
ARTICLE

Domain-Division Monte Carlo Dose Calculation Method for Particle Therapy

Kenichi L. ISHIKAWA^{1,2,3,*}, Koji NIITA⁴, Kazuo TAKEDA⁴,
Nobuhisa FUKUNISHI⁵ and Shu TAKAGI³

¹Department of Nuclear Engineering and Management, Graduate School of Engineering, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-8656, Japan

²Photon Science Center, Graduate School of Engineering, The University of Tokyo, Tokyo 113-8656, Japan

³RIKEN Computational Science Research Program, Wako 351-0198, Japan

⁴Research Organization for Information Science & Technology, Tokai 319-1106, Japan

⁵RIKEN Nishina Center for Accelerator-Based Science, Wako 351-0198, Japan

We present a Monte Carlo method using domain division, which enables dose calculations with huge voxel data under memory constraint. In this domain-division Monte Carlo (DDMC) method, we divide the entire geometry into several domains whose data size is such that Monte-Carlo simulations can be performed with a given amount of memory. We first perform a simulation for the domain to which the particle beam is incident and dump the information of exiting particles to a file. Then we perform simulations for the neighboring domains using the dump file as a source and, again, dump the information of exiting particles. We have developed and parallelized an external module to the Monte Carlo particle and heavy ion transport code PHITS that tallies the dose by repeating these procedures. The increase in calculation time due to dumping and re-reading is typically smaller than 10%, which is sufficiently small. We have also obtained good weak scaling up to 1,024 processors so far. The method we have developed is expected to open a way to the realization of large-scale Monte Carlo dose calculations on general-purpose supercomputers and in a grid environment, where only a limited amount of memory is available, instead of on specialized large-memory machines.

KEYWORDS: Monte Carlo, particle therapy, PHITS, dose calculation, domain division

I. Introduction

Particle therapy is a form of external beam radiotherapy typically using energetic proton or carbon-ion beams, and has recently been more and more used due to the excellent dose localization and, for the case of carbon-ion therapy, high relative biological effectiveness. In particle therapy, secondary particles such as neutrons, protons, alpha and heavier particles are produced through nuclear reactions. These secondary particles produced in the patient body represent an unavoidable internal radiation source. With increasing population of cancer survivors previously exposed to radiation, secondary cancer risk estimation is rapidly becoming a key issue. For the case of particle therapy, in particular, secondary neutrons can deposit dose at large distances from the tumor, which necessitates dose calculation in the entire patient body. Indeed, several groups^{1,2)} have performed dose calculation for whole-body computational voxel phantoms using Monte Carlo method, which has been widely accepted to be the way to provide the most accurate dose calculations.³⁾

With increase of CT imaging resolution, such whole-body dose calculations require a huge amount of memory. For 1 mm × 1 mm × 1 mm resolution, e.g., a whole-body phan-

tom contains more than 200 millions of voxels, requiring several gigabytes per processing element (PE). This means that a dodeca-core workstation equipped with two hexa-core CPUs needs several-tens-of-gigabyte memory, which is beyond the specification of most parallel clusters. There is also a project to construct whole-body phantoms with even finer resolution (0.3 mm × 0.3 mm × 0.3 mm), for which Monte Carlo dose calculation would be impossible not due to the limitation of available CPU time but due to that of the amount of available memory.

In this work, we describe a Monte Carlo method that enables dose calculations with huge voxel data under memory constraint by dividing the entire calculation region into multiple domains. In this method, we first perform a simulation for the domain to which the particle beam is incident and store the phase-space information of exiting particles to a file. Then we perform simulations for the neighboring domains using the stored file as a source and, again, dump the phase-space information of exiting particles. We have developed and parallelized an external module to the Monte Carlo particle and heavy ion transport code PHITS⁴⁾ that tallies the dose by repeating these procedures.

II. Method

In this method, we divide the entire geometry into several

*Corresponding author, E-mail: ishiken@riken.jp

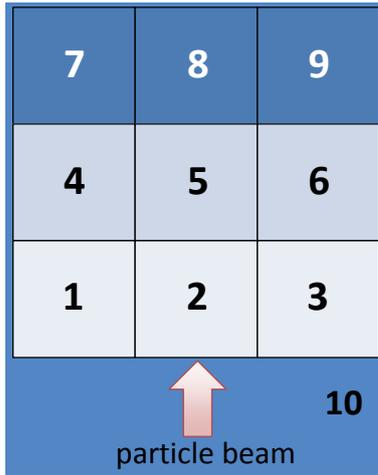


Fig. 1 Schematic of the domain division

domains whose data size is such that Monte-Carlo simulations can be performed with a given amount of memory, as is schematically depicted in **Fig. 1**. While division is made in three dimensions in practice, we show only two dimensions in this figure for simplicity. Let us assume that the particle beam is incident to domain 2. We perform dose calculation with the following procedures:

1. *Generation 1*: we first perform a simulation only for domain 2, to which the particle beam is incident, tally dose in the domain, and dump the phase-space information of exiting particles to a file. PHITS is equipped with such dumping functionality.
2. *Generation 2*: Using the phase-space file created in Generation 1 as a source (this functionality is implemented in PHITS), we perform a simulation and tally dose for each of the neighboring domains (domains 1, 3, and 5), and, again, dump the phase-space information of exiting particles.
3. *Generation 3*: Using the phase-space files created in Generation 2 as a source, we perform a simulation and tally dose for each of the neighboring domains (domains 2, 4, 6, and 8), and, again, dump the phase-space information of exiting particles.
4. *Generation 4*: Using the phase-space files created in Generation 3 as a source, we perform a simulation and tally dose for each of the neighboring domains (domains 1, 3, 5, 7, and 9), and, again, dump the phase-space information of exiting particles.
5. *Generation 5*: Using the phase-space files created in Generation 4 as a source, we perform a simulation and tally dose for each of the neighboring domains (domains 2, 4, 6, and 8), and, again, dump the phase-space information of exiting particles
6. We repeat similar procedures till there are no dumped particles left.

We have developed and parallelized an external module to the Monte Carlo particle and heavy ion transport code

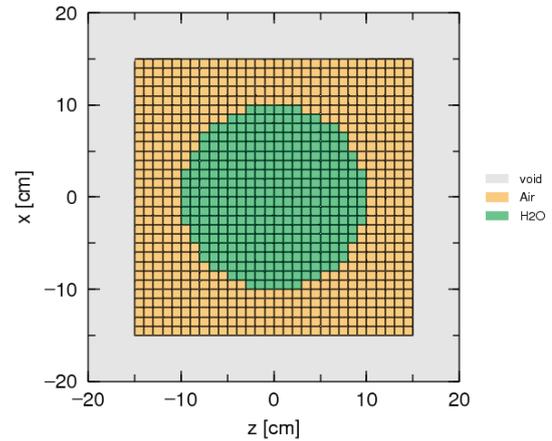


Fig. 2 Schematic representation of the water sphere phantom

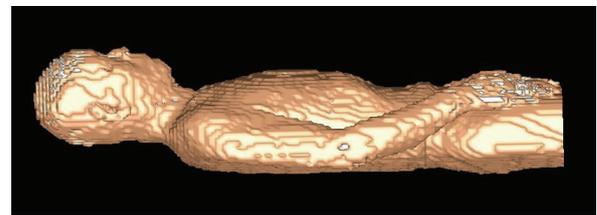


Fig. 3 Side view of the human voxel phantom

PHITS⁴) that tallies the dose by these procedures. It should be noted that, even in a parallel environment, different domains are not distributed to different PEs but simulated one after another by each PE.

We evaluate the performance of the domain-division Monte Carlo (DDMC) method using two different geometries. The one (**Fig. 2**) is a geometry in which a water sphere with a radius of 10 cm is located in the center of an air cube with an edge of 30 cm (referred to as water sphere phantom hereafter), represented with $1 \times 1 \times 1 \text{ cm}^3$ voxels. The entire region is divided into $3 \times 3 \times 3$ domains. We consider 130 MeV proton cylindrical beam (1 mm radius) incidence. The other (**Fig. 3**) is a human voxel phantom made up of $5 \times 5 \times 5 \text{ mm}^2$ voxels and twenty different materials. We consider only the upper half body and divide the entire region into $4 \times 2 \times 3$ domains. We consider 200 MeV proton pencil-beam incidence from the left-hand side of the figure.

III. Results

1. Progress of the Domain-Division Monte Carlo Simulations

In this subsection, we show how the DDMC calculation proceeds. **Figure 4** shows the dose distribution in the first, second, and third generation, for the case of 130 MeV proton incidence (10^6 histories) to the water sphere phantom. The calculation reached the tenth generation, and 93 individual Monte Carlo simulations were called in total. The simulation was performed on a single processing element. **Figure 5** compares the total doses calculated with the two methods on

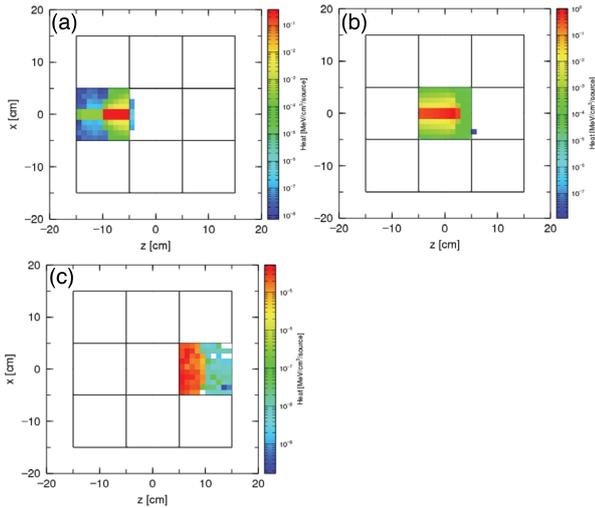


Fig. 4 Progress of the DDMC for 130 MeV proton incidence (10^6 histories) to the water phantom (see text). Physical dose distribution in the (a) first, (b) second, and (c) third generation. The false color is in a logarithmic scale and not common for the three panels.

