be further refined for soil-structure interaction problems^[38-39]. Although an
68 initial assignment of four particles in per element is often sufficient to obtain a smooth stress/strain
69 field in many cases of MPM simulations^[40], the configuration of 16 particles in each element
70 sometimes is necessary for high-speed impacting problems^[14, 41]. Therefore, most existing MPM
71 analyses were limited to small-scale problems or two-dimensional plane-strain scenarios^[14, 21, 41].

72 Parallel computation on the central processing units (CPU) or graphic processing units (GPU) is the
73 most viable option to promote the efficiency of the MPM, which often requires special treatments to
74 make the algorithm more parallelisable. Acceleration effect of the parallelisation can be significantly
75 influenced by different parallel techniques and hardware platforms. Reference [9] and [42] proposed
76 a single-CPU parallelisation scheme of the MPM with the loop-based parallel library OpenMP,
77 achieving a five-fold speedup over a sequential calculation by mobilising eight CPU cores. The
78 OpenMP-based parallelisation is quite simple by invoking an executable directive before each loop
79 operation; however, its limitation is also obvious as most commercially available CPUs have less
80 than 32 cores. Reference [43] developed a multiple-CPU parallelisation strategy using the message
81 passing interface (MPI), accelerating the computation for up to 2,500 times with 16,384 CPU cores
82 on a supercomputer. In comparison with the CPU parallelisation, the state-of-the-art GPU
83 parallelisation is more cost effective as each GPU hosts thousands of GPU cores^[44-45], but its parallel
84 techniques are more complex. Reference [46] proposed a specialised parallelisation scheme with
85 single GPU by using the compute unified device architecture (CUDA), obtaining speedups of around
86 25 times given double precision numbers were used. Limited by the memory size dedicated on the
87 GPU, the maximum number of particles allowed in the MPM model was around six million.
88 Reference [47] adopted a similar technique to parallelise an implicit MPM algorithm and applied it
89 to computer animations. Reference [48] then extended the framework to orchestrate multiple GPUs
90 on a multiple-computer cluster based on a hybrid MPI-CUDA environment, which was sensitive to
91 the data exchange between the computers through a private network. Given 16 GPUs were used on
92 four tandem computers, up to 900 times speedup was then obtained with the maximum number of
93 particles as 96 million. Recently, Reference [49] further optimised the massively parallel framework
94 of the MPM and achieved over 100 times speedup on a single GPU; however, its acceleration effect
95 on multiple-GPU platform is heavily dependent on the hardware performance for the stream event
96 synchronisation on the specific GPU device; as a result, speedup of the framework is not always so

97 high (~ 100) as presented in Reference [49], especially for medium scale problems with less than
98 500,000 particles; and much complexities were caused in the parallel scheme (such as the Particle-
99 Grid offset technique) and the data transport between the GPUs (such as the AoSoA data structure),
100 which may undermine the reliability and maintainability of the programme. Therefore, a reliable and
101 efficient MPM program based on a simple multiple-GPU parallel framework is still needed.

102 In this paper, a parallelisation strategy with multiple GPUs is developed within the CUDA
103 environment. Different to that in Reference [48], the mobilised GPUs are hosted in an identical
104 computer platform with a shared random access memory (RAM). Peer-to-Peer (P2P) communication
105 between the GPUs is performed to exchange the information of the crossing particles and ghost
106 element nodes, which is faster than the heavy send/receive operations between different computers
107 through the infiniband network in Reference [48]. Domain decomposition is performed to split the
108 whole computational task over the GPUs with a number of subdomains. The computations within
109 each subdomain are allocated on a corresponding GPU and the MPM algorithm on each GPU is
110 parallelised with the technique proposed in Reference [46] with specific improvements, which further
111 enhance the speedup and reliability of the computation. The calculation results are assembled on the
112 shared RAM of the computer through the connection with the GPU devices. In comparison to the
113 parallel framework in Reference [49], the parallel strategy in this study is more reliable and friendly
114 to the new developers of the MPM, which also presents satisfying acceleration effects. Specifically,
115 this paper includes the following contributions: (1) an efficient parallel technique is proposed to
116 invoke multiple GPUs on an identical computer using P2P communication with each other; (2) a
117 hybrid memory IO framework is developed based on the shared RAM and distributed GPU memory
118 hierarchy; (3) an enhanced ‘Particle-List’ scheme to parallelise the interpolation from particles to
119 nodes, which is also parallelised on GPUs and hence avoids the frequent data exchange between the
120 CPU and GPUs; (4) the parallelised MPM algorithm is extended from two to three dimensional,
121 which is more computationally intensive and requires larger memory space.

122 **2. Material point method**

123 **2.1 MPM program**

124 The parallelisation strategy was developed based on an in-house program, MPM-GeoFluidFlow,
125 which stems from an open-source package, Uintah (<http://uintah.utah.edu/>), and features a novel
126 contact algorithm ‘Geo-contact’^[50], as well as a particle reseeding technique^[51]. Geo-contact,
127 specialised for soil-structure interactions, was developed from the conventional contact algorithm
128 with enhancement of a penalty function^[5, 13, 50, 52-53]. The explicit updated Lagrangian calculation in
129 each incremental step was based on the uGIMP method^[40, 54]. Meshes with identical sizes of square

130 elements were used^[18, 41], and unstructured elements can be found in Reference [22, 23]. The
 131 definition of the stresses and strains followed finite strain theory taking account of the incremental
 132 rotation of the configurations between time steps for objectivity: the stresses were measured with the
 133 Cauchy stress and updated with the Jaumann rate, and the strains were calculated with the
 134 deformation gradient. Applications of the programme are mainly focused on penetrometer
 135 penetration^[51], submarine landslide^[55], and impact dynamics^[18, 41]. In this paper we only describe the
 136 framework utilised to solve the mass and momentum equations, but it can be applied
 137 straightforwardly to other boundary-value problems, such as heat flux in an energy equation^[7].

138 2.2 Governing equations

139 The formulation was derived from the conservation of mass and linear momentum balance. The
 140 conservation of mass requires that the time derivative of the mass entering or leaving a specific
 141 domain is zero, which can be written in mathematical form as

$$142 \quad \frac{\partial \rho}{\partial t} + \rho \nabla v = 0 \quad (1)$$

143 in which ρ is the material density, v is the velocity and t is the time. In the MPM, Equation (1) is
 144 satisfied naturally by discretising the objects into a cloud of Lagrangian particles with consistent
 145 masses and volumes^[33].

146 The linear momentum balance means that the time-variation of the linear momentum of a material is
 147 equal to the resultant of the internal and external forces, i.e. Newton's second law of motion:

$$148 \quad \rho \frac{\partial v}{\partial t} = \nabla \sigma + \rho b \quad (2)$$

149 in which σ is the Cauchy stress, and b is the body force. Equation (2) is the strong form of the
 150 conservation of linear momentum, which is usually difficult to achieve as a closed-form solution due
 151 to mathematical difficulties. Therefore, the weak form is derived instead, expressed as

$$152 \quad \int_V \rho u \frac{\partial v}{\partial t} dV = - \int_V \sigma \nabla u dV + \int_V \rho u b dV + \int_V u T dS \quad (3)$$

153 in which u is the virtual velocity, V and S are the volume and surface area, and T is the prescribed
 154 surface traction. Numerical integration is adopted with the simplification of lumped mass, producing
 155 a concise form

$$156 \quad m a = F^{\text{ext}} + F^{\text{int}} \quad (4)$$

157 where m is the lumped mass, a is the acceleration, F^{ext} and F^{int} are the external and internal forces,
 158 respectively.

159 2.3 Numerical procedures

160 The explicit integration scheme was adopted to solve the governing equations. The history-dependent
 161 information carried by particle p are: position X_p , mass m_p , volume V_p , density ρ , velocity v_p ,
 162 deformation rate D_p , vorticity W_p , stress σ_p , and external force f_p^{ext} . The governing equations (3) and
 163 (4) are solved on element nodes in terms of variables interpolated from the particles, i.e. mass m_i ,
 164 velocity v_i , momentum M_i , acceleration a_i , internal force F_i^{int} , external force F_i^{ext} , normal direction
 165 ω_i^{norm} and tangential direction ω_i^{tang} , where the subscript i represents the node number. For the soil-
 166 structure interaction problems, the structure is simplified as a rigid body. The main functions within
 167 each incremental step are:

168 (i) Initialisation of nodal variables. The time step always starts with the initialisation of the nodal
 169 variables of the structure and soil, which will be automatically abandoned at the end of the step.

170 (ii) Interpolation from particles to nodes. The masses and momenta of the associated particles
 171 (inherited from the previous incremental step) are interpolated to the nodes

$$172 \quad m_i = \sum_p S_{ip} m_p \quad (5)$$

$$173 \quad M_i = \sum_p S_{ip} m_p v_p \quad (6)$$

$$174 \quad \omega_i^{\text{norm}} = \frac{\sum_p \nabla S_{ip} m_p}{\left\| \sum_p \nabla S_{ip} m_p \right\|} \quad (7)$$

175 where S_{ip} and ∇S_{ip} are the shape function and its gradient at node i evaluated at particle p ,
 176 respectively^[40]; \sum_p represents the summation over all related particles. The derivation of the normal
 177 direction ω_i^{norm} in Eq. (7) can be referred to in Reference [52-53]. For the soil, the internal force is
 178 obtained

$$179 \quad F_i^{\text{int}} = -\sum_p \nabla S_{ip} \sigma_p V_p \quad (8)$$

180 The tractions on the Neumann boundary is calculated^[56-57]

$$181 \quad F_i^{\text{ext}} = \sum_p S_{ip} f_p^{\text{ext}} V_p \quad (9)$$

182 (iii) Calculate nodal velocities and accelerations. The velocities and accelerations on the background
 183 mesh can be obtained. At the commencement of the incremental step, the velocity of the node is

$$184 \quad v_i = \frac{M_i}{m_i} \quad (10)$$

185 The acceleration for the soil node from the internal and external forces can be calculated from the
 186 governing equation as

$$187 \quad a_i = \frac{F_i^{\text{int}} + F_i^{\text{ext}}}{m_i} \quad (11)$$

188 Then the nodal velocity is updated as

$$189 \quad v'_i = v_i + a_i \Delta t \quad (12)$$

190 where Δt is the time increment and determined through the Courant–Friedrichs–Lewy stability
 191 condition

$$192 \quad \Delta t = \frac{\phi h}{\sqrt{(\lambda + 2G)/\rho}} \quad (13)$$

193 where ϕ is the Courant number, h is the size of the square element, and G and λ are the Lamé’s
 194 parameters.

195 For the soil node in contact with a structure moving with a prescribed velocity v_0 , v'_i is further
 196 adjusted depending on the adopted contact algorithm ‘Geo-contact’^[50]. The soil may be in contact
 197 with the structure if the soil mass projections are non-zero within the predefined area of the structure.
 198 For a specific node i of the soil in contact, its normal relative velocity to the structure is
 199 $\Delta v_i^{\text{norm}} = (v'_i - v_0) \phi_i^{\text{norm}}$, with v_0 as the velocity of the structure. Node i of the soil can be distinguished as
 200 approaching or departing from the structure with the relative normal velocity

$$\begin{aligned} \Delta v_i^{\text{norm}} &> 0, \text{ approach} \\ \Delta v_i^{\text{norm}} &< 0, \text{ depart} \end{aligned} \quad (14)$$

201 The normal contact strategy between the soil and the structure is realised by adjusting the normal
 202 relative velocity by $\Delta v_i^{\text{norm},*}$: (i) for soil node i approaching the structure, the normal relative velocity
 203 is eliminated; and (ii) for soil node i departing from the structure, the normal relative velocity is
 204 eliminated only if no separation between the structure and the soil is considered (otherwise, the
 205 normal relative velocity is maintained).

