

# PARALLELIZATION OF RADIA MAGNETOSTATICS CODE\*

Anushka Banerjee<sup>†</sup>, Stony Brook University, Stony Brook, NY 11794, USA

Oleg Chubar, Brookhaven National Laboratory, Upton, NY 11973, USA

Gaël Le Bec, Joel Chavanne, ESRF, 38000 Grenoble, France

Boaz Nash, Christopher Hall, Jonathan Edelen, RadiaSoft LLC, Boulder, CO 80301, USA

## Abstract

Radia 3D magnetostatics code has been used for the design of insertion devices for light sources over more than two decades. The code uses the magnetization integral approach that is efficient for solving permanent magnet and hybrid magnet structures. The initial version of the Radia code was sequential, its core written in C++ and interface in the Mathematica language. This paper describes a new Python interfaced parallel version of Radia and its applications. The parallelization of the code was implemented on C++ level, where the semi-analytical calculations of interaction matrix elements and resultant magnetic fields were parallelized using the Message Passing Interface. The parallel performance results are encouraging, particularly for magnetic field calculation post relaxation where a ~600 speedup with respect to sequential execution was obtained. The new parallel Radia version facilitates designs of insertion devices and lattice magnets for novel particle accelerators.

## INTRODUCTION

The Radia code has been continuously improved for nearly two decades since its inception and development in the Insertion Devices laboratory of the European Synchrotron Radiation Facility (ESRF). From establishing boundary integral method-based calculations as an alternative and better approach than the conventional finite element, for e.g., FLUX3D [1], Radia has been widely used to compute fields for undulators. Over the last few years, Radia has not only provided better time performance and precision results when compared to FEM codes [2], it provides excellent agreement between the calculated and measured field values corresponding to relatively simple structures like quadrupole and sextupole [3], or complicated structures like the iron-dominated electromagnetic structure: SOLEIL Undulator HU256 [4].

Following the concepts of object-oriented programming, the core part of Radia is written in C++ and was initially interfaced to run serially on Wolfram Mathematica and Igor Pro only. Currently, the new, open-source version of Radia is already available in Python interface [5], with ongoing developments to make it executable on web-based graphical interface like Jupyter Notebook [6]. Radia's Python interface has been used for magnetic "cross talk" computations of the recent ESRF-Extremely Brilliant Source upgrade, where a large number of sequential calculations with different input parameters were performed in parallel at the ESRF cluster.

\* Work supported by the US DOE BES SBIR grant no. DE-SC0018556.

<sup>†</sup> anushka.banerjee@stonybrook.edu

The calculated values had good agreement with the measured values (relative errors in the  $10^{-4}$  range) [7].

## METHODS USED FOR PARALLELIZATION

With an increased development in the field of light sources, the importance of magnetostatics code has increased by manifolds compared to what it was a few decades earlier. Despite Radia's stellar performance when compared to conventional FEM based code, the target computations involving 3D magnetic simulations of the insertion devices and accelerator magnets are quite complicated and CPU-intensive, with further increasing complexities. With an aim to further improve Radia's overall performance, parallelization tasks were undertaken.

MPI is a communication protocol and is used as the industry standard for the message passing model where a certain application comprises of a set of tasks which are assigned their own local memory whose location can be in the same machine or across several machines. Data exchange to conduct the operation by tasks is established by sending and receiving messages [8]. MPI provides programmers the flexibility to use it as a low-level approach with a detailed control on the flow of data, or as a high-level programming approach with parallel libraries designed to provide optimized performance without going into the depths of the MPI algorithm [9].

Solving any 3D magnetostatics problems in Radia comprises of 3 subsequent steps: calculation of elements of a large (often tens GBs memory size) matrix called Interaction Matrix and describing magnetic interaction between 'active' sub-volumes of a magnet geometry, created by segmentation (1), performing a relaxation procedure on it to determine values of the magnetization vector in all the "active" sub-volumes (2), and computing magnetic field and/or field integrals or other characteristics of magnetic fields, created by the sub-volumes with magnetization or current density (3). The two sections— the generation of the interaction matrix (1) and the calculation of magnetic field values after relaxation (3) are 'embarrassingly parallel' algorithms and hence have been parallelized using MPI at the C++ level.

