

1 **Flexible and scalable particle-in-cell methods with adaptive**
2 **mesh refinement for geodynamic computations**

3 **Rene Gassmöller¹, Harsha Lokavarapu¹, Eric Heien², Elbride Gerry Puckett³ and**
4 **Wolfgang Bangerth⁴**

5 ¹Department of Earth and Physical Sciences, University of California, Davis

6 ²Computational Infrastructure for Geodynamics, University of California, Davis

7 ³Department of Mathematics, University of California, Davis

8 ⁴Department of Mathematics, Colorado State University, Fort Collins

9 **Key Points:**

- 10 • Particle-in-cell methods require new algorithms when used with finite element cal-
11 culations on dynamically partitioned, adaptively refined, unstructured meshes.
- 12 • We present approaches for particle generation, sorting, and hybrid load balancing in
13 hierarchically refined finite element computations.
- 14 • We show scalability and applicability of the developed methods for problems in
15 computational geodynamics.

Abstract

Particle-in-cell (PIC) methods couple mesh-based methods for the solution of continuum mechanics problems with the ability to advect and evolve properties on particles. PIC methods have a long history and numerous applications in geodynamic modeling. However, they are historically either implemented in sequential codes, or in parallel codes with structured, statically partitioned meshes. Yet, today's codes increasingly use adaptive mesh refinement (AMR) of unstructured coarse meshes, dynamic repartitioning, and scale to thousands of processors. Optimally balancing the work per processor for a PIC method in these environments is a difficult problem, and many existing implementations are not sufficient for this task. Thus, there is a need to revisit these algorithms for future applications.

Here we describe challenges and solutions to implement PIC methods in the context of large-scale parallel geodynamic modeling codes that use dynamically changing meshes. We also provide guidance for how to address bottlenecks that impede the efficient implementation of these algorithms and demonstrate with numerical tests that our algorithms can be implemented with optimal complexity and that they are suitable for large-scale, practical applications. We provide a reference implementation in ASPECT (Advanced Solver for Problems in Earth's convection), an open source code for geodynamic modeling built on the DEAL.II finite element library.

1 Introduction

Most methodologies to numerically solve flow problems are based on a continuum description in the form of partial differential equations, and include the finite element, finite volume, and finite difference methods. On the other hand, it is often desirable to couple these methods with discrete, "particle" approaches for a number of applications. These include, for example, visualization of flows, tracking interfaces and origins, or tracking the history of material. Use cases and discussions of computational methods can be found as far back as *Harlow* [1962] and are often referred to as *particle-in-cell* (PIC) methods.

Different implementations of such methods can be found in the geodynamic literature [*Poliakov and Podladchikov*, 1992; *Moresi et al.*, 2003; *Gerya and Yuen*, 2003; *McNamara and Zhong*, 2004; *Popov and Sobolev*, 2008; *Thielmann et al.*, 2014], but almost all of these methods were developed for either structured meshes and/or sequential computations. However, over the past two decades adaptive finite element methods have demon-

47 strated that they are vastly more accurate than computations on uniformly refined meshes
48 [*Carey, 1997; Ainsworth and Oden, 2000; Bangerth and Rannacher, 2003; Kronbichler*
49 *et al., 2012; Heister et al., 2017*], and have been successfully combined with PIC meth-
50 ods in other fields [*Wallstedt and Guilkey, 2010; Adams et al., 2015; Balay et al., 2018;*
51 *Almgren et al., 2013*]. While many parts of existing particle-in-cell algorithms can still
52 be used in this context, a number of new algorithmic challenges arise. The present con-
53 tribution is therefore primarily an assessment of possible algorithms when implementing
54 particle methods for computational geodynamics in the following two situations:

- 55 1. Unstructured, hierarchically refined quad-/octree, 2D/3D meshes that change dy-
56 namically and potentially utilize higher order polynomial mappings to represent
57 curved geometries;
- 58 2. Large parallel computations that run on thousands of cores, using tens of millions
59 of cells, and billions of particles.

60 Specifically, we will discuss the following components, along with an assessment of
61 their practical performance:

- 62 1. Parallel generation of particles in unstructured meshes;
- 63 2. Treatment of particles as they cross cell and processor boundaries;
- 64 3. Treatment of particles during mesh refinement and coarsening, including appropri-
65 ate load balancing.

66 Other components in our reference implementation use well understood algorithms:
67 We use standard C++-containers as data structures; integrate the particle trajectories us-
68 ing Forward-Euler, Runge-Kutta 2 or 4 integration schemes with higher order accuracy
69 in space and time; store variable scalar, vector, or tensor-valued properties on particles;
70 and transfer information between particles and mesh using simple arithmetic or harmonic
71 cell averaging schemes, or least-squares projections. Massively parallel output capability is
72 provided by the VTK [*Schroeder et al., 2006*] and HDF5 [*Folk et al., 1999*] data formats.
73 As our manuscript is focussed on the particular difficulties of combining particle methods
74 with adaptive finite element computations, we do not discuss traditional difficulties of par-
75 ticle methods, such as memory locality or particle clustering, as these have already been
76 addressed elsewhere [*Mellor-Crummey et al., 2001; Wang et al., 2015*].

77 We provide a reference implementation of the presented methods in the geodynamic
78 modeling code ASPECT [*Kronbichler et al., 2012; Bangerth et al., 2017a; Heister et al.,*
79 *2017*], and include most of the discipline-independent methods in the deal.II finite ele-
80 ment library [*Bangerth et al., 2007; Arndt et al., 2017*], thus making them available for a
81 variety of applications and scientific disciplines. Given that our implementation is based
82 on deal.II, we will henceforth only consider quadrilateral and hexahedral, hierarchically
83 refined meshes, which are balanced by a 2:1 refinement ratio between neighboring cells.
84 Exploiting these assumptions allows us to optimize our algorithms, but we believe that
85 generalizations to other situations are often straightforward.