206 The relative tangential velocity of the soil node i to the structure is

$$\Delta v_i^{\text{tang}} = (v_i' - v_0) \omega_i^{\text{tang}}$$

$$\omega_i^{\text{tang}} = \omega_i^{\text{norm}} \times \frac{(v_i' - v_0) \times \omega_i^{\text{norm}}}{|(v_i' - v_0) \times \omega_i^{\text{norm}}|} \quad (15)$$

207 where function ‘ \times ’ represents the cross product. The shear along the interface is governed by the
 208 Coulomb friction law, i.e. the adjusted tangential relative velocity $\Delta v_i^{\text{tang},*}$ is bounded by $\mu_c \Delta v_i^{\text{norm},*}$,
 209 in which μ_c is the Coulomb friction coefficient. In geotechnical applications involving soils with low
 210 permeability, a threshold value of the friction stress is usually applied for total stress analyses under
 211 undrained conditions

$$\tau = \alpha s_u \quad (16)$$

212 where τ is the maximum shear stress along the interface and α is the limiting shear stress ratio, ranging
 213 from 0 to 1. So the tangential relative velocity will be adjusted by

$$\Delta v_i^{\text{tang},*} = \min \left(\Delta v_i^{\text{tang}}, \mu_c \Delta v_i^{\text{norm},*}, \frac{\alpha s_u A_i \Delta t}{m_i} \right) \quad (17)$$

214 where A_i is the interface area represented by node i .

215 A penalty factor β_i is then introduced to the overall adjustment of the relative velocity Δv_i^{adju} to obtain
 216 a smooth reaction force

$$\Delta v_i^{\text{adju}} = \beta_i \left(\Delta v_i^{\text{norm},*} + \Delta v_i^{\text{tang},*} \right)$$

$$\beta_i = 1 - \left(\frac{\min(s_i, h)}{h} \right)^k \quad (18)$$

217 where s_i is the distance from node i to the surface of the structure and k is the penalty power. The total
 218 contact force on the structure is

$$P = \sum_i \frac{m_i \Delta v_i^{\text{adju}}}{\Delta t} \quad (19)$$

219 The new velocity of node i is $v_i^{\text{new}} = v_i' + v_i^{\text{adju}}$. Roller (Neumann) boundary condition can be imposed
 220 by removing the new nodal velocity normal to the boundary. Then, the overall acceleration for the
 221 current time step at soil node i is

$$a_i^{\text{new}} = \frac{v_i^{\text{new}} - v_i}{\Delta t} \quad (20)$$

223 (iv) Update particle state. The strains of the soil particles are calculated with the deformation gradient
 224 using an updated formulation

$$225 \quad F_p^{\text{new}} = f_p F_p \quad (21)$$

226 where f_p is the relative deformation gradient

$$227 \quad f_p = I + \sum_i \nabla S_{ip} v_i^{\text{new}} \quad (22)$$

228 with I indicating the identity matrix. The stresses and material properties of the soil particles are
 229 calculated using an elastic-perfectly plastic constitutive model with the deformation rate D_p and
 230 vorticity W_p

$$231 \quad D_p = \frac{1}{2} \left[\sum_i \nabla S_{ip} v_i^{\text{new}} + \left(\sum_i \nabla S_{ip} v_i^{\text{new}} \right)^T \right] \quad (23)$$

$$232 \quad W_p = \frac{1}{2} \left[\sum_i \nabla S_{ip} v_i^{\text{new}} - \left(\sum_i \nabla S_{ip} v_i^{\text{new}} \right)^T \right] \quad (24)$$

233 where the superscript T means the transposition of a tensor. The definition of the stresses follows
 234 finite strain theory taking account of the incremental rotation of the configurations between time steps
 235 for objectivity, the trial stresses being measured with the Cauchy stress and updated with the Jaumann
 236 rate according to

$$237 \quad \sigma_p^{\text{trial}} = \sigma_p + \Delta t \left[\left(\sigma_p W_p - W_p \sigma_p \right) + C D_p \right] \quad (25)$$

238 where C is the fourth-order stiffness tensor. The trial Cauchy stresses should satisfy the von Mises
 239 criterion

$$240 \quad f = \sqrt{2J_2} - \sqrt{\frac{2}{3}} s_u \leq 0 \quad (26)$$

241 where J_2 is the second deviatoric stress invariant. Otherwise, the trial Cauchy stresses will be updated
 242 with radial return mapping as the Mises yield surface is circular in the π plane.

243 In addition, the velocities and positions are updated by mapping the nodal accelerations and velocities

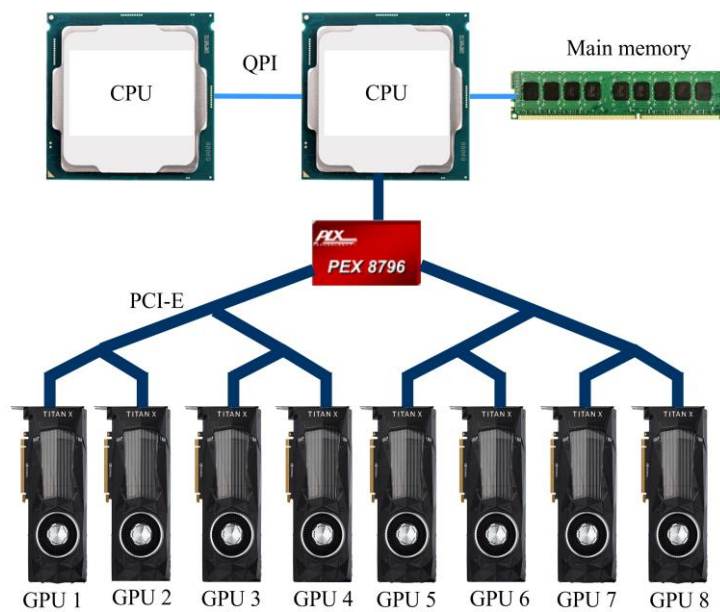
$$244 \quad v_p^{\text{new}} = v_p + \sum_i S_{ip} a_i^{\text{new}} \Delta t \quad (27)$$

$$245 \quad X_p^{\text{new}} = X_p + \sum_i S_{ip} v_i^{\text{new}} \Delta t \quad (28)$$

246 For the structure moving with a prescribed velocity v_0 , its velocity is unchanged and the new position
247 is updated by addition with $v_0\Delta t$.

248 3. Multiple-GPU platform

249 Expansion from a single-GPU into multiple-GPU parallelisation can be categorised into two
250 directions: within a single computer and across multiple computers^[48]. This paper focuses on the
251 former (Figure 1). A GPU has different memory hierarchies, such as the register, texture, constant,
252 shared, local and global memories. The global memory, the largest memory on each GPU, is the main
253 space to save the variables in the calculations. Access to the global memory from the multiprocessors
254 needs to be coalesced on aligned contiguous memory addresses to achieve a high bandwidth. The
255 GPUs are plugged on the PCI-E slots on the motherboard of the host computer and connected with
256 the shared RAM via quick-path interconnect (QPI), shaping a shared-distributed hybrid memory
257 hierarchy. Computational tasks shared between the GPU devices necessitate data manipulations
258 between the RAM and the GPU memories, which is through the PCI-Express bus and controlled by
259 the PCI-E controller element on the CPU.



260

261 Figure 1 Configuration of a multiple-GPU system with a single root complex (Reference:
262 <https://www.servethehome.com/single-root-or-dual-root-for-deep-learning-gpu-to-gpu-systems/>)

263 Conventional communication between the GPUs was traditionally in procedures with MPI and
264 CUDA commands^[48]: (i) read the information on the global memory of the local GPU; (ii) copy to
265 the RAM on the computer via PCI-E (also implicitly via CPU and QPI); (iii) write to the global
266 memory of the remote GPU from the shared RAM. Overhead on the communication is often the
267 bottleneck of parallelisation since different components are involved, especially for frequent

268 synchronisation between the GPUs in each incremental step. Also, the hybrid MPI-CUDA
269 environment causes extra complexities to the coding and increases the risk of failure in the execution.
270 In modern PCI-E architectures, GPUs connected to the same PCI-E root are allowed to access the
271 global memory of each other with the P2P communication technique without using the RAM as a
272 transit storage. Therefore, MPI send/receive manipulations as in Reference [48] can be avoided with
273 more conveniences and higher performance.

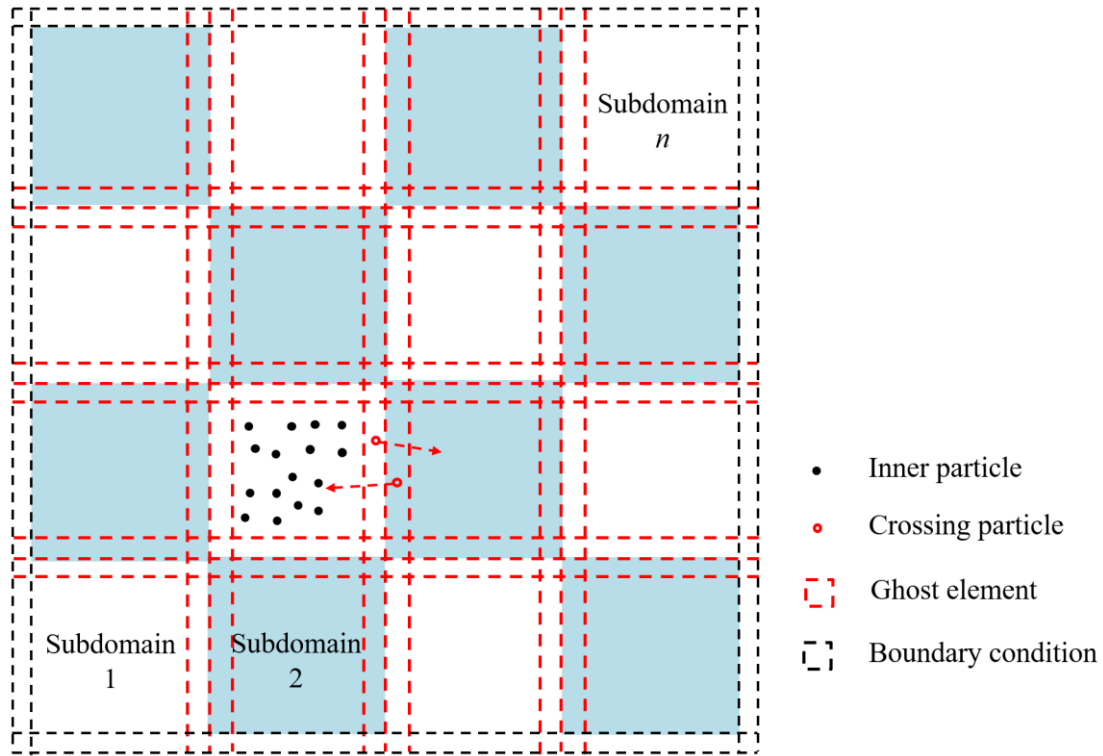
274 **4. Parallelisation of MPM**

275 Parallelisation of the in-house program MPM-GeoFluidFlow was specially tuned for a single-GPU^[46]
276 and a multiple-computer multiple-GPU frameworks^[48]. In this work, the program was parallelised on
277 the platform of a one-computer multiple-GPU system with a single-root PCI-E complex (Figure 1).
278 The parallelisation was performed purely within the CUDA environment without the MPI interface
279 required in the conventional multi-GPU parallelisation^[46, 48]. Improvements were made on the data
280 transport scheme between the shared RAM and the member GPUs, which was boosted by taking
281 advantage of the P2P technique supported on the single-root complex PCI-E architecture.
282 Parallelisation of the function ‘Interpolation from particles to nodes’ was optimised with an enhanced
283 ‘Particle-List’ scheme, which is parallelised on the GPUs. As a result, the MPM computation was
284 further accelerated and the complexity of the code was reduced. The original two dimensional
285 framework was extended to three dimensional.