Interaction matrix is a dense matrix in Radia, that may occupy a large memory, depending on geometry and its segmentation. Relaxation of this matrix may require a large number of iterations, in particular when solving complicated iron-dominated geometries [2,4]. At each of these iterations, a relatively small number of multiplications needs to be done to take into account (or update) magnetization vector in each

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2022). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

sub-volume (created by segmentation). This makes application of MPI not very efficient for parallelizing this algorithm. Instead shared memory-based parallelization, implemented using native multithreading capabilities introduced in C++11 standard library has been considered for parallelizing the specific section of the Radia code. The structure used for the tests is discussed in detail in the following section titled ‘Test Magnetic Structure’, followed by the outcomes detailed in ‘Results of Benchmarking’.

### TEST MAGNETIC STRUCTURE

To evaluate the parallel performance of the new version of the Radia 3D magnetostatics code, we used a 3D model of a 14mm-period, hybrid, in-vacuum, Cryo-cooled Permanent Magnet Undulator (CPMU) at 4mm magnetic gap (see Fig. 1). This model comprises of an upper and a lower array with each of these arrays including an alternating sequence of permanent magnets (PrFeB with 1.67 T remnant magnetization) and soft iron poles (Vanadium Permendur with ~2.29 T magnetization at saturation). The model geometry is generated by mirroring with respect to the three orthogonal symmetry planes passing through the magnetic centre. The geometry (specifically the poles) has been extensively subdivided to achieve solving accuracy better than 0.001 with respect to the peak magnetic field. 22680 independent sub-volumes are produced because of implementing subdivision before applying the symmetries, consequently yielding a dense interaction matrix of size ~2GB. The final calculated on-axis vertical field after solving the geometry for the magnetization in the sub-volumes is shown in Fig. 2.




Figure 1: Radia 3D model of the hybrid, in-vacuum, CPMU used for our parallel performance tests.

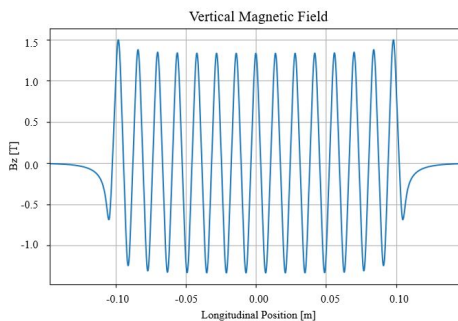


Figure 2: Calculated on-axis vertical magnetic field along the CPMU structure using Radia.

### RESULTS OF BENCHMARKING

Tests to evaluate parallel performance on many server nodes were conducted at the National Energy Research Scientific Computing Centre (NERSC), specifically the NERSC supercomputer Cori. Cori is a Cray XC40 with a theoretical peak performance of 30 petaflops. Execution times for each section were obtained at the end of each test, iteratively performed starting from 1 process (i.e., serial execution) and going up to 1024 processes and then averaged over 3 runs for each process number, thereby estimating the performance characteristics.

The parallel version of the Radia code performed well for both the cases – computation of elements of the interaction matrix and computation of magnetic field post relaxation of this matrix. For both the instances, a systematic decline in the required execution times was observed with increasing number of processes, see Fig. 3, that can be interpreted as the corresponding computational speedup achieved because of parallelizing the Radia code, leading to a maximum speedup value of ~200 for the setting up of interaction matrices, and approximately ~600 for computation of field values post relaxation, see Fig. 4.

