



Particle-in-cell Plasma Simulation Using Supercomputers Enhances Computational Physics

By Linda Barney | November 14, 2016





simulations of ultra-intense laser pulses interaction with various targets currently play a key role in discovering ways of using intense lasers for important applications ranging from creating compact ion sources for hadron therapy to fundamental experiments in quantum electrodynamics, as well as nuclear and quark-nuclear physics.

A collaborative team created a Particle-in-Cell plasma simulation tool called PICADOR that is being used in this research. PICADOR is a tool for Particle-in-Cell plasma simulation on heterogeneous cluster systems. The core of the code is developed at Lobachevsky State University of Nizhni Novgorod (Russia) by an interdisciplinary team of researchers in physics, computer science and high performance computing (<http://hpc-education.unn.ru/en/research/overview/laser-plasma>). PICADOR is used and extended in collaboration with the Institute of Applied Physics of the Russian Academy of Sciences and Chalmers University of Technology (Sweden).

“There is a growing need for carrying out large-scale 3D particle-in-cell simulations in several research directions of plasma physics,” says Dr. Iosif Meyerov, vice-head of mathematical software and supercomputing technologies department at Lobachevsky State University of Nizhni Novgorod, Russia. “Performing such simulations is possible on supercomputers with specialized parallel codes. The particle-in-cell method inherently allows massively parallel processing and thus can be efficiently implemented on supercomputers. The PICADOR code is an important tool in a multitude of research projects that can benefit from using the novel architecture of Intel systems.”

“The growth of computational power accompanied with multilevel parallelization and optimization leads to gradual extension of capabilities of particle-in-cell codes, giving access to fascinating studies that have been previously impossible,” says Dr. Arkady Gonoskov, researcher, Lobachevsky State University of Nizhni Novgorod, Institute of Applied Physics of RAS, Chalmers University of Technology, who coordinates the



Using intense lasers for accelerating ions during interaction with special targets is one of the most interesting applications of intense lasers that could have a high impact in industry, medicine and science. Although various approaches have been proposed, finding a way to accelerate ions to a sufficiently high energy remains one of the greatest challenge in this area for more than 15 years.

Recently, simulations with PICADOR code helped theoreticians from Chalmers University of Technology (Sweden) Felix Mackenroth, Arkady Gonoskov and Mattias Marklund discover and propose a new basic approach that they called Chirped-Standing-Wave Acceleration [6]. The distinguishing feature of the approach is a complete control of the acceleration process achieved via frequency variation within a laser pulse, so-called chirp. Simulations indicate the possibility of using currently available lasers for creating a source of protons with energy of 100 MeV, which is sufficient for many topical applications.


PICADOR enhances studies of highly nonlinear vacuum properties triggered by laser fields

Another interesting research activity performed based on simulations with PICADOR code is identifying how light of extraordinary intensity can produce avalanches of electron-positron pairs, creating a plasma, something which will be observable with the next generation of high intensity lasers facilities, such as XCELS [<http://www.xcels.iapras.ru/>].

According to Dr. Arkady Gonoskov, "The physics involved, as well as being computationally demanding, is extremely complex, which required developing a variety of sophisticated algorithms [A. Gonoskov et al. Phys. Rev. E 92, 023305 (2015)]. Simulations show that the collective dynamics of particles under these extreme conditions gives rise to a rich variety of unexpected phenomena, and yields new



counter-intuitively and be attracted to these regions of extreme intensity, because of so-called anomalous radiative trapping [5]. As a result, an extremely dense bunch of particles forms and interact with the strong electromagnetic field of the laser, giving rise to a unique source of energetic particles and photons.”





anomalous radiative trapping in a laser-formed dipole wave of total peak power 20 TW. Courtesy of Dr. Arkady Gonoskov, Lobachevsky State University of Nizhni Novgorod, Institute of Applied Physics RAS, Chalmers University of Technology.

What is Particle-in-Cell plasma simulation?

The Particle-in-Cell method operates on two major sets of data: an ensemble of charged particles (electrons and ions of various types) and grid values of an electromagnetic field and current density. The method has a good potential for parallelization. A key aspect of high-performance implementation of the Particle-in-Cell method is obtaining an efficient memory access pattern during the most computationally intensive Particle-Grid operations: field interpolation and current deposition.

Ensemble of particles

Grid values of field

No Particle-Particle operations


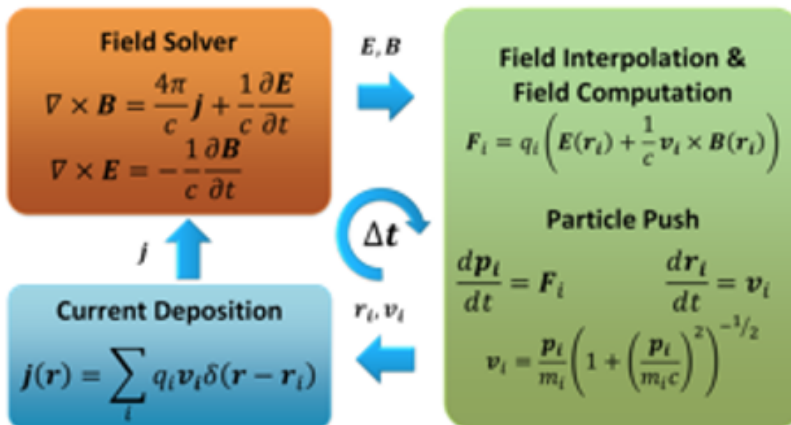


Figure 2: Particle-in-Cell plasma simulation research method. Courtesy of PICADOR team, Lobachevsky University, Institute of Applied Physics RAS, and Chalmers University of Technology.