286 **4.1 Task distribution and assembly**

287 In the pre-processing stage, the whole domain of the task is assigned over the shared RAM on the
288 computer, with the material discretised into a cloud of particles and a structured mesh constructed.
289 The history-dependent information carried on the particles and the temporary variables for the
290 element nodes are declared. Then the computational domain is evenly decomposed into a number of
291 subdomains to distribute the entire task onto the individual GPUs (Figure 2). The number of elements
292 in each subdomain is around M/n , where M is the total number of elements in the whole domain and
293 n the total number of subdomains. The discretised particles are associated with the background
294 subdomain by their locations. Variables of the particles and the element nodes in each subdomain are
295 copied from the computer RAM to the global memory of each GPU through the PCI-E and QPI.
296 Therefore, the space of the shared RAM should be larger than the total size of the global memories
297 of the hosted GPUs. Two additional layers of ghost elements in each direction are generated out of
298 the computational subdomain, which is to maintain the continuity of the information at the border of
299 the subdomain. The roller (Neumann) boundary condition is implemented on the outer boundaries of
300 the outer subdomains. Calculations of the MPM algorithm within each subdomain is parallelised on

301 the corresponding GPU within the CUDA environment. The interpolated information on the
 302 overlapping ghost element nodes from the neighbouring subdomains will be added and saved on both
 303 sides into the GPU global memories. Therefore, the neighbouring subdomains are essentially
 304 boundary conditions to each other. After the computation, the information of the particles in the
 305 subdomains is re-assembled into the RAM and will be used for post-processing.



306

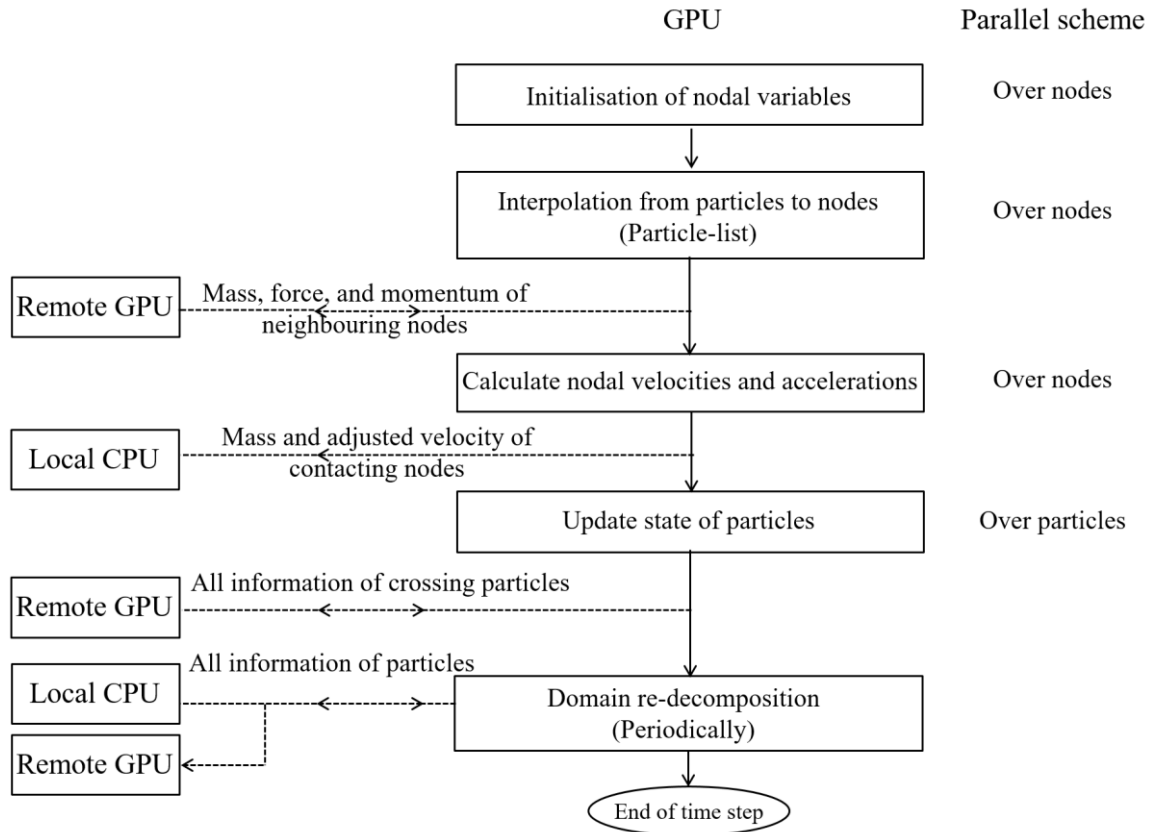
307 Figure 2 Domain decomposition in multiple-GPU parallelisation

307

308 4.2 Parallelisation on each GPU

309 In order to control the overhead of the frequent data transfer between the RAM and the GPUs, all the
 310 variables of the particles and nodes required in the essential calculations of Eqs. (5) – (28) are reserved
 311 on the global memory of the GPUs, and are accessible from the active multiprocessors. The functions
 312 ‘Initialisation of nodal variables’ and ‘Calculate nodal velocities and accelerations’ (Eqs. (10) – (20))
 313 are parallelised over the nodes, i.e. updating the information of one node on one GPU core (Figure
 314 3). In the thread for node i , only the information of the node is involved in the calculations as in Eqs.
 315 (10) – (20), which means the parallel computations in the GPU cores are independent to each other.
 316 Therefore, the two functions present relatively high parallelisability and are straightforward to
 317 parallelise over the nodes. Particularly for the soil-structure interactions by Eqs. (14) – (18), the
 318 velocity adjustment on each node, controlled by its kinematic state relative to the structure, is
 319 independent to each other and can be parallelised over the nodes. As a whole system, the total reaction
 320 force on the structure (Eq. (19)) will be collected from the contacting nodes on the GPUs and

321 superimposed on the CPU. For the function ‘Update particle state’ (Eqs. (21) – (28)), the workload is
 322 decomposed over the particles, for each of which the interpolated information are superimposed from
 323 the surrounding nodes sequentially. Different to the writing operations, which may induce data race
 324 in the memory and will be discussed later, reading data from the identical memory address by
 325 different threads is allowed in the parallelisation. Furthermore, a high bandwidth can be achieved
 326 when the multiprocessors reading the consecutive memory addresses. Hence, the function ‘Update
 327 particle state’ is expected to have a high acceleration effect by the parallelisation.

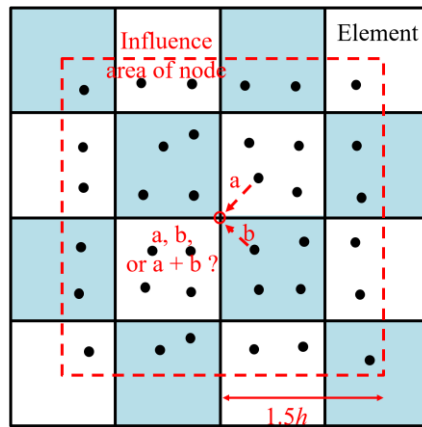


328

329 Figure 3 Essential operations in multiple-GPU parallelisation of MPM

330 In contrast, the function ‘Interpolation from particles to nodes’ is more difficult to parallelise. If the
 331 function is simply parallelised over the particles, i.e. the interpolation from each individual particle
 332 to its related nodes is configured on a thread, data race can be induced by concurrently writing the
 333 interpolation outcomes from different GPU cores into identical addresses of the global memory
 334 (Figure 4a). The data race makes the final writings to the memory unpredictable, which means the
 335 interpolations would be erroneous. Reference [47] avoided the data race by using the atomic
 336 operations, which is often supported in the modern GPUs. However, the atomic operation is not
 337 recommended as it is essentially a sequential writing action to the memory, which undermines the
 338 overall acceleration effect. In Reference [49], a special technique ‘Particle-Grid offset’ was developed
 339 to tackle the data race; however, a series of complicated operations are involved to make the algorithm

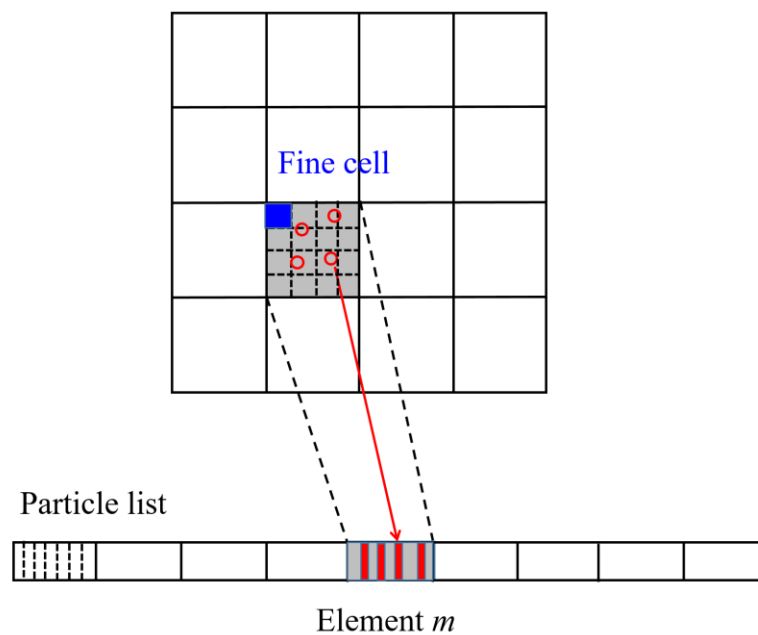
340 parallelisable. In Reference [46] and [48], the problem was solved by parallelising the function over
 341 the nodes, for each of which the interpolation outcomes from associated particles were superimposed
 342 sequentially. Before the parallel computation, a particle list for each node needs to be constructed,
 343 which was saved on the RAM and was transported to the GPU periodically. Comparing with the
 344 ‘Particle-Grid offset’ technique, the ‘Particle-List’ scheme seems to be more accessible as only two
 345 steps (generation of particle list and interpolation operation) are required. However, the updating
 346 frequency of the particle list, often determined by experience and through trial calculations, depends
 347 on the mesh size and the intrinsic characteristics of the problem to be analysed and may be very high
 348 (such as once for five incremental steps) for impact problems. In three-dimensional analysis, each
 349 node may be associated with around 216 particles ($6 \times 6 \times 6$; Figure 4a) as regulated by the uGIMP
 350 shape function^[40, 54]; therefore, a particle list relating the surrounding particles for all the nodes would
 351 exhaust the memory space on the GPU, which is typically upper bounded by 24 GB.