86 **2 Computational methods**

87 **2.1 Parallel particle generation**

88 The first step in using particles in mesh-based solvers is their creation on all in-
89 volved processors, and depending on their purpose, initial particle distributions may vary
90 widely. Two broad classes of initial distributions come to mind:

91 *Random particle positions.* Randomly chosen particle locations are often used in
92 cases where particles represent the values of a field; e.g., the origin and movement of a
93 specific type of material. In these cases, one is not interested in prescribing exact initial
94 particle locations, and randomly chosen locations are acceptable. The probability distri-
95 bution, $\rho(\mathbf{x})$, from which locations are drawn is often chosen as uniform over the domain.
96 Alternatively, one can use a higher particle density in regions of interest, for example to
97 better resolve steep gradients, which can be interpreted as equivalent to AMR in mesh-
98 based methods.

99 We propose the following algorithm, running on each processor:

- 100 1. Compute and store local cell weights as integral of $\rho(\mathbf{x})$ over each local cell.
- 101 2. Compute the global sum of the local cell weight integrals.
- 102 3. Compute the local number of particles as ratio between local and global weight
103 integral times global number of particles.
- 104 4. Compute the local starting particle index based on the partial sum of local number
105 of particles of all processes with lower rank.

- 106 5. Either: Compute the number of particles per cell by randomly drawing cells K ac-
 107 cording to their weight repeatedly and tallying up how many times each cell was
 108 selected.
- 109 6. Or: Compute the number of particles per cell according to their share of the local
 110 integral of $\rho(\mathbf{x})$
- 111 7. Generate local particles in each cell K by drawing random locations inside its axes-
 112 parallel bounding box B_K until we find a position in K (see Supplementary Text S1
 113 for details).

114 Apart from the two global reductions to determine the global weight and the local
 115 start index, all of the operations above are local to each processor. Thus, the overall run
 116 time for generating particles is proportional to the number of particles on the process with
 117 the largest number of particles, i.e., of optimal complexity in the global number of par-
 118 ticles and, if the number of particles per process is balanced, also in the number of pro-
 119 cesses. However, this balancing is often not the case in practice (see Section 2.3).

120 We note that our algorithm yields a number of particles on each process that is de-
 121 terministic. Consequently, the distribution of particles is not entirely random. However, in
 122 practice we find this does not matter for sufficiently many particles.

123 *Prescribed particle locations* An alternative to the random arrangements of parti-
 124 cles is to exactly prescribe initial locations, either algorithmically (e.g., a regular grid), or
 125 by reading locations from a file. Surprisingly, for distributed unstructured meshes this case
 126 is more computationally expensive than randomly generated particle locations.

127 Let us assume that the initial positions of all particles are given in an array $\{\mathbf{x}_k\}$,
 128 $k = 1 \dots N$. Then for each particle one has to find its surrounding cell, which in the worst
 129 case, is of complexity global number of particles times local number of cells. This is be-
 130 cause, for general unstructured meshes, we can not predict whether a given particle's lo-
 131 cation lies inside the locally owned cells without searching through all cells. This limits
 132 the usefulness of the algorithm to moderate numbers of particles. However, the algorithm
 133 can be accelerated by checking whether a particle's location lies inside the bounding box
 134 of the locally owned cells, before checking each cell.

135 On hierarchically refined meshes, one can alternatively find the cell K by finding
 136 the coarse level cell in which it is located, and then recursively searching through its chil-


137 dren. This reduces the complexity to the global number of particles times the logarithm
138 of the local number of cells. However, it only works if child cells occupy the same vol-
139 ume as their parent cell; this condition is often not met when using nonlinear polynomial
140 mappings to represent curved geometries.

141 In the paragraphs above we assume that the particle positions are known in the
142 global coordinate system, and we have to search for the surrounding cell. If however, the
143 particle coordinates are known in the *local* cell coordinate system (e.g. the center), then
144 the algorithm is much simpler. A loop over all cells and all local particle coordinates that
145 are then mapped into the real space [as used in *Puckett et al., 2017*] will generate the par-
146 ticles in the optimal order, and will be cheap.

147 **2.2 Transport between cells and subdomains**

148 PIC codes in geodynamics contain a time integration in which one computes a ve-
149 locity field (usually on some grid), and then moves the particles with the flow field. To
150 parallelize these computations the grid is usually fully distributed, which means each pro-
151 cess only knows about local cells and one layer of “ghost” cells around the local domain.
152 Thus, after each particle movement the new particle location is either inside its original
153 cell K or in a different cell K' . To be able to transfer data between particle and grid we
154 then need to find its new cell that may be owned by the same processor or a different
155 one. The challenge, in the context of adaptive, distributed meshes lies in constructing al-
156 gorithms that can efficiently search for the new surrounding cells of particles, as well as
157 potentially transfer the particle to a different processor. In practice, communication pat-
158 terns that cover the exchange of particles between processes that own adjacent parts of the
159 mesh are often sufficient to implement, i.e., using point-to-point messages. In particular,
160 this is possible if the (ODE) time step is chosen such that the CFL number is less than or
161 equal to one, because then particles travel no more than one cell diameter in each step.
162 Following these arguments our reference implementation employs the following algorithm,
163 executed for each particle that is not in its old cell:

- 164 1. Search for the locally known current cell K' .
- 165 2. If K' is owned by the current process, mark the particle as being in K' .
- 166 3. If K' is in a ghost cell owned by the process p , mark the particle for transmission
167 to p .



186 **Figure 1.** In 2:1 balanced quadtree meshes finding the new cell K' for a particle that has left its old cell K
 187 is a nontrivial problem. Limiting the search to the cells that contain the vertex of the old cell that is closest
 188 to the new particle position (left panel) reduces the search cost. Note that sorting neighbor cells according to
 189 angle between a and $b_{K'}$ (right panel, see main text for definitions) correctly predicts the search order (red
 190 numbers) in most cases, while a simpler criterion like particle–cell–center distance mispredicts the new cell
 191 (center panel).

168 4. If K' cannot be found, mark particle for deletion.

169 After the algorithm has finished, (1) all particles marked for transmission are com-
 170 municated to their neighbors that now own them, (2) all particles that have been lost or
 171 communicated are removed from local storage, and (3) all particles with a new cell as-
 172 sociation (local or communicated) are reinserted into local storage. This bulk handling is
 173 advantageous, since particles of the same cell tend to move into the same neighbor cells,
 174 and a collective insertion reduces copies and reallocation of memory.