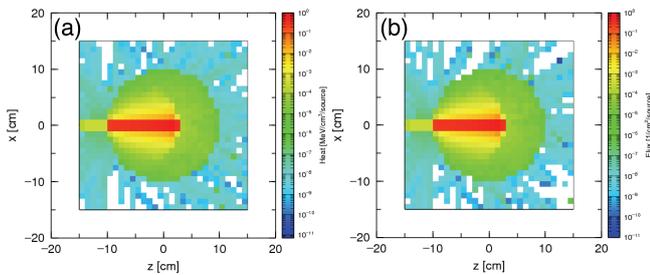


Fig. 5 Dose distribution in the water phantom (see text), calculated with (a) the ordinary Monte Carlo method and (b) the domain-division Monte Carlo method, with 10^6 histories

a single processing element. We can clearly see that both methods give the same results within a statistical error. **Figure 6** shows the results for the human voxel phantom. The parallel simulation was performed on 32 PEs.

2. Computational Time Increase Due to Dumping and Re-Reading

It is expected that the DDMC calculation takes longer than the ordinary Monte Carlo simulation due to dumping and re-reading of phase-space files. In **Table 1**, we examine increase in computational time for the case of Figs. 4 and 5. Although 93 individual PHITS processes were called, involving 92 processes of dumping and re-reading, the computational time increased only by 8%. **Table 2** shows the results of a similar comparison for the voxel phantom. Again, the increase in computational time is typically smaller than 10%, while it depends on the voxel data size and the number of histories. These values are sufficiently small for the DDMC method to be useful in practical situations, especially if we consider that the cost of additional memory is usually higher than that of additional CPUs.

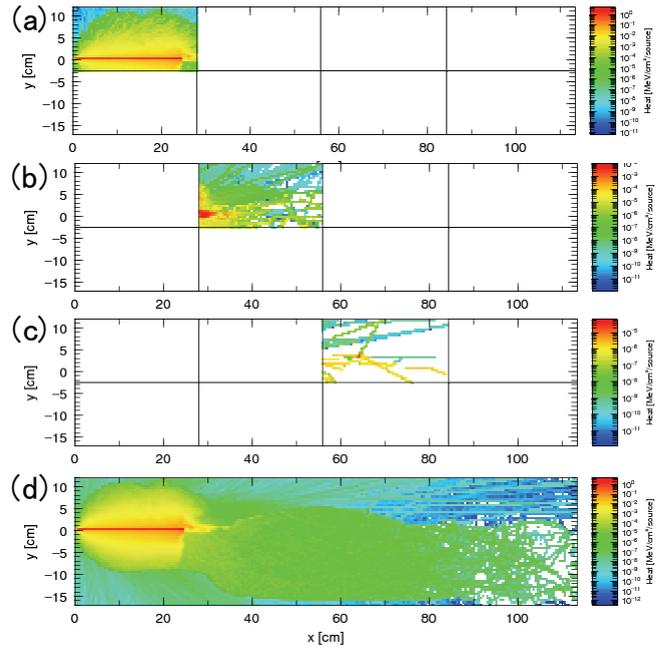


Fig. 6 Progress of the DDMC for 200 MeV proton incidence (1.6×10^6 histories) to the human voxel phantom. A pencil beam is incident from the left-hand side. *xy* tally physical dose distribution in the (a) first, (b) second, and (c) third generations as well as (d) the total dose distribution. The false color is in a logarithmic scale and not common for the four panels.