352

353


(a) Data race



354

355

(b) Particle list



(c) Parallelisation by layers

Figure 4 Parallelisation of the function ‘Interpolation from particles to nodes’

356

357

358

359

360

361

362

363

364

365

366

367

368

369

370

371

372

373

374

375

376

377

378

379

In this work, the ‘Particle-List’ scheme was improved by generating the particle list for each element rather than the node, which means the total memory requirement by the particle list is $O(N_p)$ and much smaller than that by the original scheme $O(216N_i)$, where N_p and N_i are the total number of the particles and nodes, respectively. The generation of the particle list is parallelised with the GPUs instead of the CPU sequential operations in the original scheme. A fully engaged element often accommodates 4 or 16 particles^[40, 41]. To avoid the potential data race, an expanded particle list of each element is adopted, which was developed from a similar method used in the creation of contact pair list in the discrete element method^[58]. The original elements are evenly divided into a number of finer cells (Figure 4b). Each fine cell corresponds to a specific memory address of the particle list, and accommodates only one particle or less^[59]. Then, the particle sorting operations can be parallelised across the GPU cores over the particles, in which each memory address is written with only one or less particle ID without the risk of data race. The number of the fine cells in each element, dependent on the smallest distance between the particles controlled by the strains, can be determined through trial calculations and $8 \times N_{PPC}$ is often acceptable, where N_{PPC} represents the particle number in per element. For specific problems with extreme strains of material, the number of the fine cells can be increased and a novel technique in Reference [60] reseeding the particles in the elements is also suggested. The particle list is sparse due to the large number of empty fine cells out of the engaged ones (Figure 4b), which still means a heavy memory requirement. An additional compression step is then performed to obtain a dense particle list. The enhanced ‘Particle-List’ scheme is fully performed on the GPUs and need no data transport between the RAM and the GPUs. Therefore, improvement of the speedup of the parallelisation is expected with the enhanced scheme.

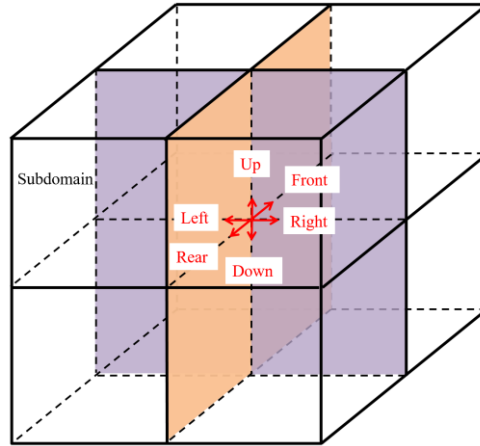
380 The interpolations from the particles to nodes are then parallelised over the nodes. The particles on
381 the particle list of the surrounding elements within the influence range of each node are involved in
382 the interpolation on one GPU thread. In consideration of the heavy requirement of memory space for
383 the particle list of large-scale problems, the generation of particle list and the interpolations are
384 performed in layers of nodes (Figure 4c).

385 **4.3 GPU communications**

386 In order to keep the continuity of the stress and strain field between the connected subdomains, the
387 interpolated variables on the border nodes within the neighbouring subdomains, including the mass,
388 momentum and internal force, by the function ‘Interpolation from particles to nodes’ is superimposed
389 at each incremental step (Figure 3). The operations are implemented with the direct P2P operations
390 between the GPU memories for the single-root complex PCI-E framework, which owns much higher
391 bandwidths (5 GB/s) than that used in Reference [48] through the RAM and the infiniband network
392 (1.25 GB/s) within MPI environment (Figure 1). The send and receive operations can be performed
393 in bi-direction with a modern PCI-E, which doubles the intrinsic bandwidth of the data migration to
394 10 GB/s. Furthermore, the transfer of data between the GPUs can be hidden by the essential
395 calculations on each GPU, which means overhead on the synchronisation process is virtually nil and
396 a perfect scaling may be possible. In comparison, the data transfers within MPI environment accounts
397 for about 5% of the total computational effort^[48]. The exchanged information of the ghost element
398 nodes from the neighbouring GPUs is added to the counterparts on the local GPU (Figure 2).

399 The particles may move across the subdomains and hence migrate from a GPU to its neighbour, for
400 which all the information of the migration particles should be transferred to the new subdomain with
401 P2P operations similar to that exchange the neighbouring nodal information. The subdomains
402 communicate with their neighbours in six directions (left-right; front-rear; up-down) (Figure 5a). For
403 most ghost elements (Figure 5b) and particle migrations (Figure 5c), send/receive operations are
404 performed between two subdomains in two directions; there are also some corner ghost elements are
405 shared by four or eight subdomains and some particles moves to unconnected subdomains (such as
406 upper-right subdomain), which need more than one synchronisation step (Figure 5). The particle
407 migration can be time-consuming if performed in every incremental step. If a particle moves from
408 the current subdomain to the neighbouring one between the particle migration operations, the
409 interpolations from the particle are allocated to the ghost element nodes of the current subdomain,
410 which will be synchronised between the neighbouring subdomains as described previously. Therefore,
411 the interpolation results are accurate if only the particles do not move across the outer layer of ghost
412 elements of the current subdomain. The one layer of ghost elements essentially functions as a
413 buffering zone of the particle migration between the subdomains. Therefore, the particle migration is


414 conducted for once in a number of steps in this study, which is controlled by the time of particles
 415 move across the outer ghost element and depends on the mesh and particle discretization. In the
 416 calculations, the number of step to migrate the particles can be determined by an experimental
 417 computation, which is often selected as 5 in this study.



418

419

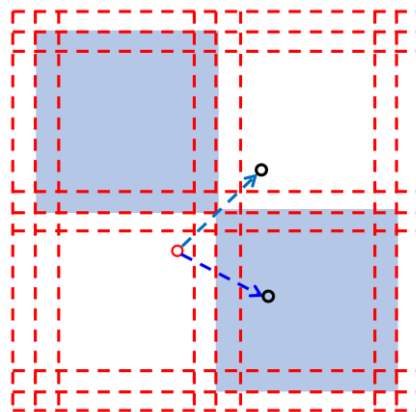
(a) GPU communications between subdomains



420

421

(b) Corner ghost elements



422

423

(c) Particle migration

424

Figure 5 GPU communications

425 In the calculations, void areas tend to be formed in the subdomains due to the large deformation of
426 materials, e.g. in the mini-slump test in Section 5.1, which idle the computational resource. The
427 migration of the particles between the subdomains may also undermine the workload balance among
428 the GPUs. A simple procedure of domain re-decomposition is then performed to updates the
429 subdomain dimensions periodically at intervals of a large number of incremental steps, such as 50,000.
430 The communication overhead for the neighbouring nodes and the migration particles between the
431 subdomains may be non-ignorable but acceptable in many cases, which can be compensated by the
432 substantial calculations in each GPU. The information of all the particles in the subdomains are
433 gathered from the GPUs' global memories to the shared RAM of the computer. The upper and lower
434 boundaries of the material are derived from the particle coordinates, based on which the boundaries
435 of the computational domain in the following steps are updated. Then the task distribution operations
436 in Section 4.1 is re-performed by evenly decompose the computational domain. The present domain
437 de-composition procedure is mainly to solve the problem of void areas, and somehow mitigate the
438 workload imbalance between the GPUs. More advanced technique with adaptive subdomain sizes to
439 balance the workloads between the GPUs will be developed in the future work.

440 **5. Performance assessment**

441 The accuracy and convergence for the standard MPM algorithm with explicit calculations have been
442 assessed in Reference [32, 35], which may be inherited here due to the trivial modifications. The
443 benchmark cases, mini-slump test and cone penetration test, were simulated to assess the acceleration
444 effect of the multiple-GPU parallelisation scheme. The parallel computations were performed on a
445 single-computer server, which hosts 8 NVIDIA Titan Xp GPUs based on a single-root complex
446 (Figure 1) and 2 Intel Xeon E5-2687WV4 CPUs. Each CPU has 12 cores with frequency of 3.0 GHz;
447 on each GPU, a total number of 3840 cores are accessible and the dedicated global memory is 12 GB;
448 The RAM space of the server is 256 GB. The operating system was Ubuntu 18.04, the C++ compiler
449 was gcc 5.3.0 and the GPU compiler was CUDA v8.0.44. All the computations were based on double-
450 precision numbers to guarantee the accuracy.

451 The soil was considered to be an elastic-perfectly plastic material and regulated with the von Mises
452 yield criterion. The Poisson's ratio of the soil was selected as 0.49. The time step was calculated by

$$453 \quad \Delta t \leq \frac{\alpha d}{\sqrt{(\lambda + 2G)/\rho}} \quad (29)$$

454 where G is the shear modulus of soil, λ is the Lamé constant, d represents the mesh size, and α is the
455 Courant number. The 'speedup' factor was to characterise the acceleration effect of the parallel

456 computations: $\text{Speedup} = T_{\text{Sequential}} / T_{\text{Parallel}}$, where $T_{\text{Sequential}}$ and T_{Parallel} are the runtimes of the CPU
457 sequential and GPU parallel calculations over a number of incremental steps (Appendixes A and B).

458 **5.1 Mini-slump test**

459 Submarine landslides are known as one of the most hazardous threat to the submarine structures,
460 featuring enormous volumes of sediments running at very high velocities and reaching very far runout
461 distances before final deposition^[61]. The sliding behaviour of the soil can be studied in laboratory
462 through the mini-slump test, in which the soil is released from a cylinder and then runout along a flat
463 base. The test performed by Reference [62] with the remoulded soil from Heimdal, Norway was
464 simulated using the MPM (Figure 6a). The cylinder had a height H of 120 mm and a diameter of 100
465 mm. The mechanical behaviour of the soil was considered by a Herschel-Bulkley (H-B) model^[63]

$$466 \quad s_u = s_{u0} + K\gamma^n \quad (30)$$


467 where s_u is the undrained shear strength of the soil, s_{u0} is the threshold shear strength $s_{u0} = 200$ Pa, γ
468 the shear strain rate, K the consistency coefficient and n the shear-thinning index. In the experimental
469 test, the parameters in Eq. 2 were determined as $s_{u0} = 200$ Pa, $K = 15$ Pa·s ^{n} , and $n = 0.35$. In the MPM
470 analysis, the mesh size d was selected as $H/60$, which was validated to be sufficiently fine in viscous
471 sliding problems of soil as it presents similar results with a finer mesh $H/120$ ^[18, 41]. The Young's
472 modulus was 500 times the undrained shear strength s_u . In total, there were 393,216 slurry particles.
473 The flat base was assumed as a no-slip boundary. The time step Δt was determined with a Courant
474 number of 0.3. Another numerical simulation using the computational fluid dynamics (CFD) method
475 was also performed by Reference [64]. Figure 6b shows the final morphologies of the soil predicted
476 by the CFD and MPM analyses, which matches well with each other. The slumped width of the soil
477 was 0.12 m in the experiment and the final height was 0.04 m, which are close to that of the numerical
478 predictions. The profile of the soil at 0.3 s in the MPM simulation was shown in Figure 6c.