175 The vast majority of particles remain in the current cell, end up in a new local cell
 176 (option 2), or a cell owned by another process (option 3). A few cases, however, do not
 177 fall in these categories. First, the ODE integrator error during particle movement can
 178 carry a particle over a processor boundary, and out of the one-cell ghost layer. Second,
 179 the integrator error can transport a particle across a geometry boundary, after which it
 180 is not contained in any cell. For a benchmark model setup (see Supplementary Dataset
 181 S1) we have found that only a negligible fraction of particles is lost because of these two
 182 mechanisms. As expected, an explicit Euler scheme loses significantly more particles than
 183 the RK2 methods, while decreasing the time step significantly reduces the loss. Given the
 184 small overall loss and added computational expense to reduce the timestep, dropping parti-
 185 cles that fall out of bounds (option 4) seems like a reasonable approach to us.

192 The algorithm above requires finding the cell a particle is in now (Step 1). As dis-
 193 cussed in Section 2.1, without additional information this requires $O(N_{\text{local cells}})$ operations,
 194 all of which are expensive. Furthermore, because many of the local particles cross to a
 195 different cell, this step is not of optimal – i.e., $O(N)$ for N particles – complexity. While
 196 the tree structure of the mesh makes it possible to implement global tree-search algorithms
 197 with logarithmic complexity in the number of cells [Isaac *et al.*, 2015], we found that the
 198 algorithm spends the majority of its works on determining whether a particle is inside a
 199 cell K' , i.e., inverting the mapping of K' for the position of the particle. Since in our ap-
 200 plication the vast majority of particles only cross from one cell to its neighbors, we can
 201 accelerate the global algorithms significantly by first searching all neighbor cells in an or-
 202 der that makes it likely that we find the correct one early. Only the very small fraction
 203 that does not end up in a neighbor then requires an expensive search over all cells. We
 204 note that for problems without this local property other algorithms might be more appro-
 205 priate [e.g. Mirzadeh *et al.*, 2016; Burstedde, 2018].

206 Following some experimentation, we found that the following strategy to pre-sort local
 207 neighbor cells works best (see also Fig. 1): Let \mathbf{p}' be the particle’s current position,
 208 K the known previous cell of the particle, \mathbf{v} be the vertex of K closest to \mathbf{p}' , and $\mathbf{c}_{K'}$ be
 209 the center of the potential new cell K' , which is a vertex neighbor of K adjacent to vertex
 210 \mathbf{v} (see Fig. 1). Let $\mathbf{a} = \mathbf{p}' - \mathbf{v}$ be the normalized vector from the closest vertex of K to
 211 the particle, and $\mathbf{b}_{K'} = \mathbf{c}_{K'} - \mathbf{v}$ be the normalized vector from the closest vertex to the
 212 center of cell K' . Then we search through all K' in the order of descending scalar product
 213 $\mathbf{a} \cdot \mathbf{b}_{K'}$ (Fig. 1, right panel). In other words, cells with a center in the direction of the par-
 214 ticle movement are checked first. This algorithm is somewhat similar to the one proposed
 215 by [Capodaglio and Aulisa, 2017], with the difference that we know our particle is in one
 216 of the neighbors of the old cell, and we therefore search through a sorted list of neighbor
 217 cells, instead of along a computed search path through multiple cells. While there are cor-
 218 ner cases in which our algorithm fails to find the new cell in the first try, in practice more
 219 than 98% of the particles moving to a new cell in the models discussed in Section 3 are
 220 found immediately. The rest of the particles is found in at most 2 (2D) or 4 (3D) searches,
 221 except if the particle left the immediate neighbors of the old cell as discussed above. Sim-
 222 pler criteria – like searching by distance between particle and cell center (Fig. 1, middle
 223 panel) – fail more often, in particular for adaptively refined neighbors.

224 If particles crossed a process boundary they are communicated to neighboring pro-
225 cesses in two steps. First, two integers are exchanged between every neighbor and the
226 current process, representing the number of particles that will be sent and received. In a
227 second step every process transmits the serialized particle data and receives its respective
228 data from its neighbors. This allows us to implement all communications as non-blocking
229 point-to-point MPI transfers, only generating $O(1)$ transmissions and $O(N_{\text{local particles}})$ data
230 per process. Since we already determined which ghost cell contains this particle on the
231 old process, we also transmit this information. Because ghost cells are guaranteed to ex-
232 ist on the owning process we thus avoid another search for the enclosing cell on the new
233 process.

234 **2.3 Handling adaptively refined, dynamically changing meshes**

235 In the current context, adaptively refined, dynamically changing meshes present two
236 particular challenges.

237 *Mesh refinement and repartitioning* Typically, refinement and coarsening happens
238 in two steps: First, cells are refined or coarsened separately on each process, and particles
239 are distributed to the children of their previous cell (upon refinement), or are merged to
240 the parent of their previous cell (upon coarsening). The second step of mesh adaptation
241 consists of redistributing the resulting mesh to achieve an efficient parallel load distribu-
242 tion [Burstedde et al., 2011; Bangerth et al., 2011]. To keep this process simple we append
243 the serialized particle data to other data already attached to a cell (such as vertex locations
244 and values of field based solution variables), and transmit all data at the same time. We
245 can therefore utilize existing software for parallel mesh handling [Burstedde et al., 2011],
246 which uses well-optimized bulk communication patterns, and thereby avoid sending parti-
247 cles individually or having to re-join particles with their cells.

248 *Load balancing* The mesh repartitioning discussed in the previous paragraph is de-
249 signed to redistribute work equally among all available processes. For mesh-based meth-
250 ods, this typically means equilibrating the number of cells each process “owns”. On the
251 other hand, in the context of PIC methods for adaptive meshes, the number of particles
252 per cell frequently ranges from zero to a few hundred. Consequently, the described pro-
253 cess leads to unbalanced workloads during particle-related parts of the code. Conversely,
254 rebalancing the mesh to equilibrate the number of particles leaves the mesh-based algo-

255 rithms with unbalanced workloads. Both situations reduce the overall parallel efficiency of
256 the code.

