



HIJING++, a Heavy Ion Jet INteraction Generator for the High-Luminosity Era of the LHC and Beyond [†]

Gábor Bíró ^{1,2,*} , Gábor Papp ¹ , Gergely Gábor Barnaföldi ² , Dániel Nagy ¹,
Miklos Gyulassy ^{2,3,4,5} , Péter Lévai ², Xin-Nian Wang ^{3,4} and Ben-Wei Zhang ³

¹ Institute for Physics, Eötvös Loránd University, 1/A Pázmány P. Sétány, H-1117 Budapest, Hungary; pg@ludens.elte.hu (G.P.); nagy.dani@wigner.mta.hu (D.N.)

² Wigner Research Centre for Physics of the Hungarian Academy of Sciences, 29-33 Konkoly-Thege Miklós Str, H-1121 Budapest, Hungary; barnafoldi.gergely@wigner.mta.hu (G.G.B.); gyulassy@phys.columbia.edu (M.G.); levai.peter@wigner.mta.hu (P.L.)

³ Key Laboratory of Quark & Lepton Physics (MOE) and Institute of Particle Physics, Central China Normal University, Wuhan 430079, China; xnwang@lbl.gov (X.-N.W.); bwzhang@mail.ccnu.edu.cn (B.-W.Z.)

⁴ Nuclear Science Division, MS 70R0319, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

⁵ Pupin Lab MS-5202, Department of Physics, Columbia University, New York, NY 10027, USA

* Correspondence: biro.gabor@wigner.mta.hu; Tel.: +36-30-308-6742

[†] Presented at Hot Quarks 2018-Workshop for Young Scientists on the Physics of Ultrarelativistic Nucleus-Nucleus Collisions, Texel, The Netherlands, 7–14 September 2018.

Published: 10 April 2019



Abstract: HIJING++ (Heavy Ion Jet INteraction Generator) is the successor of the widely used original HIJING, developed almost three decades ago. While the old versions (1.x and 2.x) were written in FORTRAN, HIJING++ was completely rewritten in C++. During the development we keep in mind the requirements of the high-energy heavy-ion community: the new Monte Carlo software have a well designed modular framework, therefore any future modifications are much easier to implement. It contains all the physical models that were also present in it's predecessor, but utilizing modern C++ features it also includes native thread based parallelism, an easy-to-use analysis interface and a modular plugin system, which makes room for possible future improvements. In this paper we summarize the results of our performance tests measured on 2 widely used architectures.

Keywords: high energy physics; heavy-ion; Monte Carlo; event generator; parallel computing; HIJING

PACS: 24.10.Lx; 25.75.-q; 25.75.Ag; 25.75.Dw

1. Introduction

During the approaching Long Shutdown 2 (LS2) of the Large Hadron Collider (LHC) in 2019–2020 many technical improvement will occur in the accelerator complex, in the detector and in the data acquisition systems. These will result in a huge increase of the number of expected collisions per second and also the amount of measured data per event will grow rapidly. This period is the forerunner of the next generation of particle accelerators, such as the High-Luminosity LHC (HL-LHC) or the Future Circular Collider (FCC), where we will accumulate high-energy experimental data in a higher rate than ever. In parallel we need to improve also the numerical tools in order to be able to keep up the requisites of the high-precision era.

The new HIJING++ heavy-ion Monte Carlo framework is written from scratch with a modular, effective C++ structure and with built-in CPU based parallelism in order to fulfill these requirements. Though the program flow is based on the original FORTRAN HIJING[1,2], the design is completely revised so the main components of the program can work together effectively. Such components are the most recent versions of PYTHIA8 [3] (used for the hard scattering processes and for the hadronization), LHAPDF6 [4], the GNU Scientific Library [5,6] (utilizing the VEGAS multi-dimensional Monte Carlo integration), and the CERN ROOT [7] data analysis software along with the HijAnalysis data collector framework.

HIJING++ is intended to work effectively regarding different aspects, not just based on the raw performance of the CPU. As an example, it is possible to replace any of the main components, such as the jet quenching and shadowing algorithms, in a convenient, well defined way, without modifying the core code. An another built-in feature is the above mentioned HijAnalysis framework, which adds the possibility to define any kind of data collecting objects, such as ROOT TTrees, histograms or simple ASCII files to collect all final state particles event-by-event. Utilizing modern C++ features, the result of a run will be data structures that can be further processed in a convenient way.