479

480

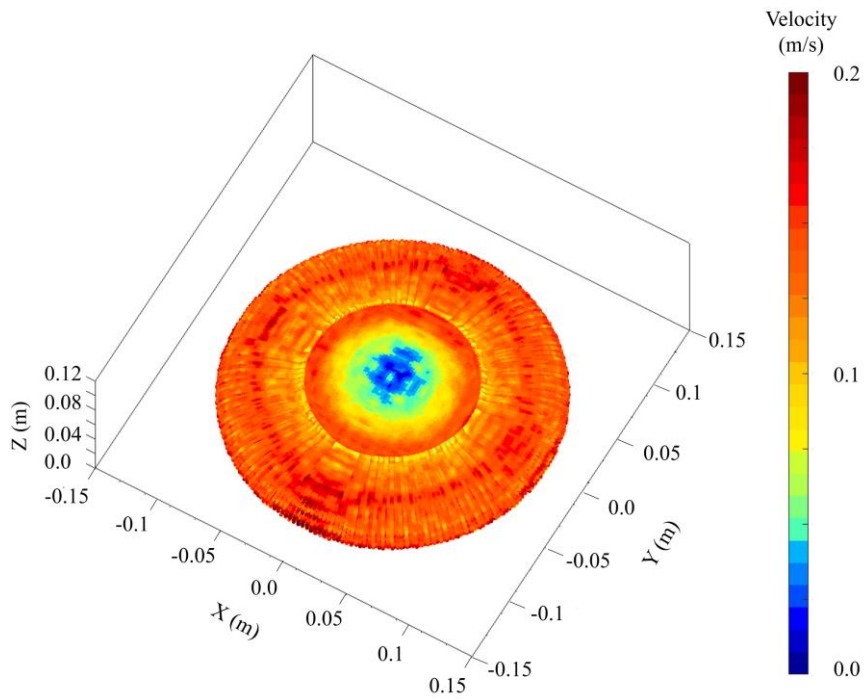
(a) Laboratory test^[65]



481

482

(b)



483

484

(c) 0.3 s

485

Figure 6 Mini-slump test and runout morphologies

486

487

488

489

490

491

492

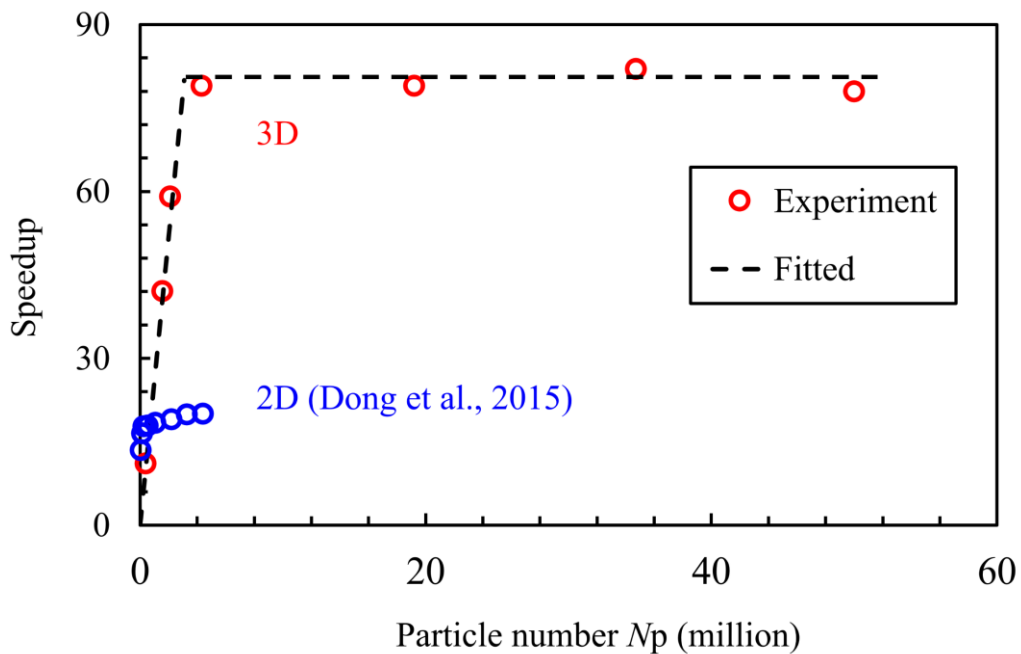
493

To investigate the acceleration effect of the GPU parallel framework, the particle number N_p was increased from 393,216 to 1,572,864, 2,102,203, 4,343,424, 19,232,055, 34,744,320, 50,035,200 and 121,065,216. The total memory space engaged in the simulation is proportional to the particle number, which increases from 0.08 GB (393,216 particles) to 23.5 GB (121,065,216 particles). In Reference [46], the maximum particle number for a 2D model with a memory size of 4 GB was 50,035,200, and will be much less for a 3D model than the counterpart in this study. The reason is that the size of the particle list for the element nodes in Reference [46] was nearly 25 times of the total particle number, which is avoided in this study by establishing a particle list for the elements in each layer. Therefore,

494 the maximum particle number accommodated in 8 GPUs can be up to 400,035,200 as the concern of
495 many large-scale geotechnical problems. A total runtime within 100 incremental steps was recorded
496 for each case with the CPU sequential and GPU parallel simulations, in which the acceleration effect
497 by the GPU parallel strategy is clearly demonstrated (Table 1). Within the CPU sequential
498 calculations, the runtime is linearly proportional to the scale of the case in terms of particle number.
499 The speedup linearly increases for less than 4,000,000 particles (Figure 7); if the computational scale
500 is enlarged further with more particles, the GPU seems to be fully loaded, and the acceleration effect
501 presents a good scaling behaviour and converges to an average speedup of about 80. The average
502 speedup of ~ 80 in this study is much higher than that in Reference [35] of around 20 for two main
503 reasons: the GPU Titan Xp in this study outperforms that GTX 780M in Reference [46]; the
504 parallelisation schemes are optimised in this study in terms of particle list and memory access. Among
505 all the functions of the MPM, the function ‘Interpolation from particles to nodes’ consumes the most
506 computational efforts for around 70%, which is due to the non-coalesced memory access when
507 writing the interpolations to the nodal addresses. The overhead on the establishment of the particle
508 list takes less than 2% of the total runtime, much less than that in Reference [46] and [48] of around
509 15%, as the operations are fully moved and parallelised onto the GPU. The remaining 28% of the
510 computations is mainly on the function ‘Update particle state’, while that on the functions
511 ‘Initialisation of nodal variables’ and ‘Calculate nodal velocities and accelerations’ are ignorable.
512 Specifically, the average speedup for the function ‘Interpolation from particles to nodes’ is around
513 65, while it increases to about 170 for the functions ‘Initialisation of nodal variables’ and ‘Calculate
514 nodal velocities and accelerations’. Therefore, the function ‘Interpolation from particles to nodes’ is
515 the bottleneck of the parallelisation of the MPM. Additional calculations were performed for the case
516 with 4,343,424 particles by using naïve atomic operations to parallelise the function ‘Interpolation
517 from particles to nodes’, which presents very low speedup of less than 5 due to the heavy writing
518 conflicts between the neighbouring particles. Further experiments show that the speedup with the
519 naïve atomic operations is not stable, which varies with data structure of the particles. Due to the very
520 low efficiency, atomic operations are not recommended in many parallel algorithms.

521 The acceleration effect of the multiple-GPU parallelisation over the CPU sequential computations is
522 shown in Table 1 and Figure 8. The total runtime for 100 incremental steps rather than one step was
523 recorded for each case to avoid the random errors during recording. For example, in the case with
524 393,216 particles, the average runtime for the cases with GPU is smaller than 0.2 s, which will be
525 significantly influenced by a small random error. The average runtime for an incremental step can be
526 easily obtained from the total value for 100 steps. The performance of the GPU is maximised when
527 it is fully loaded with workload. Given one or two GPUs are invoked, the GPU seems to be fully


528 loaded with 4 million particles: for the cases with < 4 million particles on each GPU, the speedup
529 increases with the particle number when using an identical number of GPUs. The maximum speedup
530 with less than two GPUs is around $80N_{\text{gpu}}$ with N_{gpu} as the number of GPUs, which is similar to the
531 average speedup predicted previously with one GPU. When four or eight GPUs are mobilised, the
532 GPU is not fully loaded even with 30 million particles on each GPU. The reason is not very clear but
533 is inferred to be related to the scheduling elements in the CPU. The overall speedup with less than 8
534 GPUs is fitted as $71.5N_{\text{gpu}}$. Although it is not meaningful to compare the speedups of the multiple-
535 GPU parallel schemes based on different hardware and software, the speedups of about $59N_{\text{gpu}}$ and
536 $110N_{\text{gpu}}$ in Reference [48] and [49] were obtained. In consideration of the convenient implementation
537 and high reliability of the present parallel framework, the speedup of $71.5N_{\text{gpu}}$ is quite satisfactory.
538 Due to the optimisation of communication between the neighbour domains as described in Section
539 3.4, overhead on the synchronisation of the ghost nodal information, the particle migration, and the
540 domain re-decomposition is ignorable with less than 2% when comparing to that on the essential
541 computations. Therefore, the parallel framework of the MPM with multiple-GPU in this study
542 presents good behaviour in terms of hardware communications, which has been the common problem
543 of many numerical methods. Also, the advantage of the P2P technique based on the single-root
544 complex structure of the PCI-E is well presented.



545

546

Figure 7 Speedups of GPU parallelisation with one GPU



547

548

Figure 8 Speedups of GPU parallelisation with multiple-GPU

549

Table 1 Speedups of GPU parallel simulations in 100 incremental steps for mini-slump test cases

N_p	Memory size (GB)	CPU sequential	1 GPU		2 GPUs		4 GPUs		8 GPUs	
		Runtime (s)	Runtime (s)	Speedup	Runtime (s)	Speedup	Runtime (s)	Speedup	Runtime (s)	Speedup
393,216	0.08	180	16	11	17	11	18	10	19	9
1,572,864	0.3	760	18	42	18	42	19	40	19	40
2,102,203	0.5	921	16	59	16	59	16	59	15	61
4,343,424	0.9	1,890	24	79	20	95	20	95	19	99
19,232,055	3.2	9,552	121	79	61	157	51	189	35	276
34,744,320	6.4	19,100	232	82	127	150	85	225	55	347
50,035,200	10.2	27,400	350	78	171	160	118	232	75	363
121,065,216	23.5	78,700	—	—	—	—	261	301	137	572

550

551

5.2 Cone penetration test

552

553

554

555

556

557

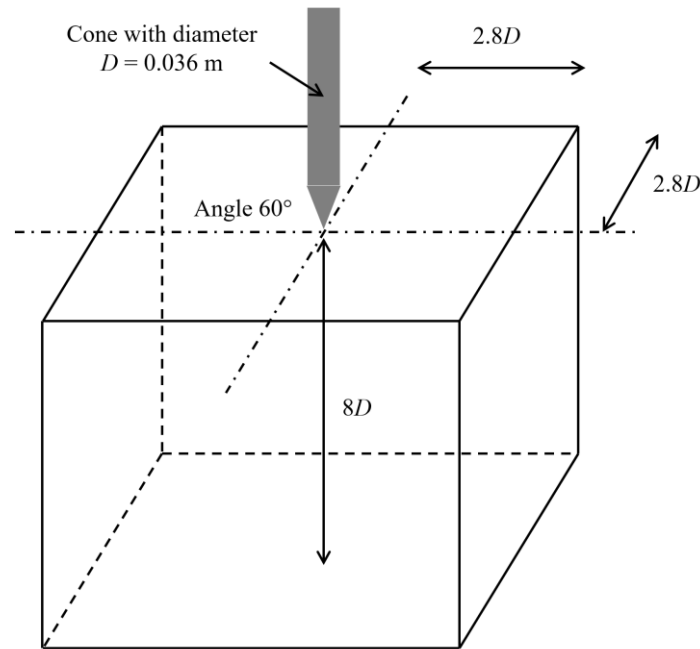
558

559

560

The cone penetrometer has been considered as the most widely used in-situ geotechnical instrument to obtain the sequence and the physical and mechanical properties of the subsurface strata. For the cone penetrated in pure clays, the penetration resistance is related with the undrained shear strength of the soil s_u through the calibration of a bearing factor N_k , which was often investigated with theoretical, experimental and numerical analyses^[65-67]. The numerical model of a cone penetration test used in Reference [68] with a large deformation finite element (LDFE) method was duplicated in this study. The standard cone had a diameter of $D = 35.7$ mm and its tip had an angle of 60° , as shown in Figure 9. Quarter of the model was simulated by taking advantage of its symmetry to save the runtime. In Reference [69], a smaller model with a wedge of 20° in the rotational direction was

561 simulated, which was also proven to be accurate to represent the full model of the test. Herein, the
 562 parallel efficiency is the main concern, therefore, the quarter model was used and the wedge model
 563 will be adopted in the future applications. The chamber extensions on the horizontal and vertical
 564 directions were $2.8D$ and $8D$, respectively. The mesh size $d = D/36$, which is satisfactorily fine to
 565 achieve a convergent prediction of N_k . In each element fully occupied by the soil, 2×2 particles were
 566 configured prior to the calculation. In total, 24,000,000 soil particles were discretised. The cone was
 567 assumed to be rigid and smooth. The penetration speed of the cone was taken as $2.8D/s$ (0.1 m/s),
 568 which was verified as sufficiently low to use the dynamic formulation to simulate the quasi-static
 569 process^[69]. The submerged density of the soil was 1500 kg/m^3 . The geostatic stresses induced by the
 570 self-weight of soil were not considered. The clay had a uniform shear strength of $s_u = 10 \text{ kPa}$ and a
 571 soil rigidity index $G/s_u = 100$.