257 The only approach to restore perfect scalability is to partition cells differently for the
258 mesh-based and particle-based parts of the code. On the other hand, one can not avoid
259 transporting all mesh and particle data during these rebalancing steps, because each phase
260 of the algorithm might require all data from the other. Consequently, the amount of data
261 that has to be transported twice per time step is significant.

262 In practice, some level of imbalance can often be tolerated. One can work with the
263 following compromise solutions:

- 264 1. *Repartition the mesh according to the combined particle and cell load* (“*Balanced*
265 *repartitioning*”). Instead of estimating the workload of each cell during the rebal-
266 ancing step as constant (pure mesh-based methods) or proportional to the number
267 of particles in a cell (pure particle-based methods), one can estimate it as an appro-
268 priately weighted sum of the two. The resulting mesh is optimal for neither of the
269 two phases, but is better balanced than either of the extremes (see Section 3.2 and
270 Supporting Figure S1).
- 271 2. *Ignore imbalance*. As long as the number of particles is small one may simply ig-
272 nore the imbalance. A typical case is when particles are only used to output in-
273 formation for a few specific points of interest, e.g. an accumulated strain profile
274 through a subducting slab.
- 275 3. *Adjust particle density to mesh during particle generation*. The particle density can
276 be chosen to follow the mesh resolution, if the region of highest mesh resolution
277 is known in advance. This is most useful for tracking pre-existing interfaces. The
278 higher particle density close to the interface then not only increases the accuracy in
279 regions of interest, but it also improves parallel efficiency and scalability.
- 280 4. *Adjust mesh to particle density*. Instead of prescribing the particle density following
281 the mesh, the mesh resolution can also be adjusted to the particle distribution. As
282 in the previous alternative, the alignment of mesh and particle density yields better
283 parallel efficiency and scaling.
- 284 5. *Adjust particle density to mesh by particle population management*. In cases of a pri-
285 ori unknown regions of high mesh density it can be necessary to manage the par-
286 ticle density actively during the model run. This includes removing particles from

287 regions with high particle density or adding particles in regions of low density. If
288 done appropriately, the result will be a mesh where the average number of particles
289 per cell is managed so that it remains approximately constant.

290 While the last three approaches lead to better scalability, they may of course not suit
291 the problem one originally wanted to solve. On the other hand, generating additional parti-
292 cles upon refinement of a cell, and thinning out particles upon coarsening, is a common
293 strategy in existing codes [*Popov and Sobolev, 2008; Leng and Zhong, 2011*]. We also
294 note that while load balancing is particularly important for dynamically changing adap-
295 tive meshes, it is also beneficial for uniform meshes if the particle distribution happens to
296 be non-uniform.

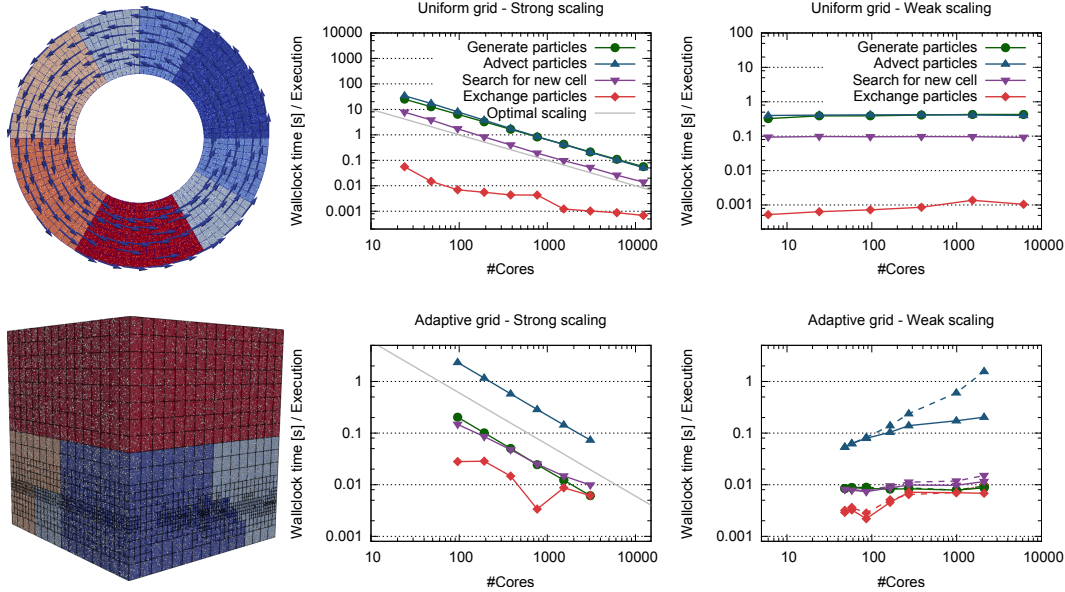
297 **3 Scalability**

298 To verify our claims of performance and scalability, we show that our algorithms
299 scale well to typical model sizes in computational geodynamics. Technical information
300 about the used hardware, and the definition of the timing events is provided in Supple-
301 mentary Text S2. Additional benchmarks confirming the correctness of the implemented
302 advection schemes is provided in Supporting Text S3, and Figure S1.

303 **3.1 Uniform meshes**

311 We first show scalability using a two-dimensional benchmark case with a static and
312 uniformly refined mesh. We employ a circular-flow setup in a spherical shell, with no flow
313 across the boundary. Particles are distributed randomly with uniform density (see Fig. 2,
314 top left), and are advected using a RK2 integration scheme.

315 The top row of Fig. 2 shows excellent weak and strong scaling over at least three or-
316 ders of magnitude of model size. For a fixed problem size (strong scaling), we use a mesh
317 with $786,432 = 12 \cdot 256^2$ cells and $1.536 \cdot 10^7$ particles. Increasing the number of processes
318 from 12 to 12,288 shows an almost perfect decrease in wall time for all operations, despite
319 the rather small problem each process has to deal with for large numbers of processes.
320 Note that the scaling of the “Exchange particles” event is likely specific to the used net-
321 work topology and probably shows the transition from a large-throughput large-latency
322 mode of transfer to a small message-size small-latency transfer.