In the following section we present the results of the performance tests of the pre-release version of HIJING++, taking advantage of these features.

2. Results

We have already presented preliminary physics and performance results in Ref. [8,9]. Here we summarize the benchmark tests measured on two different machines.

2.1. Benchmark Setups

In order to measure the performance in a real case situation, we calculated 6 different histograms to collect various quantities of the current run, such as the impact parameter, number of binary collisions, event multiplicity, p_T spectra and pseudorapidity distributions of different identified hadrons with various binnings. We performed each run several times in order to reduce fluctuations. The main parameters of the different run setups are summarized in Table 1 [10,11].

Table 1. The run setups and their main parameters.

Collision System	Event Number	(n)PDF
pp	10^6	CT14n1o [10]
p-Pb	10^4	CT14n1o (for protons), EPPS16n1o_CT14n1o_Pb208 [11] (for lead nuclei)
Pb-Pb	10^3	EPPS16n1o_CT14n1o_Pb208

The tests were made on 2 commonly used, typical architectures, whose parameters are listed in Table 2 [12,13]. These setups represents common use cases in the heavy-ion community: CPUs with lower TDP values (*thermal design power*—the higher the value, the larger the power consumption and performance) and its variants are widely used in recent laptops and ultrabooks, while CPUs with higher TDP are common in desktop computers or larger workstations, clusters.

Table 2. The computer architectures [12,13] used to measure run performance.

CPU Type	Release Year	Number of Cores (Threads)	Base (Turbo) Frequency	TDP	RAM
Intel(R) Core(TM) i5-8250U	Q3'17	4 (8)	1.6 GHz (3.4 GHz)	15 W	8 GB
Intel(R) Xeon(TM) E3-1231 v3	Q2'14	4 (8)	3.4 GHz (3.8 GHz)	80 W	32 GB

2.2. Results

The results of the benchmarking runs for the two different CPUs are shown on Figure 1. As expected, the measured times show significant differences between the two systems: using the CPU with the lower TDP value (*upper panels*) by increasing the number of threads the total runtime decreases significantly until $N_{thread} = 4$, then the speedup gained from the multiple threads is compensated by the fact that more CPU cores have to share the same amount of energy, resulting in a decrease of the CPU frequency. In accordance with this, the initialization time increases slightly with the increasing thread number. In contrast to these, on the *lower panels* the results achieved with the higher performance desktop/server CPU are shown, where the speedup is more significant with the higher number of threads. In this case, the initialization time increases with a much lower rate. The reason is that this CPU doesn't have to decrease the performance when we are operating with multiple cores.