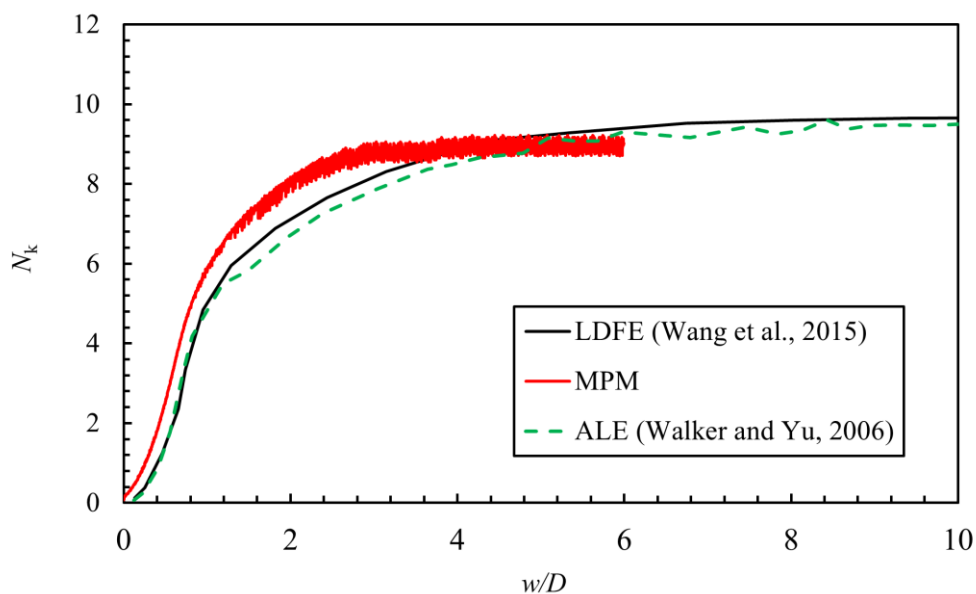
572
 573 Figure 9 Setup of the cone penetration test

574 Profile of the bearing factor N_k is plotted versus the normalised penetration depth w/D as shown in
 575 Figure 10, in which w is the penetration depth of the pipeline. The bearing capacity increases with
 576 the penetration of the cone, which stabilises at about 8.97 with $w/D = 3$. The profile obtained from
 577 the MPM analyses has some high-frequency fluctuations of around 1.5% due to the particles below
 578 the cone crossing the element boundaries in the penetration process, which can be mitigated with
 579 finer meshes without affecting the steady values^[51]. The prediction by the MPM (8.97) is slightly
 580 lower than those by LDFE (9.65) and arbitrary Lagrangian Eulerian (ALE; 9.47) methods with
 581 discrepancies within 7.5%. Predictions of the bearing capacity factor of 11.1 and 9.7 are also available
 582 with the coupled Eulerian-Lagrangian method^[68] and the strain path method^[65], respectively.

583 Therefore, the bearing capacity factor obtained with the MPM is reliable. The velocity magnitude
584 induced by the cone penetration at $w = 3.47D$ is shown in Figure 11.

585 In the calculations, the total runtimes for 100 incremental step were recorded for the CPU sequential
586 and multiple-GPU parallel computations (Table 2). The speedup with one-GPU parallelisation is
587 about 88, which increases to 382 with 8 GPUs invoked. To investigate the acceleration effect of the
588 GPU parallel framework, the particle number N_p was modified from 24,000,000 (6 GB) to 5,000,000,
589 10,000,000, 50,000,000 (12 GB), and 100,000,000. The maximum speedup of 89 is achieved for all
590 the cases with one GPU as the GPU is fully loaded with more than 4 million particles (Figure 12a).
591 For all the cases, the maximum speedups with different number of GPUs are fitted as $85N_{\text{gpu}}$ (Figure
592 12b), which is higher than that for the mini-slump test cases. That is due to the larger void areas, with
593 none essential computations, in the mini-slump test cases when the slurry expands on the base.

594 To summarise from the two benchmarks, very high speedups ($> 71.5N_{\text{gpu}}$) are expected with the
595 multiple-GPU parallelisation over the conventional CPU sequential calculations. However, it is also
596 noteworthy that the real-time speedup for specific dynamic problems (such as slurry slump) with very
597 large deformation of the material may be undermined by many factors, such as the percentage of void
598 elements and the bandwidth of the P2P channels in the computer. The particle migrations between
599 the GPUs bring extra complexity and may affect the robustness of the program. That means
600 optimisation of the program itself can be very time-consuming. Also, the multiple-GPU parallelised
601 program suffers from a lower portability since its software framework is dedicated to the hardware
602 platform. Therefore, the topic of multiple-GPU parallelisation of the MPM remains to be open in a
603 near future.



604
605 Figure 10 Profile of resistance to the cone penetration




Figure 11 Velocity distribution in soil at $w/D = 3.47$

Table 2 Speedups of multiple-GPU parallel simulations in 100 incremental step for cone penetration


N_p	Memory size (GB)	CPU sequential	1 GPU		2 GPUs		4 GPUs		8 GPUs	
		Runtime (s)	Runtime (s)	Speedup	Runtime (s)	Speedup	Runtime (s)	Speedup	Runtime (s)	Speedup
5,000,000	1.2	2,882	32	89	22	128	15	188	10	282
10,000,000	2.6	5,670	66	86	40	140	29	195	18	320
24,000,000	6	13,640	155	88	82	166	59	232	36	382
50,000,000	12	27,333	318	86	165	166	102	267	54	505
100,000,000	24	55,612	—	—	—	—	158	352	92	604

6. Conclusions

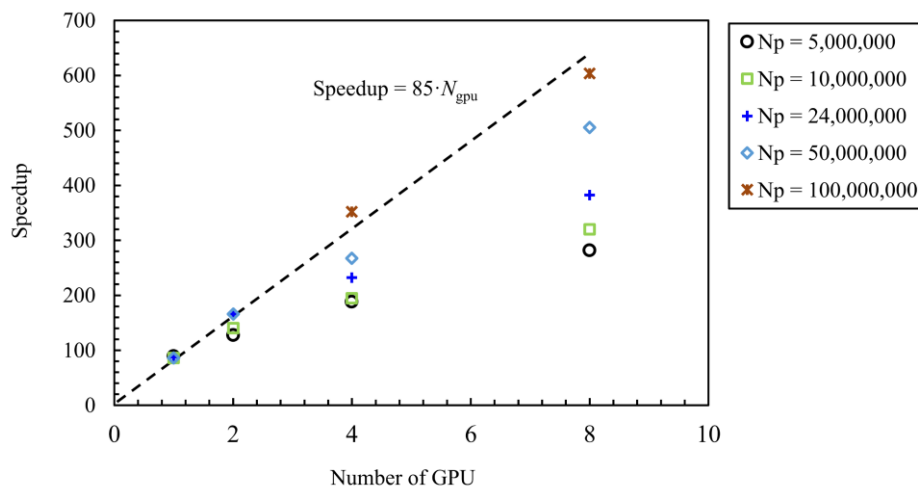
As one of the arbitrary Lagrangian-Eulerian methods, the material point method (MPM) owns intrinsic advantages in simulation of large deformation problems by combining the merits of the Lagrangian and Eulerian approaches. Significant computational intensity is involved in the calculations of the MPM due to its very fine mesh needed to achieve a high accuracy. Considering the limitations with the CPU and single-GPU performance, multiple-GPU parallelisation provides a promising means to boost the computational efficiency of the MPM. In this study, a new multiple-GPU parallel strategy was developed based on a single-root complex architecture of the computer within a CUDA environment. Peer-to-Peer (P2P) communication between the GPUs was performed to exchange the information of the crossing particles and ghost element nodes, which is faster than the heavy send/receive operations between different computers through the infiniband network. Domain decomposition is performed to split the whole computational task over the GPUs with a

622 number of subdomains. Within each GPU, a particle list was constructed for each node to avoid the
 623 data race when parallelising the ‘Interpolation from particles to nodes’.

624 The acceleration effect of the parallelisation was evaluated with two benchmarks cases, mini-slump
 625 test and cone penetration test. The maximum speedups with 1 GPU was 88, and increased to 604
 626 using 8 GPUs. Among all the functions of the MPM, the function ‘Interpolation from particles to
 627 nodes’ consumes the most computational efforts for around 70%, which is due to the non-coalesced
 628 memory access when writing the interpolations to the nodal addresses. The overhead on the
 629 establishment of the particle list takes less than 2% of the total runtime, much less than that in Dong
 630 et al. (2015) and Dong and Grabe (2018) of around 15%, as the operations are fully moved and
 631 parallelised onto the GPU. The remaining 28% of the computations is mainly on the function ‘Update
 632 particle state’, while that on the functions ‘Initialisation of nodal variables’ and ‘Calculate nodal
 633 velocities and accelerations’ are ignorable.



634
 635 (a) One GPU ($N_p = 24,000,000$)



636
 637 (b) Multiple GPUs

638 Figure 12 Speedups of GPU parallelisation with multiple-GPU

639 **Acknowledgements**

640 This paper was supported by the National Natural Science Foundations of China (Grant No. 51909248)
641 and the Open Research Fund of State Key Laboratory of Coastal and Offshore Engineering, Dalian
642 University of Technology (Grant No. LP2012).

643 This work was also supported by the NVIDIA Corporation with the donation of the GPUs Geforce
644 Titan Xp and GeForce Titan V.

645 The authors would also like to acknowledge the valuable input of Dr. Xinlei Wang at Shenzhen Zenus
646 software technology Ltd. (formerly PhD student at Zhejiang University) through personal
647 communications.

648 **Data Availability Statement**

649 The data that support the findings of this study are available from the corresponding author upon
650 reasonable request.