304 **Figure 2.** Scaling of algorithms. Top row: Results for a uniformly refined mesh. Bottom row: Results
 305 for an adaptively refined mesh. Left column: Model geometry and initial parallel partition. Center column:
 306 Strong scaling for a constant number of cells and particles. Top right: Weak scaling for a uniform mesh with
 307 a constant number of cells and particles per process. Bottom right: Weak scaling for an adaptive mesh with
 308 a fixed (though increasingly unbalanced) number of cells and particles per process. The dashed models use
 309 the common cell load balancing, while the solid models use balanced repartitioning as described in Subsec-
 310 tion 2.3.

323 Keeping the number of cells and particles per core fixed and increasing the problem
 324 size and number of processes accordingly (weak scaling, Fig. 2, top right panel), the wall-
 325 clock time stays constant between 6 and 6,144 processes. In this test each process owns
 326 512 cells and $1.0 \cdot 10^4$ particles. Each refinement step leads to four times as many cells,
 327 and consequently processes. 6,144 cores was the last multiple to which we had access for
 328 timing purposes. Results again show excellent scalability, even to large problem sizes, in
 329 this case approximately 3 million cells and 61 million particles.

3.2 Adaptively refined meshes

331 Discussing scalability for *adaptive* meshes is more complicated because increasing
 332 level of refinement does not create a predictable number of cells. We apply the same
 333 particle distribution and integration as for the uniform mesh case, but use a model setup
 334 based on the benchmarks presented in [van Keken et al., 1997], extended to three spatial

335 dimensions. Specifically, we use a rectangular domain $[0, 0.9142] \times [0, 1] \times [0, 1]$ that con-
336 tains a sharp non-horizontal interface separating a less dense lower layer from a denser
337 upper layer. The shape of the interface then leads to a Rayleigh-Taylor instability. For the
338 strong scaling tests, we create an adaptive mesh of at most 256^3 cells, retaining fine cells
339 only in the vicinity of the interface. This mesh consists of approximately 1,000,000 cells,
340 and we generate approximately 30 million, uniformly distributed particles, and run this
341 setup on increasing numbers of processors.

342 The results in Fig. 2 show that *strong* scaling for the adaptive grid case is nearly as
343 good as for the uniform grid case, decreasing the total runtime essentially linearly from 96
344 to 3,072 cores. The small worse-than-linear component of the cell-search algorithm seems
345 to be related to the imbalance between particles and cells that will be further discussed
346 in the weak scaling results, but since this part is one order of magnitude cheaper than the
347 particle advection it will only limit the scalability beyond 10,000 cores. As for the uni-
348 form mesh the “Exchange particles” algorithm shows some variations, likely caused by the
349 interaction between the allocated compute nodes and the network topology used for the
350 tests. Because this scaling test actually solves for the Stokes solution on the finite element
351 mesh we are more restricted in the number of possible model sizes compared to the syn-
352 thetic test for uniform meshes above. Increased memory consumption excludes very small
353 core numbers and limited scaling of the Stokes solver for very small number of degrees of
354 freedom per core limits the maximum number of cores. Nevertheless, 100 to 3000 cores
355 is the most common model size for our application and increasing or decreasing the model
356 size has not revealed significant changes to the scaling behavior outside of the here pre-
357 sented range.

358 Setting up weak scaling tests requires further consideration. Since we can not pre-
359 dict the number of cells for a given number of mesh refinements, we use a 16^3 mesh and
360 adaptively refine it a variable number of times taking note of the resulting numbers of
361 cells. We then run this model series with increasing number of cores to keep the number
362 of cells per process approximately constant at 550 cells per process. Each of the models
363 uses ≈ 25 times as many particles as cells, uniformly distributed across the domain.

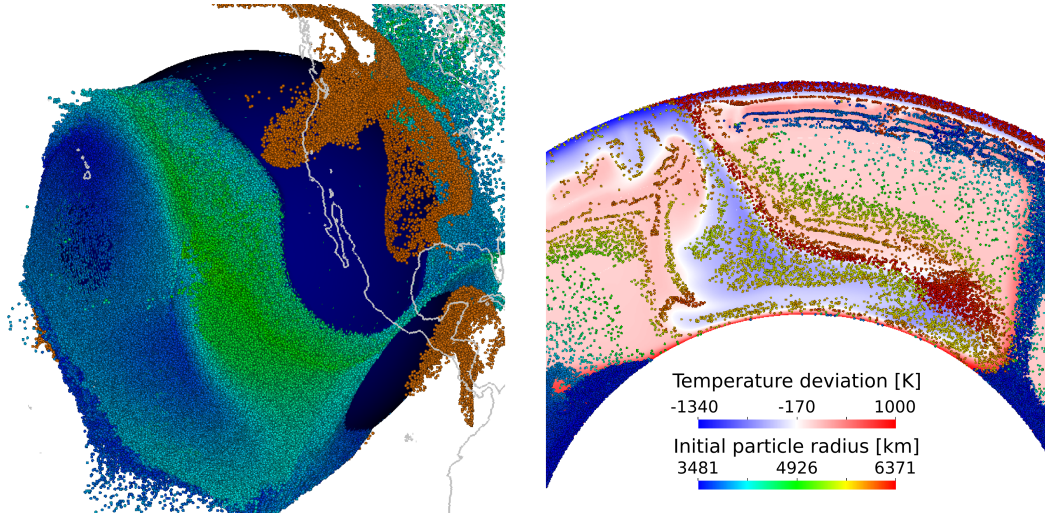
364 The weak scaling results are more difficult to interpret than the strong scaling case.
365 In a first series, we only strive to balance the number of cells per process (Option 2 in
366 Subsection 2.3). However, because the particle density is constant while cell sizes increas-

367 ingly vary, the imbalance in the number of particles per process grows with the size of the
368 model. This is easily seen in the bottom left panel of Fig. 2 in which all four processes
369 own the same number of cells, but vastly different volumes and consequently numbers of
370 particles. Therefore, the run time for some parts of the algorithm – in particular for par-
371 ticle advection – grows as the model size increases (dashed lines, bottom right panel of
372 Fig. 2).