Providing 3D simulation



supports dynamic load balancing. Each Message Passing Interface (MPI) process handles a part of a simulation area (domain) using a multicore CPU or an Intel Xeon Phi processor via Intel MPI. All MPI exchanges occur only between processes handling neighboring domains. A CPU + Intel Xeon Phi processor heterogeneous mode is supported.

The team developed the PICADOR particle simulation code over a period of five years with increased tuning occurring over the past two years. The team originally designed PICADOR code to use multi-process parallelism driven by Intel MPI to use cluster execution. Later work was done on the code to make it both parallel and vector processing friendly. The PICADOR code is currently under development and is undergoing constant improvements.

The team made substantial efforts on PICADOR code so that it would run efficiently on Intel or other hardware. Specific work was done to optimize code to take advantages of the hardware features of Intel Xeon processors and both Intel Xeon Phi processors and coprocessors to speed up processing [3, 4]. A specific goal was to have better memory utilization of cache and memory subsystems. Because memory can't typically catch up to the speed due to the technology limitations, the team worked to improve memory locality, enhancing scaling efficiency and vectorization, including auto-vectorization by an Intel compiler via `#pragma simd` [4]. A specialized optimization technique used "supercells" that aggregate several particle cells together so that super cell information is processed together—thus improving processing speed [4].

Simulation modeling with supercomputers running Intel Xeon Phi processors

A vast majority of plasma simulations using PICADOR are run on Intel processors, including Intel Xeon Phi processors because the platform provides the potential for running parallelizable codes. In their research, the team used the Intel Xeon Phi coprocessor 7120 (formerly known as Knight's Corner) compared to Intel Xeon Phi



performance of plasma simulation compared to other Intel processors and coprocessors. It is beneficial for both main scenarios of using PICADOR.

- For small/moderate size problems that require parameter scanning, involving lots of independent runs with different sets of parameters, an increase in processing speed directly transforms into speedup of research time.
- For large-scale simulations, often run using all available computational resources of a cluster system, researchers can reduce the computational noise by using more particles without an increase in CPU-hours. For cluster systems available to the users of PICADOR, there is an inherent limit on CPU-hours to use for a single run—72 hours is the limit. If research is being done on a system using Intel's Intel Xeon Phi processor 7250, a researcher typically would not want to solve the largest-for-CPU problem faster, but to instead increase accuracy by choosing a larger grid or more particles, so that it is the largest-solvable problem. Thus, for extreme-sized problems, using the Intel Xeon Phi processor helps improve accuracy while staying within the CPU-hour limit.

A significant benefit of optimizing PICADOR code to run on Intel Xeon Phi processors is that optimization is beneficial for Intel Xeon Phi coprocessors and other Intel processors as well. This essentially allows maintaining a single version of the code for all Intel hardware (except minor intrinsic-based pieces).

PICADOR software code optimization techniques

In developing PICADOR code, the team used Intel MPI and OpenMP* software packages for parallelism. Intel® Vectorization Advisor and Intel VTune Amplifier tools were used to analyze and optimize code. Recommendations for optimizing the code to run on Intel Xeon and Intel Xeon Phi processors include:

First, do appropriate optimization (memory, vectorization) on the CPU. Try to



reciprocal instruction features. The team found that a “gather/scatter-based implementation” becomes beneficial with AVX-512.

Speed-up in plasma simulation using Intel Xeon Phi processors

PICADOR testing used a frozen plasma benchmark problem with a 40×40×40 grid, 50 particles per cell and 1000 time steps that can be solved on a single CPU or an Intel Xeon Phi processor. The research uses a run time of the baseline version running a single MPI process per device with one OpenMP thread per core on CPU and four OpenMP threads per core on the Intel Xeon Phi processor. Using a rebuilding of PICADOR code, testing resulted in a 2.43 times speedup on the Intel Xeon Phi processor 7250 as compared to running the same code on an Intel Xeon Phi coprocessor 7120 [4].

After further optimization, the speedup on the Intel Xeon Phi processor 7250 was 3.47 times compared to running on an Intel Xeon Phi coprocessor 7120 and was 2.35 times faster compared to running on an Intel Xeon processor E5-2697 v3 (formerly known as Haswell) [4].

The team also used supercells but found that using larger supercells results in decreasing the number of independent sub-problems solved in parallel using OpenMP. The team found that for their code, empirically the best supercell size on the Intel Xeon Phi processor 7250 was two, on the Intel Xeon Phi coprocessor 7120 and the Intel Xeon processor E5-2697 v3 it was two for particle push and four for current deposition. Supercells are beneficial for all three platforms in question, with a 1.21 x speedup on the Intel Xeon processor E5-2697 v3 and 1.32 x speedup on the Intel Xeon Phi processor 7250 and the Intel Xeon Phi coprocessor 7120 [4].