651 **References**

- 652 1. Di Y, Yang J, Sato T. An operator-split ALE model for large deformation analysis of geomaterials. *International*
653 *Journal for Numerical and Analytical Methods in Geomechanics*. 2007; 31: 1375-1399.
- 654 2. Nazem M, Sheng D, Carter JP, Sloan SW. Arbitrary Lagrangian–Eulerian method for large-strain consolidation
655 problems. *International Journal for Numerical and Analytical Methods in Geomechanics*. 2008; 32: 1023-1050.
- 656 3. Wang D, White DJ, Randolph MF. Large-deformation finite element analysis of pipe penetration and large-
657 amplitude lateral displacement. *Canadian Geotechnical Journal*. 2010; 47(8): 842–856.
- 658 4. Zhang Z, Pan Y, Wang J, Zhang H, Chen Z, Zheng Y, Ye H. A total-Lagrangian material point method for coupled
659 growth and massive deformation of incompressible soft materials. *International Journal for Numerical Methods*
660 *in Engineering*. 2021; DOI: <https://doi.org/10.1002/nme.6787>.
- 661 5. Sulsky D, Zhou SJ, Schreyer HL. Application of a particle-in-cell method to solid mechanics. *Comput Phys*
662 *Commun*. 1995; 87: 236–52.
- 663 6. Harlow FH. The particle-in-cell computing method for fluid dynamics. *Methods Comput Phys*. 1964; 3: 319–43.
- 664 7. Ma S, Zhang X, Lian Y, Zhou X. Simulation of high explosive explosion using adaptive material point method.
665 *Computer Modeling in Engineering & Sciences*. 2009; 39(2):101-123.
- 666 8. Nairn JA. Material point method calculations with explicit cracks. *Computer Modeling in Engineering & Sciences*.
667 2003; 4 (6): 649–664.
- 668 9. Huang P, Zhang X, Ma S, Huang X. Contact algorithms for the material point method in impact and penetration
669 simulation. *International Journal for Numerical Methods in Engineering*. 2011; 85(4): 498-517.
- 670 10. de Vaucorbeil A, Nguyen VP, Hutchinson C. A total-Lagrangian material point method for solid mechanics
671 problems involving large deformations. *Computer Methods in Applied Mechanics and Engineering*. 2019; 360:
672 112783.
- 673 11. Nguyen VP, de Vaucorbeil A, Nguyen CT, Mandal TK. A generalized particle in cell method for explicit solid

- 674 dynamics. *Computer Methods in Applied Mechanics and Engineering*. 2020; 371: 113308.
- 675 12. de Vaucorbeil A, Nguyen VP. Modelling contacts with a total Lagrangian material point method. *Computer*
676 *Methods in Applied Mechanics and Engineering*. 2020; 373: 113503.
- 677 13. York AR, Sulsky D, Schreyer HL. Fluid–membrane interaction based on the material point method. *International*
678 *Journal for Numerical Methods in Engineering*. 2000; 48(6): 901–924.
- 679 14. Stomakhin A, Schroeder C, Chai L, Teran J, Selle A. A material point method for snow simulation. *ACM Trans*
680 *Graph (TOG)*. 2013; 32(4): 102.
- 681 15. Jiang C, Gast T, Teran J. Anisotropic elastoplasticity for cloth, knit and hair frictional contact. *ACM Transactions*
682 *on Graphics*. 2017; 36(4): 152.
- 683 16. Fei Y, Guo Q, Wu R, Huang L, Gao M. Revisiting integration in the material point method: a scheme for easier
684 separation and less dissipation. *ACM Trans. Graph*. 2021; 40(4): 109.
- 685 17. Soga K, Alonso E, Yerro A, Kumar K, Bandara S. Trends in large-deformation analysis of landslide mass
686 movements with particular emphasis on the material point method. *Géotechnique*. 2015; 66(3): 248–273.
- 687 18. Dong Y, Wang D, Randolph MF. Investigating of impact forces on pipeline by submarine landslide using
688 material point method. *Ocean Engineering*. 2017; 146: 21–28.
- 689 19. Pinyol NM, Alvarado M, Alonso EE, Zabala F. Thermal effects in landslide mobility. *Géotechnique*. 2018; 68(6):
690 528-545.
- 691 20. Zhao E, Dong Y, Tang Y, Sun J. Numerical Investigation of Hydrodynamics and Local Scour around Submarine
692 Pipeline under Joint Effect of Solitary Wave and Current. *Ocean Engineering*. 2021; 222: 108553.
- 693 21. Coetzee CJ, Vermeer PA, Basson AH. The modelling of anchors using the material point method. *International*
694 *Journal for Numerical and Analytical Methods in Geomechanics*. 2005; 29(9): 879–895.
- 695 22. Ceccato F, Bisson A, Cola S. Large displacement numerical study of 3D plate anchors. *European Journal of*
696 *Environmental and Civil Engineering*. 2017; 1–19.
- 697 23. Wang L, Coombs WM, Augarde CE, Cortis M, Brown MJ, Brennan AJ, Knappett JA, Davidson C, Richards D,
698 White DJ, Blake AP. An efficient and locking-free material point method for three-dimensional analysis with
699 simplex elements. *International Journal for Numerical Methods in Engineering*. 2021; 122(15): 3876-3899.
- 700 24. Mast CM, Arduino P, Mackenzie-Helnwein P, Gregory RM. Simulating granular column collapse using the
701 Material Point Method. *Acta Geotechnica*. 2015; 10: 101–116.
- 702 25. Yuan W, Wang H, Zhang W, Dai B, Liu K, Wang Y. Particle finite element method implementation for large
703 deformation analysis using Abaqus. *Acta Geotechnica*. 2021; 12: 1-14.
- 704 26. Zhang W, Zhong Z, Peng C, Yuan W, Wu W. GPU-accelerated smoothed particle finite element method for large
705 deformation analysis in geomechanics. *Computers and Geotechnics*. 2021; 129: 103856.
- 706 27. Abe K, Kenichi S, Samila B. Material point method for coupled hydromechanical problems. *Journal of*
707 *Geotechnical and Geoenvironmental Engineering*. 2014; 140(3): 04013033.
- 708 28. Yerro A, Alonso EE, Pinyol NM. The material point method for unsaturated soils. *Géotechnique*. 2015; 65(3):
709 201-217.
- 710 29. Bandara S, Ferrari A, Laloui L. Modelling landslides in unsaturated slopes subjected to rainfall infiltration using
711 material point method. *International Journal for Numerical and Analytical Methods in Geomechanics*. 2016;
712 40(9): 1358–1380.
- 713 30. Troncone A, Conte E, Pugliese L. Analysis of the slope response to an increase in pore water pressure using the
714 material point method. *Water*. 2019; 11(7): 1446.

- 715 31. Wallstedt PC, Guilkey JE. Improved Velocity Projection for the Material Point Method. *CMES-Computer*
716 *Modeling in Engineering & Sciences*. 2007; 19(3): 223–232.
- 717 32. Wallstedt PC, Guilkey JE. An evaluation of explicit time integration schemes for use with the generalized
718 interpolation material point method. *Journal of Computational Physics*. 2008; 227: 9628-9642.
- 719 33. Nair J, Hammerquist C. Material point method simulations using an approximate full mass matrix inverse.
720 *Computer Methods in Applied Mechanics and Engineering*. 2021; 377(2): 113667.
- 721 34. Buzzi O, Pedroso DM, Giacomini A. Caveats on the implementation of the generalized material point method.
722 *Computer Model Eng Sci*. 2008; 31(2): 85–106.
- 723 35. Steffen M, Kirby RM, Berzins M. Analysis and reduction of quadrature errors in the material point method
724 (MPM). *Int J Numer Meth Eng*. 2008; 76(6): 922–48.
- 725 36. Yerro A, Alonso EE, Pinyol NM. Run-out of landslides in brittle soils. *Computers and Geotechnics*. 2016; 80:
726 427-439.
- 727 37. Charlton TJ, Coombs WM, Augarde CE. iGIMP: An implicit generalised interpolation material point method for
728 large deformations. *Computers and Structures*. 2017; 190: 108-125.
- 729 38. Gan Y, Sun Z, Chen Z, Zhang X, Liu Y. Enhancement of the material point method using B-spline basis functions.
730 *International Journal for Numerical Methods in Engineering*. 2018; 113: 411-431.
- 731 39. Sun Z, Gan Y, Huang Z, Zhou X. A local grid refinement scheme for B-spline material point method.
732 *International Journal for Numerical Methods in Engineering*. 2020; 121(11): 2398-2417.
- 733 40. Bardenhagen SG, Kober EM. The generalized interpolation material point method. *Computer Model Engineering*
734 *and Science*. 2004; 5(6): 477–96.
- 735 41. Dong Y, Wang D, Randolph MF. Quantification of impact forces on fixed mudmats from submarine landslides
736 using the material point method. *Applied Ocean Research*. 2020; 146: 21-28.
- 737 42. Zhang Y, Zhang X, Liu Y. An alternated grid updating parallel algorithm for material point method using
738 OpenMP. *Computer Modeling in Engineering & Sciences*. 2010; 69(2): 143–165.
- 739 43. Parker SG. A component-based architecture for parallel multi-physics PDE simulation. *Future Generation*
740 *Computer Systems*. 2006; 22(1): 204–216.
- 741 44. Stantchev G, Dorland W, Gumerov N. Fast parallel particle-to-grid interpolation for plasma PIC simulations on
742 the GPU. *Journal of Parallel and Distributed Computing*. 2008; 68(10):1339–1349.
- 743 45. Gibson MJ, Keedwell EC, Savić DA. An investigation of the efficient implementation of cellular automata on
744 multi-core CPU and GPU hardware. *Journal of Parallel and Distributed Computing*. 2015; 77: 11-25.
- 745 46. Dong Y, Wang D, Randolph MF. A GPU parallel computing strategy for the material point method. *Computers*
746 *and Geotechnics*. 2015; 66: 31–38.
- 747 47. Gao M, Wang X, Wu K, Pradhana A, Sifakis E, Yuksel C, Jiang C. GPU optimization of material point methods.
748 *ACM Transactions on Graphics (TOG)*. 2018; 37(6), 1-12.
- 749 48. Dong Y, Grabe J. Large scale parallelisation of the material point method with multiple GPUs. *Computers and*
750 *Geotechnics*. 2018; 101: 149-158.
- 751 49. Wang X, Qiu Y, Slattery SR, Fang Y, Li M, Zhu S, Zhu Y, Tang M, Manocha D, Jiang C. A massively parallel
752 and scalable multi-GPU material point method. *ACM Transactions on Graphics (TOG)*. 2020; 39(4): 30.
- 753 50. Ma J, Wang D, Randolph MF. A new contact algorithm in the material point method for geotechnical
754 simulations. *International Journal for Numerical and Analytical Methods in Geomechanics*. 2014; 38(11): 1197-
755 1210.