373 As discussed in Section 2.3, this effect can be addressed by balancing cell and parti-
374 cle numbers. The solid lines in the bottom right panel of Fig. 2 show that with appro-
375 priately chosen weights, the increase in runtime can be reduced from a factor of 30 to a
376 factor of 4. To achieve this, we introduce a cost factor W for each particle. The total cost
377 of each cell in load balancing is then one (the cost of the field-based methods per cell)
378 plus W times the number of particles in this cell. $W = 0$ implies that we only consider the
379 number of cells for load balancing, whereas $W = \infty$ only considers the number of parti-
380 cles. In practice, one will typically choose $0 \leq W < 1$; for realistic applications, we found
381 $W = 0.01$ to be adequate. On the other hand, computational experiments suggest that it is
382 not important to *exactly* determine the optimal value since the overall runtime varies only
383 weakly in the vicinity of the minimum (see Supporting Figure S1).

384 **4 Example application: Convection in the Earth’s mantle**

385 We illustrate the applicability of our algorithms to realistic applications by model-
386 ing compressible Stokes flow in the Earth’s mantle constrained by known movements of
387 the tectonic plates at the surface for the past 250 million years. The equations we solve
388 and the model setup are identical to a previously published model [Heister *et al.*, 2017],
389 but enhanced by adding 4.8 million particles, which are used to track material movement
390 over time. The particles are generated randomly with a uniform distribution, are integrated
391 with a RK2 integration scheme, and in order to enforce balanced parallel workloads we
392 limit the maximum number of particles per cell to 25 and remove additional particles dy-
393 namically during the model run. Therefore, at the final time regions with coarse cells have
394 a lower particle density than finely resolved regions (see right panel of Fig. 3). As the
395 number of particles is relatively small, it was not necessary to use balanced repartition-
396 ing to improve load balancing. Material properties such as density and heat capacity are
397 computed from a database for basaltic and harzburgitic rocks, following [Nakagawa *et al.*,
398 2009], and the viscosity is based on a published viscosity model incorporating mineral



402 **Figure 3.** Illustration of a 3D mantle convection model with particles. Left: Subducting plates below the
 403 Western United States (brown particles) push material at the core-mantle boundary (dark blue sphere) towards
 404 the west. Only a selection of particles is shown, and each is colored by the distance from its initial position
 405 (blue: small to green: large). Right: Vertical slice through the subduction zone. All particles close to the
 406 slice are shown, and they are colored by the radius of their initial position (red: surface; blue: core-mantle
 407 boundary).

399 physics properties, geoid deformation, and seismic tomography [*Steinberger and Calder-*
 400 *wood, 2006*]. The prescribed surface velocities use reconstructions of past plate movement
 401 on Earth [*Seton et al., 2012*].

408 In the first time steps of this example model (before the number of particles is in-
 409 fluenced by particle deletion) particle advection takes approximately 2.0 s per time step,
 410 particle cell-search requires 1.4 s per time step, particle generation is a one time process
 411 requiring 6.7 s, and particle communication was negligible, compared to a total time per
 412 time step of 26 s. A linear extrapolation to a larger number of particles (e.g. 20 per cell,
 413 as needed for active particles) would suggest a total particle cost of about 50 % of the
 414 total runtime, although this is highly simplified as for more particles a balanced reparti-
 415 tioning strategy could save significant amounts of runtime.

416 Fig. 3 shows a part of the example model, the present-day state of the Farallon sub-
 417 duction zone below the Western United States. Particles that are initially close to the core-
 418 mantle boundary are colored by the displacement they have experienced. This reveals that
 419 the Farallon slab (orange) has primarily pushed the easternmost material. Particles in the

420 Central Pacific have not moved significantly, illuminating the limited influence of the West
421 Pacific subduction zones.

422 **5 Conclusions**

423 In this article, we have presented strategies for implementing PIC methods in com-
424 putational geodynamic problems that use unstructured adaptive meshes. We have de-
425 scribed our algorithms for the parallel generation of particles including both random and
426 prescribed particle locations, and how utilizing information about the neighbors of cells
427 can efficiently help to predict the owning cell of a particle. We discussed different load
428 balancing techniques during mesh repartitioning and explained how balanced repartition-
429 ing can improve scalability significantly even in the presence of imbalanced workloads
430 such as the ones that occur when combining unstructured AMR and PIC methods. Finally,
431 we have documented in scaling tests and application examples that the expected optimal
432 complexities can indeed be realized in practice. While there is certainly room for opti-
433 mization in the presented algorithms, we are convinced that the present state allows for
434 useful combination of unstructured AMR and PIC techniques in geodynamic modeling
435 codes. Our implementation is freely available as part of the ASPECT and DEAL.II soft-
436 ware.

437 **Acknowledgments**

438 All models were computed with the open-source software ASPECT [*Bangerth et al., 2017b,*
439 <http://aspect.geodynamics.org>] published under the GPL2 license, and the
440 necessary data to reproduce the models is included in the supplementary material. We
441 thank the Computational Infrastructure for Geodynamics (<http://geodynamics.org>) which
442 is funded by the National Science Foundation under awards EAR-0949446 and EAR-
443 1550901.

444 R. Gassmüller and W. Bangerth were partially supported by the National Science
445 Foundation under award OCI-1148116 as part of the Software Infrastructure for Sustained
446 Innovation (SI2) program; and by the Computational Infrastructure in Geodynamics ini-
447 tiative (CIG), through the National Science Foundation under Award No. EAR-0949446
448 and The University of California – Davis. E. G. Puckett was supported by the National
449 Science Foundation under Award ACI-1440811 as part of the SI2 Scientific Software Ele-
450 ments (SSE) program.

451 The computational resources were provided by the North-German Supercomputing
 452 Alliance (HLRN) as part of the project bbk00003 “Plume-Plate interaction in 3D mantle
 453 flow – Revealing the role of internal plume dynamics on global hot spot volcanism”.