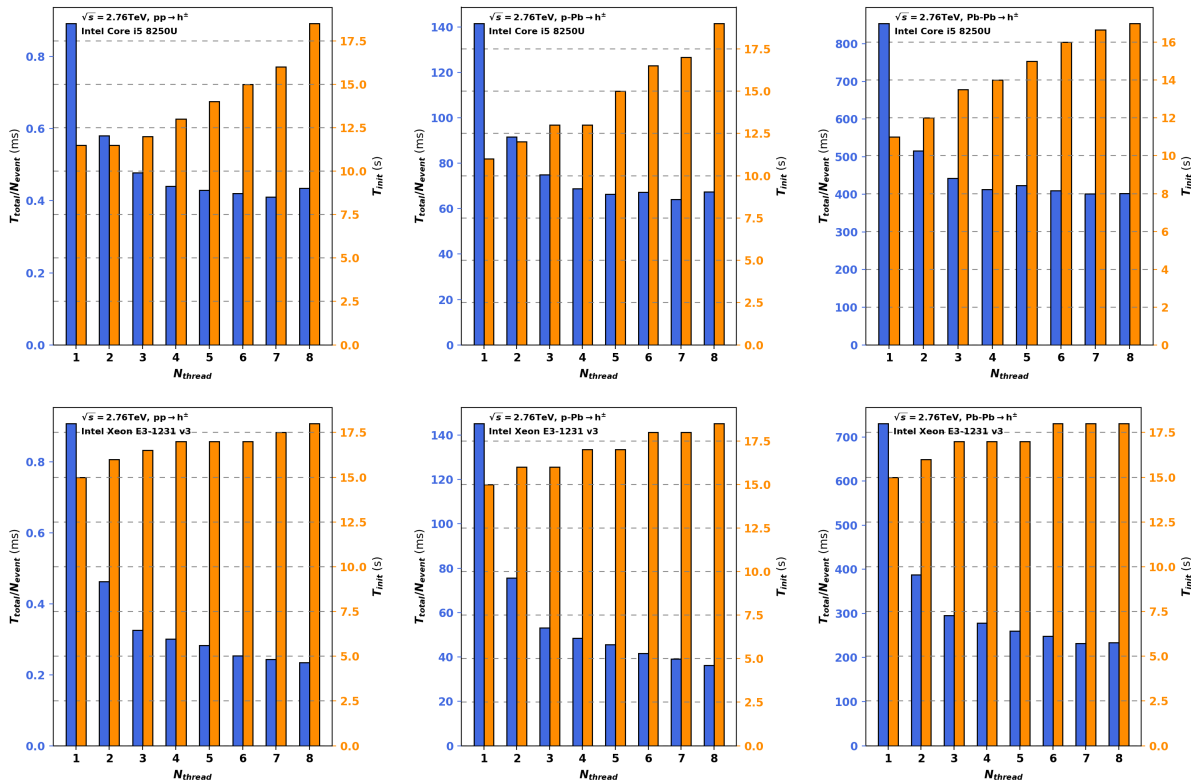


Figure 1. The total runtime normalized with the event number (blue bars) and the initialization time (orange bars) versus the CPU threads, measured on an Intel(R) Core(TM) i5-8250U CPU (**upper panels**) and on an Intel(R) Xeon(TM) E3-1231 v3 CPU (**lower panels**). The run parameters: $\sqrt{s} = 2.76$ ATeV proton-proton (**left panels**), proton-lead (**middle panels**) and lead-lead collisions (**right panels**), using (nuclear) parton distribution functions CT14n1o (for protons) and EPPS16n1o_CT14n1o_Pb208 (for lead nuclei), defining 6 different histogram analysis objects.

By fitting the measured results with Amdahl's law [14] we can determine the maximum theoretical speedup compared to the single thread run that can be achieved on the specific architecture:

$$\text{Speedup}(N_{threads}) = \frac{N_{threads}}{1 + \alpha(N_{threads} - 1)} \quad (1)$$

where α is the non-parallelizable part of the code. According to the results summarized in Table 3 the scalability on the higher performance CPU is better, the non-parallelizable parts (such as the thread managing system itself) result in a lower α value. However, using 3–4 threads HIJING++ runs more efficiently also with the low TDP CPU, resulting in a considerably reduced runtime.

Table 3. The maximum theoretical speedup compared to the single thread and the non-parallelizable part α of the code.

CPU	Speedup			α		
	pp	p-Pb	Pb-Pb	pp	p-Pb	Pb-Pb
Intel(R) Core(TM) i5-8250U	2.6	2.7	2.6	0.38	0.37	0.38
Intel(R) Xeon(TM) E3-1231 v3	6.4	6.6	4.5	0.16	0.15	0.22

In order put the performance of HIJING++ into context, we measured and compared the (single thread) runtime of PYTHIA8.2 and HIJING v2.552. We found that HIJING++ is ~30% faster than PYTHIA8.2 and ~50% slower than HIJING v2.552.

This is not a surprising result, because the published FORTRAN HIJING was originally written with single precision floating point numbers: on one hand, this can lead to significant numerical errors (especially at LHC energies) when performing calculations with frequently occurring small quantities like $\sim \frac{m_q}{\sqrt{s}} \ll 1$, where m_q is the mass of a given quark species and \sqrt{s} is the center-of-mass energy. On the other hand, we measured the effect of modifying the FORTRAN HIJING into double precision, and we found that in such case it's runtime scales up by a factor of 4.

