

This is a repository copy of *Performance of second order particle-in-cell methods on modern many-core architectures*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/136576/>

Version: Accepted Version

Conference or Workshop Item:

Brown, Dominic, Bettencourt, Matthew T., Wright, Steven A. orcid.org/0000-0001-7133-8533 et al. (2 more authors) (Accepted: 2017) Performance of second order particle-in-cell methods on modern many-core architectures. In: IOP Computational Plasma Physics Conference, 20-22 Nov 2017. (In Press)

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Performance of Second Order Particle-in-Cell Methods on Modern Many-Core Architectures

Dominic A. S. Brown¹, Matthew T. Bettencourt², Steven A. Wright¹, John P. Jones³,
and Stephen A. Jarvis¹

¹Department of Computer Science, University of Warwick, UK

²Electromagnetic Theory, Sandia National Laboratories, Albuquerque, NM

³UK Atomic Weapons Establishment, Aldermaston, UK

Abstract

The emergence of modern many-core architectures that offer an extreme level of parallelism makes methods that were previously infeasible due to computational expense now achievable. Particle-in-Cell (PIC) codes often fail to fully leverage this increased performance potential due to their high use of memory bandwidth. The use of higher order PIC methods may offer a solution to this by improving simulation accuracy significantly for an increase in computational intensity when compared to their first order counterparts. This greater expense is accompanied with only a minor increase in the amount of memory throughput required during the simulation.

In this presentation we will show the performance of a second order PIC algorithm. Our implementation uses second order finite elements and particles that are represented with a collection of surrounding ghost particles. These ghost particles each have associated weights and offsets around the true particle position and therefore represent a charge distribution. We test our PIC implementation against a first order algorithm on various modern compute architectures including Intel's Knights Landing (KNL) and NVIDIA's Tesla P100.

Our preliminary results show the viability of second order methods for PIC applications on these architectures when compared to previous generations of many-core hardware. Specifically, we see an order of magnitude improvement in performance for second order methods between the Pascal and Kepler GPU architectures, despite only a $4\times$ improvement in theoretical peak performance between the architectures. Although these initial results show a large increase in runtime over first order methods, we hope to be able to show improved scaling behaviour and increased simulation accuracy in the future.




Figure 1: Execution time of a two stream problem using 1 million particles, 500 steps