454 **References**

- 455 Adams, M., P. O. Schwartz, H. Johansen, P. Colella, T. J. Ligocki, D. Martin, N. Keen,
 456 D. Graves, D. Modiano, B. Van Straalen, et al. (2015), Chombo software package for
 457 amr applications-design document, *Tech. rep.*
- 458 Ainsworth, M., and J. T. Oden (2000), *A Posteriori Error Estimation in Finite Element*
 459 *Analysis*, John Wiley and Sons.
- 460 Almgren, A. S., J. B. Bell, M. J. Lijewski, Z. Lukić, and E. Van Andel (2013), Nyx: A
 461 massively parallel amr code for computational cosmology, *The Astrophysical Journal*,
 462 765(1), 39.
- 463 Arndt, D., W. Bangerth, D. Davydov, T. Heister, L. Heltai, M. Kronbichler, M. Maier, J.-P.
 464 Pelteret, B. Turcksin, and D. Wells (2017), The deal.II library, version 8.5, *Journal*
 465 *of Numerical Mathematics*, doi:10.1515/jnma-2016-1045.
- 466 Balay, S., S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin,
 467 V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes,
 468 R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, and
 469 H. Zhang (2018), PETSc users manual, *Tech. Rep. ANL-95/11 - Revision 3.9*, Argonne
 470 National Laboratory.
- 471 Bangerth, W., and R. Rannacher (2003), *Adaptive Finite Element Methods for Differential*
 472 *Equations*, Birkhäuser Verlag.
- 473 Bangerth, W., R. Hartmann, and G. Kanschat (2007), deal.II – a general purpose object
 474 oriented finite element library, *ACM Trans. Math. Softw.*, 33(4), 24.
- 475 Bangerth, W., C. Burstedde, T. Heister, and M. Kronbichler (2011), Algorithms and data
 476 structures for massively parallel generic adaptive finite element codes, *ACM Trans.*
 477 *Math. Softw.*, 38(2).
- 478 Bangerth, W., J. Dannberg, R. Gassmüller, T. Heister, et al. (2017a), ASPECT: Advanced
 479 Solver for Problems in Earth’s ConvecTion, User Manual, doi:10.6084/m9.figshare.
 480 4865333, doi:10.6084/m9.figshare.4865333.
- 481 Bangerth, W., J. Dannberg, R. Gassmüller, T. Heister, , et al. (2017b), Aspect v1.5.0 [soft-
 482 ware], doi:http://doi.org/10.5281/zenodo.344623.

- 483 Burstedde, C. (2018), Parallel tree algorithms for AMR and non-standard data access,
484 *ArXiv e-prints*.
- 485 Burstedde, C., L. C. Wilcox, and O. Ghattas (2011), p4est: Scalable algorithms for par-
486 allel adaptive mesh refinement on forests of octrees, *SIAM J. Sci. Comput.*, *33*(3), 1103–
487 1133, doi:10.1137/100791634.
- 488 Capodaglio, G., and E. Aulisa (2017), A particle tracking algorithm for parallel finite ele-
489 ment applications, *Computers & Fluids*, *159*, 338–355.
- 490 Carey, G. F. (1997), *Computational Grids: Generation, Adaptation and Solution Strategies*,
491 Taylor & Francis.
- 492 Folk, M., A. Cheng, and K. Yates (1999), HDF5: A file format and I/O library for high
493 performance computing applications, in *Proc. ACM/IEEE Conf. Supercomputing (SC'99)*.
- 494 Gerya, T. V., and D. A. Yuen (2003), Characteristics-based marker-in-cell method with
495 conservative finite-differences schemes for modeling geological flows with strongly vari-
496 able transport properties, *Physics of the Earth and Planetary Interiors*, *140*(4), 293–318.
- 497 Harlow, F. (1962), *The particle-in-cell method for numerical solution fo problems in fluid*
498 *dynamics*.
- 499 Heister, T., J. Dannberg, R. Gassmüller, and W. Bangerth (2017), High accuracy mantle
500 convection simulation through modern numerical methods – II: realistic models and
501 problems, *Geophys. J. Int.*, *210*(2), 833–851, doi:https://doi.org/10.1093/gji/ggx195.
- 502 Isaac, T., C. Burstedde, L. C. Wilcox, and O. Ghattas (2015), Recursive Algorithms for
503 Distributed Forests of Octrees, *SIAM Journal on Scientific Computing*, *37*(5), C497–
504 C531, doi:10.1137/140970963.
- 505 Kronbichler, M., T. Heister, and W. Bangerth (2012), High accuracy mantle convection
506 simulation through modern numerical methods, *Geophysics Journal International*, *191*,
507 12–29.
- 508 Leng, W., and S. Zhong (2011), Implementation and application of adaptive mesh refine-
509 ment for thermochemical mantle convection studies, *Geochemistry, Geophysics, Geosys-*
510 *tems*, *12*(4), doi:10.1029/2010GC003425, q04006.
- 511 McNamara, A. K., and S. Zhong (2004), Thermochemical structures within a spherical
512 mantle: Superplumes or piles?, *Journal of Geophysical Research*, *109*(B7), 1–14, doi:
513 10.1029/2003JB002847.
- 514 Mellor-Crummey, J., D. Whalley, and K. Kennedy (2001), Improving memory hierarchy
515 performance for irregular applications using data and computation reorderings, *Intern-*