3. Summary and Conclusions

We presented the results of the performance benchmarks of the new HIJING++ heavy-ion Monte Carlo event generator using different CPUs and collision systems. Utilizing the built-in CPU parallelization and analysis frameworks HIJING++ provides a significant decrease in the necessary computation time which is especially important at higher performance architectures. In the future developments further optimizations are planned to improve the scalability.

Author Contributions: G.B. developed the software framework and wrote the first version of the manuscript. Authors G.P., G.G.B., M.G., X.N.W., B.W.Z. and P.L. supervised the development, provided theoretical background and reviewed the manuscript. D.N. developed the benchmarking framework.

Funding: This research was funded by Hungarian–Chinese cooperation grant No. MOST 2014DFG02050 and Wigner HAS-OBOR-CCNU grant; OTKA grants K120660, K123815, THOR COST action CA15213. Author G.B. acknowledge the support of Wigner Data Center and Wigner GPU Laboratory.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Wang, X.N.; Gyulassy, M. Hijing: A Monte Carlo model for multiple jet production in pp, pA, and AA collisions. *Phys. Rev. D* **1991**, *44*, 3501.
2. Deng, W.T.; Wang, X.N.; Xu, R. Hadron production in p+ p, p+ Pb, and Pb+ Pb collisions with the HIJING 2.0 model at energies available at the CERN Large Hadron Collider. *Phys. Rev. C* **2011**, *83*, 014915.
3. Sjöstrand, T. An Introduction to PYTHIA 8.2. *Comput. Phys. Commun.* **2015**, *191*, 159.
4. Buckley, A.; Ferrando, J.; Lloyd, S.; Nordström, K.; Page, B.; Rüfenacht, M.; Schönherr, M.; Watt, G. LHAPDF6: parton density access in the LHC precision era. *Eur. Phys. J. C* **2015**, *75*, 132.
5. Galassi, M.; Davies, J.; Theiler, J.; Gough, B.; Jungman, G.; Alken, P.; Booth, M.; Rossi, F.; Ulerich, R. *GNU Scientific Library Reference Manual*, 3rd ed.; Network Theory Ltd.: Bristol, UK, 2009; ISBN 0954612078.
6. Lepage, G.P. A new algorithm for adaptive multidimensional integration. *J. Comput. Phys.* **1978**, *27*, 192–203.

7. Available online: <https://root.cern.ch/> (accessed on 25 October 2018).
8. Barnaföldi, G.G.; Bíró, G.; Gyulassy, M.; Haranózó, S.M.; Lévai, P.; Ma, G.; Papp, G.; Wang, X.N.; Zhang, B.W. First Results with HIJING++ in High-Energy Heavy-Ion Collisions. *Nucl. Part. Phys. Proc.* **2017**, *289*, 373–376.
9. Papp, G.; Barnaföldi, G.G.; Bíró, G.; Gyulassy, M.; Harangozó, S.M.; Ma, G.; Lévai, P.; Wang, X.N.; Zhang, B.W. First Results with HIJING++ on High-energy Heavy Ion Collisions. *arXiv* **2018**, arXiv:1805.02635.
10. Dulat, S.; Hou, T.J.; Gao, J.; Guzzi, M.; Huston, J.; Nadolsky, P.; Pumplin, J.; Schmidt, C.; Stump, D.; Yuan, C.P. New parton distribution functions from a global analysis of quantum chromodynamics. *Phys. Rev. D* **2016**, *93*, 033006.
11. Eskola, K.J.; Paakkinen, P.; Paukkunen, H.; Salgado, C.A. EPPS16: nuclear parton distributions with LHC data. *Eur. Phys. J. C* **2017**, *77*, 163.
12. Available online: <https://ark.intel.com/products/80910/Intel-Xeon-Processor-E3-1231-v3-8M-Cache-3-40-GHz-> (accessed on 25 October 2018).
13. Available online: <https://ark.intel.com/products/124967/Intel-Core-i5-8250U-Processor-6M-Cache-up-to-3-40-GHz-> (accessed on 25 October 2018).
14. Amdahl, G.M. Validity of the single processor approach to achieving large scale computing capabilities. In Proceedings of the AFIPS Conference, Atlantic City, NJ, USA, 18–20 April 1967; Volume 30, p. 483.



© 2019 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).