- 756 51. Dong Y. Reseeding of particles in the material point method for soil-structure interactions. *Computers and*
757 *Geotechnics*. 2020; 127: 103716.
- 758 52. Bardenhagen SG, Brackbill JU, Sulsky D. The material-point method for granular materials. *Computer Methods*
759 *in Applied Mechanics and Engineering*. 2000; 187(3-4): 529-541.
- 760 53. Bardenhagen SG, Guilkey JE, Roessig KM, Brackbill JU, Witzel WM. An improved contact algorithm for the
761 material point method and application to stress propagation in granular material. *Computer Modeling in*
762 *Engineering & Sciences*. 2001; 2(4): 509-522.
- 763 54. Zhang D, Ma X, Giguere PT. Material point method enhanced by modified gradient of shape function. *Journal*
764 *of Computational Physics*. 2011; 230(16): 6379-6398.
- 765 55. Dong Y, Wang D, Randolph M. Runout of submarine landslide simulated with material point method. *Journal of*
766 *Hydrodynamics*. 2017; 29(3): 438-444.
- 767 56. Bing Y, Cortis M, Charlton TJ, Coombs WM, Augarde CE. B-spline based boundary conditions in the material
768 point method. *Computers & Structures*. 2019; 212: 257-274.
- 769 57. Cortis M, Coombs WM, Augarde CE, Brown M, Brennan A, Robinson S. Imposition of essential boundary
770 conditions in the material point method. *International Journal for Numerical Methods in Engineering*. 2018;
771 113(1): 130-152.
- 772 58. Nishiura D, Matsuo MY, Sakaguchi H. ppohDEM: computational performance for open source code of the
773 discrete element method. *Computer Physics Communications*. 2014; 185(5): 1486-1495.
- 774 59. Nishiura D, Sakaguchi H. Parallel-vector algorithms for particle simulations on shared-memory multiprocessors.
775 *Journal of Computational Physics*. 2011; 230(5): 1923-1938.
- 776 60. Dong Y. Reseeding of particles in the material point method for soil-structure interactions. *Computers and*
777 *Geotechnics*. 2020; 127: 103716.
- 778 61. Locat J, Lee HJ. Submarine landslides: advances and challenges. *Canadian Geotechnical Journal*. 2002; 39(1):
779 193-212.
- 780 62. Thakur V, Degago SA. Quickness of sensitive clays. *Géotechnique Letters*. 2012; 2(3): 87-95.
- 781 63. Boukpeti N, White DJ, Randolph MF, Low HE. Strength of fine-grained soils at the solid-fluid transition.
782 *Géotechnique*. 2012; 62(3): 213-226.
- 783 64. Fornes P, Bihs H, Thakur VKS. Implementation of non-Newtonian rheology for Debris Flow simulation with
784 REEF3D. IAHR World Congress, 2017.
- 785 65. Teh CI, Houlsby GT. An analytical study of the cone penetration test in clay. *Géotechnique*. 1991; 41(1): 17-34.
- 786 66. Lu Q, Randolph MF, Hu Y, Bugarski IC. A numerical study of cone penetration in clay. *Géotechnique*. 2004;
787 54(4): 257-267.
- 788 67. Walker J, Yu HS. Analysis of the cone penetration test in layered clay. *Géotechnique*. 2010; 60(12): 939-948.
- 789 68. Wang D, Bienen B, Nazem M, Tian Y, Zheng J, Pucker T, Randolph MF. Large deformation finite element
790 analyses in geotechnical engineering. *Computers and Geotechnics*. 2015; 65: 104-114.
- 791 69. Ceccato F, Beuth L, Vermeer PA, Simonini P. Two-phase Material Point Method applied to the study of cone
792 penetration. *Computers and Geotechnics*. 2016; 80: 440-452.

793
794
795
796

797

798

799

Appendix A Snippets of pseudo code for multi-GPU parallelisation

800

Main Functions:

```
for (int GPU_ID = 0; GPU_ID < N_Gpus; GPU_ID++)
//N_Gpus is the total number of GPUs available
{
```

```
    cudaSetDevice(GPU_ID);
```

//Function (i): Initialisation of nodal variables

```
    int block_size = 32;
    int n_blocks = int(N_G[GPU_ID] / 32);
    //N_G is the total number of element node in each subdomain
```

```
    Initialisation_of_Nodal_Variables << <n_blocks, block_size> >> (GPU_ID,  $m_i$ ,  $v_i$ ,  $M_i$ ,  $a_i$ ,
 $F_i^{\text{int}}$ ,  $F_i^{\text{ext}}$ ,  $\omega_i^{\text{nomm}}$ ,  $\omega_i^{\text{tang}}$ );
```

801

//Function (ii): Interpolation from particles to nodes

```
    for (int Index_Layer = 0; Index_Layer < N_G_Z[GPU_ID]; Index_Layer++)
    // N_G_Z is the total number of layers of nodes along the height of the model
    {
```

```
        block_size = 32;
        n_blocks = int(N_P[GPU_ID] / 32);
        //N_P is the total number of particles in each subdomain
```

```
        Generate_Particle_List << <n_blocks, block_size> >> (GPU_ID, Index_Layer,
Particle_List,  $X_p$ );
```

```
        block_size = 32;
        n_blocks = int(GGx[GPU_ID] * GGy[GPU_ID] / 32);
        // GGx and GGy are the total number of nodes in X and Y directions, respectively
```

```
        Interpolation_From_Particles_To_Nodes << <n_blocks, block_size> >>
(GPU_ID, Index_Layer, Particle_List,  $X_p$ ,  $m_p$ ,  $V_p$ ,  $\rho$ ,  $v_p$ ,  $\sigma_p$ ,  $f_p^{\text{ext}}$ ,  $m_i$ ,  $M_i$ ,  $F_i^{\text{int}}$ ,  $F_i^{\text{ext}}$ ,  $\omega_i^{\text{nomm}}$ );
    }
```

802

//Function (iii): Calculate nodal velocities and accelerations

```
    int block_size = 32;
    int n_blocks = int(N_G[GPU_ID] / 32);
```

```
    Calculate_Nodal_Velocities_and_Accelerations << <n_blocks, block_size> >> (GPU_ID,
 $m_i$ ,  $v_i$ ,  $M_i$ ,  $a_i$ ,  $F_i^{\text{int}}$ ,  $F_i^{\text{ext}}$ ,  $\omega_i^{\text{nomm}}$ ,  $\omega_i^{\text{tang}}$ );
```

803

804

805

806

//Function (iv): Update particle state

block_size = 32;

n_blocks = int(N_P[GPU_ID] / 32);

//N_P is the total number of particles in each subdomain

Update_Particle_State << <n_blocks, block_size>> (GPU_ID, v_i , a_i , X_p , m_p , V_p , ρ , v_p , D_p ,

W_p , σ_p);

}

807

808

Individual function:

```
__global__ void Initialisation_of_Nodal_Variables (GPU_ID,  $m_i$ ,  $v_i$ ,  $M_i$ ,  $a_i$ ,  $F_i^{\text{int}}$ ,  $F_i^{\text{ext}}$ ,  $\omega_i^{\text{nom}}$ ,  
 $\omega_i^{\text{tang}}$ )  
{  
    int i = blockIdx.x * blockDim.x + threadIdx.x;  
    if(i < N_G)  
    {  
        //initialisation operations  
    }  
}  
809 }  
810
```

```

__global__ void Generate_Particle_List (GPU_ID, Index_Layer, Particle_List, Xp)
{
    int i = blockIdx.x * blockDim.x + threadIdx.x;
    if(i < N_P)
    {
        if ((Xp_X[i] >= Lower_X) && (Xp_X[i] <= Upper_X) &&
            (Xp_Y[i] >= Lower_Y) && (Xp_Y[i] <= Upper_Y) &&
            (Xp_Z[i] >= Lower_Z) && (Xp_Z[i] <= Upper_Z) )
            // Lower and Upper are the lower and upper coordinates of the influence area of
the node layer, respectively
            {
                Index_X = int((Xp_X[i] - Lower_X) / h / h_cell);
                Index_Y = int((Xp_Y[i] - Lower_Y) / h / h_cell);
                Index_Z = int((Xp_Z[i] - Lower_Z) / h / h_cell);
                // h and h_cell are the sizes of the element and fine cell, respectively

                List_Position = Index_Z*FCy*FCx + Index_Y*FCx + Index_X;
                // FCx and FCy are the total numbers of the fine cell in X and Y directions,
811 respectively

                Particle_List[List_Position] = i;
            }
        }
    }
812 }
813

```

```

__global__ void Interpolation_From_Particles_To_Nodes (GPU_ID, Index_Layer, Particle_List,
 $X_p, m_p, V_p, \rho, v_p, \sigma_p, f_p^{\text{ext}}, m_i, M_i, F_i^{\text{int}}, F_i^{\text{ext}}, \omega_i^{\text{nomm}}$ )

```

```
{
```

```
    int i = blockIdx.x * blockDim.x + threadIdx.x;
```

```
    if(i < GGx*GGy)
```

```
    {
```

```
        int NODE = Index_Layer*GGy*GGx + i;
```

```
        for(/*cycle of the elements around node i*/)

```

```
        {
```

```
            for(/*cycle of the particle list of element*/)

```

```
            {
```

```
                Execute: Eqs. (5) – (9)
```

```
            }
```

```
        }
```

```
    }
```

```
814 }
```

```
815
```

```

__global__ void Calculate_Nodal_Velocities_and_Accelerations (GPU_ID,  $m_i, v_i, M_i, a_i, F_i^{\text{int}}, F_i^{\text{ext}}, \omega_i^{\text{nomm}}, \omega_i^{\text{tang}}$ )

```

```
{
```

```
    int i = blockIdx.x * blockDim.x + threadIdx.x;
```

```
    if(i < N_G)
```

```
    {
```

```
        Execute: Eqs. (10) – (20)
```

```
    }
```

```
816 }
```

```
817
```

```
__global__ void Update_Particle_State (GPU_ID,  $v_i$ ,  $a_i$ ,  $X_p$ ,  $m_p$ ,  $V_p$ ,  $\rho$ ,  $v_p$ ,  $D_p$ ,  $W_p$ ,  $\sigma_p$ )
{
    int i = blockIdx.x * blockDim.x + threadIdx.x;
    if(i < N_P)
    {
        Execute: Eqs. (21) – (28)
    }
}
818
819
820
821
```

822

Appendix B Snippets of pseudo code for CPU sequential computations

823

//Function (i): Initialisation of nodal variables

Initialisation_of_Nodal_Variables ($m_i, v_i, M_i, a_i, F_i^{\text{int}}, F_i^{\text{ext}}, \omega_i^{\text{norm}}, \omega_i^{\text{tang}}$);

//Function (ii): Interpolation from particles to nodes

Interpolation_From_Particles_To_Nodes ($X_p, m_p, V_p, \rho, v_p, \sigma_p, f_p^{\text{ext}}, m_i, M_i, F_i^{\text{int}}, F_i^{\text{ext}}, \omega_i^{\text{norm}}$);

//Function (iii): Calculate nodal velocities and accelerations

Calculate_Nodal_Velocities_and_Accelerations ($m_i, v_i, M_i, a_i, F_i^{\text{int}}, F_i^{\text{ext}}, \omega_i^{\text{norm}}, \omega_i^{\text{tang}}$);

//Function (iv): Update particle state

Update_Particle_State ($v_i, a_i, X_p, m_p, V_p, \rho, v_p, D_p, W_p, \sigma_p$);

824

825

826

Individual function:

void Initialisation_of_Nodal_Variables ($m_i, v_i, M_i, a_i, F_i^{\text{int}}, F_i^{\text{ext}}, \omega_i^{\text{norm}}, \omega_i^{\text{tang}}$)

{

for(*/*cycle of the nodes*/*)

{

//initialisation operations

}

}

827

828


```

void Interpolation_From_Particles_To_Nodes ( $X_p, m_p, V_p, \rho, v_p, \sigma_p, f_p^{\text{ext}}, m_i, M_i, F_i^{\text{int}}, F_i^{\text{ext}}, \omega_i^{\text{norm}}$ )
{
    for(/*cycle of the particles*/)
    {
        for(/*cycle of related nodes*/)
        {
            Execute: Eqs. (5) – (9)
        }
    }
829 }
830
831

void Calculate_Nodal_Velocities_and_Accelerations ( $m_i, v_i, M_i, a_i, F_i^{\text{int}}, F_i^{\text{ext}}, \omega_i^{\text{norm}}, \omega_i^{\text{tang}}$ )
{
    for(/*cycle of the nodes*/)
    {
        Execute: Eqs. (10) – (20)
    }
832 }
833
834

void Update_Particle_State ( $v_i, a_i, X_p, m_p, V_p, \rho, v_p, D_p, W_p, \sigma_p$ )
{
    for(/*cycle of the particles*/)
    {
        Execute: Eqs. (21) – (28)
    }
835 }
836
837

```