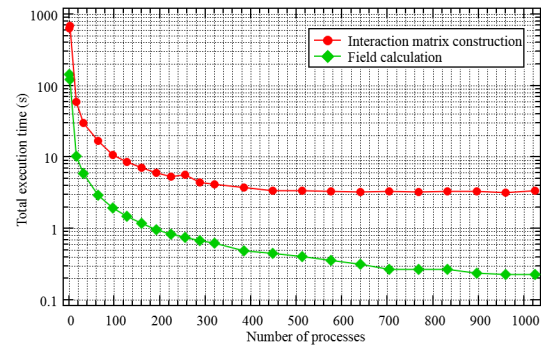


Figure 3: Mean time to set up the interaction matrix and calculate magnetic field, both as a function of the number of MPI processes.

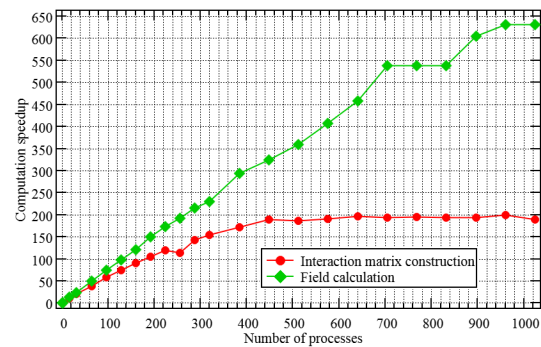


Figure 4: Total speedup in time values to set up interaction matrix and calculate field vs. number of MPI processes.

The performance corresponding to computation of elements of the interaction matrix demonstrates significant

speed gain only until number of parallel processes is lesser than 400, whereas no such parallel performance saturation is observed for the function of field calculation after relaxation. Thus, profiling of the parallel performance was implemented at the C++ level, where time stamps were introduced at the beginning and end of specific functions in the code, to evaluate their corresponding execution times. Functions specifically related to communication and exchange of data between processes were profiled, namely 'MPI\_Send' and 'MPI\_Recv'. Primarily, 'MPI\_Send' is executed by workers after they have computed the data packets to be sent to the master, and 'MPI\_Recv' is performed by the master to receive the data packets sent by each worker. The plots demonstrating the total time required for the execution of 'MPI\_Send' by each worker and 'MPI\_Recv' by the master are given in Fig. 5 for setting up of interaction matrix and Fig. 6 for calculation of field post relaxation.




Figure 5: Total time to execute 'MPI\_Send' by each worker, 'MPI\_Recv' by master, and complete execution of interaction matrix generation vs. number of MPI processes.

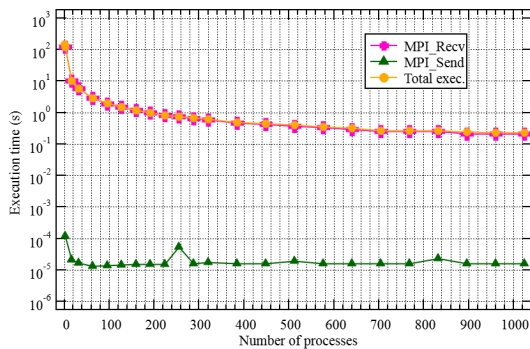


Figure 6: Total time to execute 'MPI\_Send' by each worker, 'MPI\_Recv' by master, and complete execution of field calculation vs. number of MPI processes.

For interaction matrix construction (Fig. 5), it can be observed that at some point near the saturation, the time spent by a worker in 'MPI\_Send' exceeds the total time spent by the master in 'MPI\_Recv', such a crossover is absent in case of the field calculation post relaxation, where the time spent by a worker in 'MPI\_Send' is always much smaller than the time spent by the master in 'MPI\_Recv', as observed

in Fig. 6. The role of the master is to receive data after the calculation by workers. Thus, initially (for lesser number of workers), for interaction matrix generation, master time is mostly spent waiting for workers to finish the computation. But with increasing number of workers, the computation time decreases (i.e., each worker executes lesser number of calculations) and eventually master is unable to receive quickly enough all the computed data, thus delaying the workers who have already finished computing and instead get queued to send the data. This saturation is absent in the case of calculation of the field after the relaxation, when each worker performs large number of summing up operations to account for contributions from all sub-volumes of the 3D geometry to the field value at a given point in space. This is confirmed by the fact that in our tests, the time spent by a worker on calculation of one element of the interaction matrix was only  $6.6 \mu\text{s}$ , while the calculation of the magnetic field after the relaxation by a worker at one observation point took a much longer time, 22.5 ms. These observations explain the earlier saturation of the overall speedup curve in the case of the interaction matrix generation, as compared to the field calculation after the relaxation.