“Parallel processing benefits from hardware instructions that are only available in Intel Xeon Phi processors. These special instructions make it



Zakhar Matveev, Intel Product Architect, Consultant on HPC FOR

project

How HPC will extend plasma simulation in the future

A growing demand for large-scale simulation in many areas of computational physics imposes high requirements for both supercomputing hardware and software for scientific computing. Such software must be highly scalable to match modern many-core architectures. However, with the growing number of cores and Single Instruction, Multiple Data (SIMD) width, it becomes difficult to efficiently utilize even a single device, let alone a whole supercomputer.

“Thus, we believe it is very important for modern and future many-core architectures to be efficiently programmable via standard programming languages and parallel programming libraries and technologies. Otherwise, the growing gap between HPC experts and researchers (physicists, biologists, etc.) would make it nearly impossible for most scientific computing codes to utilize computational resources of supercomputers and immense amounts of computational time and energy would be effectively wasted,” says Meyerov.

“One challenge is making progress with vectorization and parallelization. From Intel Advisor reports, it was clear that efficiency must be improved using the ‘gather and scatter’ technique because even though you can do computations in parallel, there is still a serialized bottleneck in memory. Using supercells is one way to improve this issue but it requires more work. Changes must be made in algorithms to make use of vector processing units available in modern highly parallel platforms like Intel Xeon Phi processors to make them more vector friendly. While changes are relatively easy, it still requires code modifications. Other future enhancements for HPC systems must also include thread and vector-level parallelism optimization and effective memory utilization,” says Zakhar Matveev, Intel Product Architect.



Surmin. Particle-in-cell plasma simulation on heterogeneous cluster systems. Journal of Computational Science. 3 (6), 2012. P. 474-479.

[2] A. Gonoskov, S. Bastrakov, E. Efimenko, A. Ilderton, M. Marklund, I. Meyerov, A. Muraviev, A. Sergeev, I. Surmin, E. Wallin. Extended Particle-in-Cell Schemes for Physics in Ultrastrong Laser Fields: Review and Developments. Physical Review E. 92 (2), 2015. P. 023305.

[3] I.A. Surmin, S.I. Bastrakov, E.S. Efimenko, A.A. Gonoskov, A.V. Korzhimanov, I.B. Meyerov. Particle-in-Cell laser-plasma simulation on Xeon Phi coprocessors. Computer Physics Communications. 202, 2016. P. 204–210.

[4] I. Surmin, S. Bastrakov, Z. Matveev, E. Efimenko, A. Gonoskov, I. Meyerov. Co-design of a particle-in-cell plasma simulation code for Intel Xeon Phi: a first look at Knights Landing. <https://arxiv.org/abs/1608.01009>

[5] A. Gonoskov, A. Bashinov, I. Gonoskov, C. Harvey, A. Ilderton, A. Kim, M. Marklund, G. Mourou, and A. Sergeev. Anomalous Radiative Trapping in Laser Fields of Extreme Intensity, Physical Review Letters, 113, 014801 – Published 2 July 2014. <http://journals.aps.org/prl/abstract/10.1103/PhysRevLett.113.014801>


[6] F. Mackenroth, A. Gonoskov, and M. Marklund. Chirped-Standing-Wave Acceleration of Ions with Intense Lasers, Physical Review Letters, 117, 104801 – Published 30 August 2016.

<http://journals.aps.org/prl/abstract/10.1103/PhysRevLett.117.104801>


Linda Barney is the founder and owner of Barney and Associates, a technical/marketing writing, training and web design firm in Beaverton, OR.




Sandia to receive Fujitsu 'green' processor



MIT joins White House supercomputing effort to speed up search for COVID-19 solutions



Early research on existing drug compounds via supercomputing could combat coronavirus



Creating the heart of a quantum computer: developing qubits



NEED R&D WORLD NEWS IN A MINUTE?

We Deliver!

R&D World Enewsletters get you caught up on all the mission critical news you need in research and development. Sign up today.






R&D world

@RandDWorld

Silicon Carbide — “Silicon Carbide, excellent for high temperature applications”

Silicon Carbide is among the hardest of ceramics, and retains hardness and strength at elevated temperatures which translates into among the best wear resistance also: bit.ly/3jxZpm1



[Embed](#)

[View on Twitter](#)

R&D WORLD DIGITAL ISSUES



[VIEW BACK ISSUES](#)

[SUBSCRIBE](#)



SUBSCRIBE TO R&D WORLD MAGAZINE

ENEWS SIGN UP

CONTACT US

DRUG DISCOVERY & DEVELOPMENT

PHARMACEUTICAL PROCESSING

2020 GLOBAL FUNDING FORECAST

Copyright © 2020 WTWH Media LLC. All Rights Reserved. The material on this site may not be reproduced, distributed, transmitted, cached or otherwise used, except with the prior written permission of WTWH Media

[Privacy Policy](#) | [Advertising](#) | [About Us](#)