Table 1 Comparison of the computational time between the ordinary and domain-division Monte Carlo simulations for the water sphere phantom

Method	Number of domains	Number of generations	Number of called PHITS processes	Computational time (sec)
Ordinary	1	1	1	1174
Domain-division	27	10	93	1268

Table 2 Comparison of the computational time between the ordinary and domain-division Monte Carlo simulations for the voxel phantom

Number of histories	Computational time (h:mm:ss)		Increase (%)
	Ordinary	Domain-division	
5×10^4	0:05:28	0:06:20	15.9
1.024×10^6	1:10:56	1:15:48	6.9
2.048×10^6	2:19:39	2:27:56	5.9
4.096×10^6	4:37:02	4:51:59	5.4
8.192×10^6	9:12:32	9:40:10	5.0

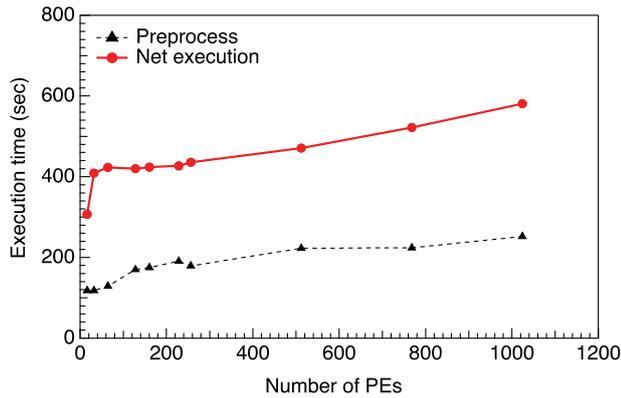


Fig. 7 Measured weak scaling between 16 and 1,024 PEs

3. Weak Scaling of Parallel Computation

As shown in Fig. 6, the DDMC dose calculation is successfully done on parallel computers. In this subsection, let us examine the weak scaling, defined as how the solution time varies with the number n of PEs for a fixed problem size (number of histories) per PE. Since parallelized Monte Carlo simulations require virtually no communication between PEs, we usually expect good weak scaling.

In Fig. 7 we show the results of the measurements using the QUEST PC clusters at RIKEN Computation Science Research Program. 5×10^4 histories per PE were simulated for the same condition as for Fig. 6. The time elapsed for preprocessing (including the input and nuclear data file transfer) and for the net calculation (including the output file transfer). The former depends on the network bandwidth and architecture of the system. From the curve for the net execution time, we see that, although the DDMC method exhibits a relatively good weak scaling, the solution time gradually increases with n .

One of the reasons is that the number of simulated generations is not necessarily the same for all the PEs; it can happen that there are no more exiting particles at the end of, say, the i -th generation and the execution is terminated for one PE, while there are still a few exiting particles, necessitating the next generation, for another PE. Such rare events that require many generations are more likely to appear for a larger number of histories, i.e., for a larger number of PEs. An improvement in load balancing and/or, considering a plausibly negligible contribution to the total dose from such histories, an introduction of some sort of cutoff generation number would be desirable in this respect. Another possible cause of the weak scaling degradation is the loss induced when multiple processing elements access files common to all the PEs such as a material information data file, material

color specification file, parameter data file, tally data files, and voxel data files.

IV. Conclusion

We have presented a new type of a Monte Carlo method, which divides the entire geometry into several domains and performs simulation for each domain successively and iteratively. We have successfully developed and parallelized an external module to PHITS, which delivers the same results as ordinary Monte Carlo runs within a statistical error. The overhead due to the dumping and re-reading of phase-space files is reasonably low. We have observed a good weak scaling in general, though performance degradation is seen for an increasing number of PEs and, thus, further tuning is desirable.

The method we have developed is expected to open a way to the realization of large-scale Monte Carlo dose calculations using high-resolution CT data and computational voxel phantoms on general-purpose supercomputers and in a grid environment, where only a limited amount of memory is available, without need of specialized large-memory machines.

Acknowledgment

This research was supported by Research and Development of the Next-Generation Integrated Simulation of Living Matter, a part of the Development and Use of the Next-Generation Supercomputer Project of MEXT, Japan, and also partially supported by RIKEN President's Discretionary Fund (Strategic Programs for R&D). KLI acknowledges financial support by the Ministry of Education, Science, Sports and Culture, Grant-in-Aid for Exploratory Research, 21656240, 2009-2010.

References

- 1) T. Sato, Y. Kase, R. Watanabe, K. Niita, L. Sihver, "Biological dose estimation for charged-particle therapy using an improved PHITS code coupled with a microdosimetric kinetic model," *Radiat. Res.*, **171**, 107-117 (2009).
- 2) B. Athar, K. Henker, O. Jäkel, N. Bassler, H. Paganetti, "Comparison of out-of-field neutron equivalent doses in scanning carbon and proton therapies for cranial fields," *Med. Phys.*, **37**, 3281 (2010).
- 3) H. Jiang, H. Paganetti, "Adaptation of GEANT4 to Monte Carlo dose calculations based on CT data," *Med. Phys.*, **31**, 2811-2818 (2004).
- 4) K. Niita, T. Sato, H. Iwase, H. Nose, H. Nakashima, L. Sihver, "PHITS-a particle and heavy ion transport code system," *Radiat. Meas.*, **41**, 1080-1090 (2006).