Scalability of shared memory parallelization when compared to MPI, is weaker, but the resource sharing is better optimized. Preliminary tests to estimate the parallel performance of the relaxation of interaction matrix yielded a speedup of approximately 10 times compared to sequential execution of the code. The estimated overall speedup for the test geometry is  $\sim 41$  when using 400 processes, the extent of parallelization dependent on the number of field calculations performed. Further algorithm-specific implementations are required to reach higher levels of parallelism.

## CONCLUSION

The upgraded version of parallel Radia is now available for the Python interface, and it offers considerable speedups in different types of 3D magnetostatics computations [5]. Efforts are ongoing to improve its parallel performance in order to efficiently meet the computation requirements of the next generation light sources.

## REFERENCES

- [1] O. Chubar, P. Elleaume, and J. Chavanne, "A three-dimensional magnetostatics computer code for insertion devices", *J. Synchrotron Radiat.*, vol. 5, pp. 481-484, 1998. doi:10.1107/S0909049597013502
- [2] P. Elleaume, O. Chubar, and J. Chavanne, "Computing 3D Magnetic Fields from Insertion Devices", in *Proc. 17th Particle Accelerator Conf. (PAC'97)*, Vancouver, Canada, May 1997, paper 9P027, pp. 3509-3511. doi:10.1109/PAC.1997.753258
- [3] A. Andersson, O. Chubar, and L. J. Lindgren, "3D Calculations for the MAX II Lattice Magnets", in *Proc. 6th European Particle Accelerator Conf. (EPAC'98)*, Stockholm, Sweden, Jun. 1998, paper THP26F, p. 1207.
- [4] O. Chubar, C. Benabderrahmane, J. Chavanne, P. Elleaume, O. Marcouill , and F. Marteau, "Application of Finite Volume

This is a preprint — the final version is published with IOP

Content from this work may be used under the terms of the CC BY 4.0 licence (  2022). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI

Integral Approach to Computing of 3D Magnetic Fields Created by Distributed Iron-dominated Electromagnet Structures”, in *Proc. 9th European Particle Accelerator Conf. (EPAC'04)*, Lucerne, Switzerland, Jul. 2004, paper WEPKF033, pp. 1675-1677.

- [5] Radia, <https://github.com/ochubar/Radia>
- [6] C. Hall *et al.*, “Recent developments to the Radia magneto-statics code for improved performance and interface”, in *The 14th International Conference on Synchrotron Radiation Instrumentation (SRI 2021)*, Hamburg, Germany, Apr. 2022. unpublished.
- [7] Gaël Le Bec, Joël Chavanne, Simone Liuzzo, and Simon White, “Cross talks between storage ring magnets at the Extremely Brilliant Source at the European Synchrotron Radiation Facility”, *Phys. Rev. Accel. Beams*, vol. 24, p. 072401, 2021. doi:10.1103/PhysRevAccelBeams.24.072401
- [8] Sol Ji Kang, Sang Yeon Lee, and Keon Myung Lee, “Performance comparison of OpenMP, MPI, and MapReduce in practical problems”, *Advances In Multimedia*, vol. 2015, 2015. doi:10.1155/2015/575687
- [9] William D. Gropp and Ewing Lusk “High-Level Programming in MPI”, in *European Parallel Virtual Machine/Message Passing Interface Users' Group Meeting*, p. 27, 2003. doi:10.1007/978-3-540-39924-7\_7