- 516 *tional Journal of Parallel Programming*, 29(3), 217–247.
- 517 Mirzadeh, M., A. Guittet, C. Burstedde, and F. Gibou (2016), Parallel level-set methods
518 on adaptive tree-based grids, *Journal of Computational Physics*, 322, 345–364, doi:10.
519 1016/J.JCP.2016.06.017.
- 520 Moresi, L., F. Dufour, and H. B. Muhlhaus (2003), A Lagrangian integration point finite
521 element method for large deformation modeling of viscoelastic geomaterials, *J. Comp.*
522 *Ph.*, 184, 476–497.
- 523 Nakagawa, T., P. J. Tackley, F. Deschamps, and J. A. Connolly (2009), Incorporating self-
524 consistently calculated mineral physics into thermochemical mantle convection simula-
525 tions in a 3-D spherical shell and its influence on seismic anomalies in Earth’s mantle,
526 *Geochemistry, Geophysics, Geosystems*, 10(3).
- 527 Poliakov, A., and Y. Podladchikov (1992), Diapirism and topography, *Geophysical Journal*
528 *International*, 109(3), 553–564.
- 529 Popov, A. A., and S. V. Sobolev (2008), SLIM3D : A tool for three-dimensional thermo-
530 mechanical modeling of lithospheric deformation with elasto-visco-plastic rheology,
531 *Physics of the Earth and Planetary Interiors*, 171, 55–75, doi:10.1016/j.pepi.2008.03.007.
- 532 Puckett, E. G., D. L. Turcotte, Y. He, H. Lokavarapu, J. M. Robey, and L. H. Kellogg
533 (2017), New numerical approaches for modeling thermochemical convection in a com-
534 positionally stratified fluid, *Physics of the Earth and Planetary Interiors*, doi:https:
535 //doi.org/10.1016/j.pepi.2017.10.004.
- 536 Schroeder, W., K. Martin, and B. Lorensen (2006), *The Visualization Toolkit: An Object-*
537 *Oriented Approach to 3D Graphics*, 3rd ed., Kitware, Inc.
- 538 Seton, M., R. Müller, S. Zahirovic, C. Gaina, T. Torsvik, G. Shephard, a. Talsma, M. Gur-
539 nis, M. Turner, S. Maus, and M. Chandler (2012), Global continental and ocean
540 basin reconstructions since 200Ma, *Earth-Science Reviews*, 113(3-4), 212–270, doi:
541 10.1016/j.earscirev.2012.03.002.
- 542 Steinberger, B., and A. R. Calderwood (2006), Models of large-scale viscous flow in the
543 Earth’s mantle with constraints from mineral physics and surface observations, *Geo-*
544 *physical Journal International*, 2, 1461–1481, doi:10.1111/j.1365-246X.2006.03131.x.
- 545 Thielmann, M., D. A. May, and B. J. P. Kaus (2014), Discretization errors in the hybrid
546 finite element particle-in-cell method, *Pure and Applied Geophysics*, 171, 2165–2184.
- 547 van Keken, P. E., S. D. King, H. Schmeling, U. R. Christensen, D. Neumeister, and M.-P.
548 Doin (1997), A comparison of methods for the modeling of thermochemical convection,

549 *Journal of Geophysical Research: Solid Earth*, 102(B10), 22,477–22,495, doi:10.1029/
550 97JB01353.

551 Wallstedt, P., and J. Guilkey (2010), A weighted least squares particle-in-cell method for
552 solid mechanics, *International Journal for Numerical Methods in Engineering*, 85(13),
553 1687–1704.

554 Wang, H., R. Agrusta, and J. Hunen (2015), Advantages of a conservative velocity inter-
555 polation (cvi) scheme for particle-in-cell methods with application in geodynamic mod-
556 eling, *Geochemistry, Geophysics, Geosystems*, 16(6).

Figure 1.



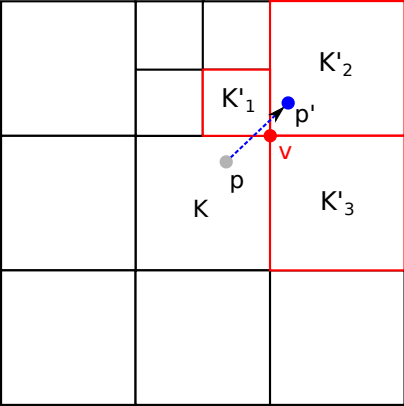


Figure 2.


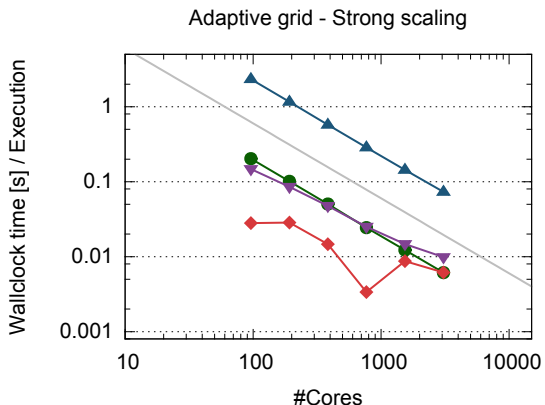

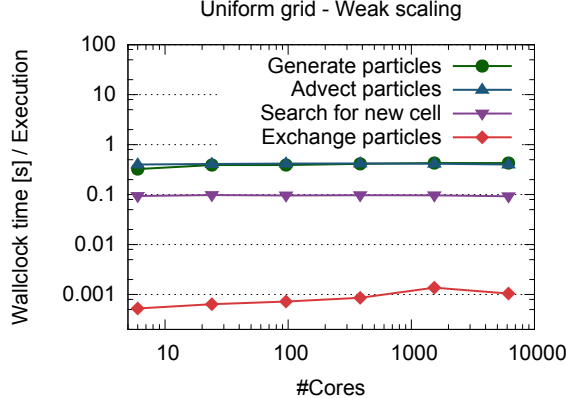
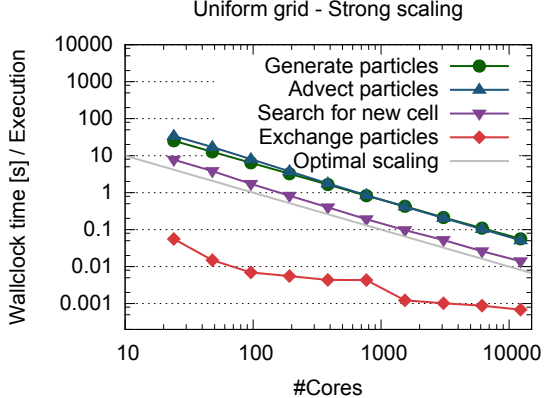
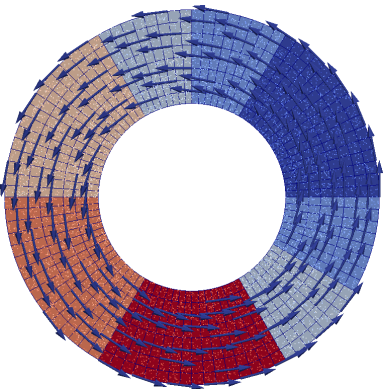


Figure 3.

