

# Up to 700k GPU cores, Kepler, and the Exascale future for simulations of star clusters around black holes

Peter Berczik<sup>1,2,3</sup>, Rainer Spurzem<sup>1,3</sup>,  
Long Wang<sup>4</sup>, Shiyan Zhong<sup>1</sup>, Siyi Huang<sup>1</sup>, Maxwell Xu Tsai<sup>1</sup>,  
Gareth Kennedy<sup>1</sup>, Shuo Li<sup>1</sup>, Luca Naso<sup>1</sup>, Changhua Li<sup>1</sup>,  
Alexander Veles<sup>2</sup>, Igor Zinchenko<sup>2</sup>

<sup>1</sup>National Astronomical Observatories, Chinese Academy of Sciences, 20A Datun Rd., Chaoyang Distr.,  
Beijing 100012, P.R. China

<sup>2</sup>Main Astronomical Observatory, National Academy of Sciences of Ukraine, MAO/NASU, 27 Akademika  
Zabolotnoho St. 03680 Kyiv, Ukraine

<sup>3</sup>Zentrum für Astronomie, Astronomisches Rechen-Institut, Univ. of Heidelberg, Monchhofstr. 12-14, 69120  
Heidelberg, Germany

<sup>4</sup>Kavli Inst. for Astronomy and Astrophysics, Peking University, Yi He Yuan Rd. 5, Haidian Distr., Beijing  
100871, P.R. China

berczik@mao.kiev.ua, berczik@nao.cas.cn, spurzem@nao.cas.cn

**Abstract.** We present direct astrophysical N-body simulations with up to a few million bodies using our parallel MPI/CUDA code on large GPU clusters in China, Ukraine and Germany, with different kinds of GPU hardware. These clusters are directly linked under the Chinese Academy of Sciences special GPU cluster program in the cooperation of ICCS (International Center for Computational Science). We reach about the half the peak Kepler K20 GPU performance for our  $\varphi$ -GPU<sup>1</sup> code [2], in a real application scenario with individual hierarchically block time-steps with the high (4<sup>th</sup>, 6<sup>th</sup> and 8<sup>th</sup>) order Hermite integration schemes and a real core-halo density structure of the modeled stellar systems. The code and hardware are mainly used to simulate star clusters [23, 24] and galactic nuclei with supermassive black holes [20], in which correlations between distant particles cannot be neglected.

## Keywords

N-body simulation, parallel computing, many core, GPU acceleration, star clusters, galactic nuclei, black hole physics, astrophysics.

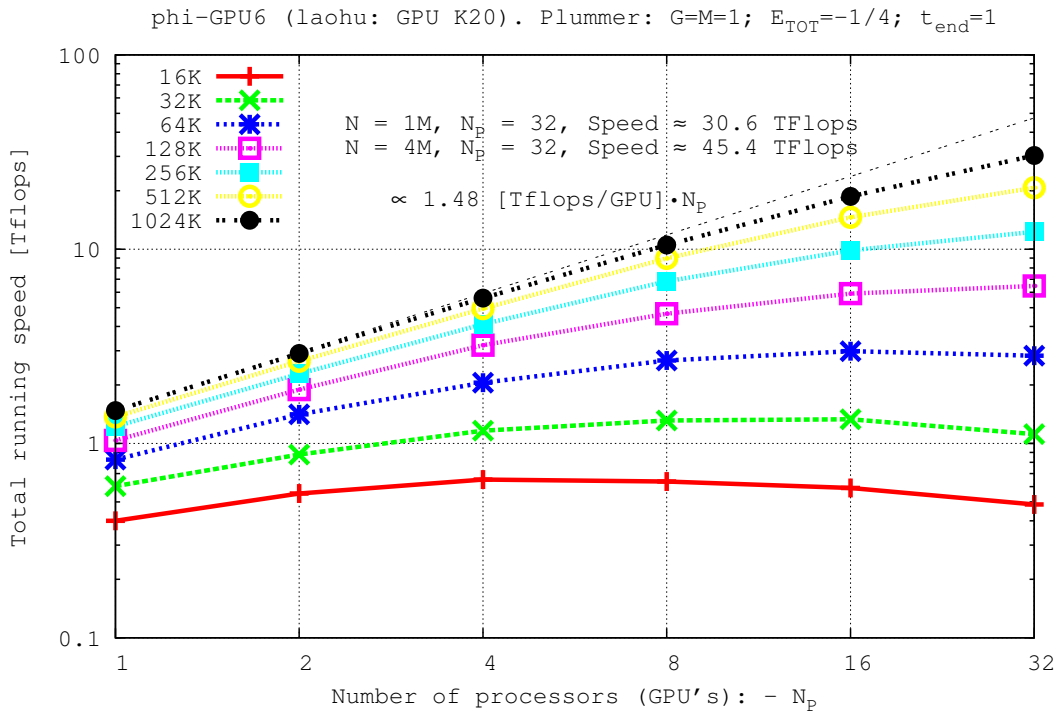
## 1 Introduction

Many, if not all galaxies harbor supermassive black holes [8, 10]. If galaxies merge, which is quite common in the process of hierarchical structure formation in the universe, their black holes sink to the center of the merger remnant and form a tight binary [7, 9, 18, 19]. Depending on initial conditions and time supermassive black hole binaries are prominent gravitational wave sources, if they ultimately come close together and coalesce [11, 13, 14]. We model such systems as gravitating N-body systems (stars) with two or more massive bodies (black holes), including if necessary relativistic corrections to the classical Newtonian gravitational forces, and model their gravitational radiation emission directly from the simulation [12, 15, 16, 17].

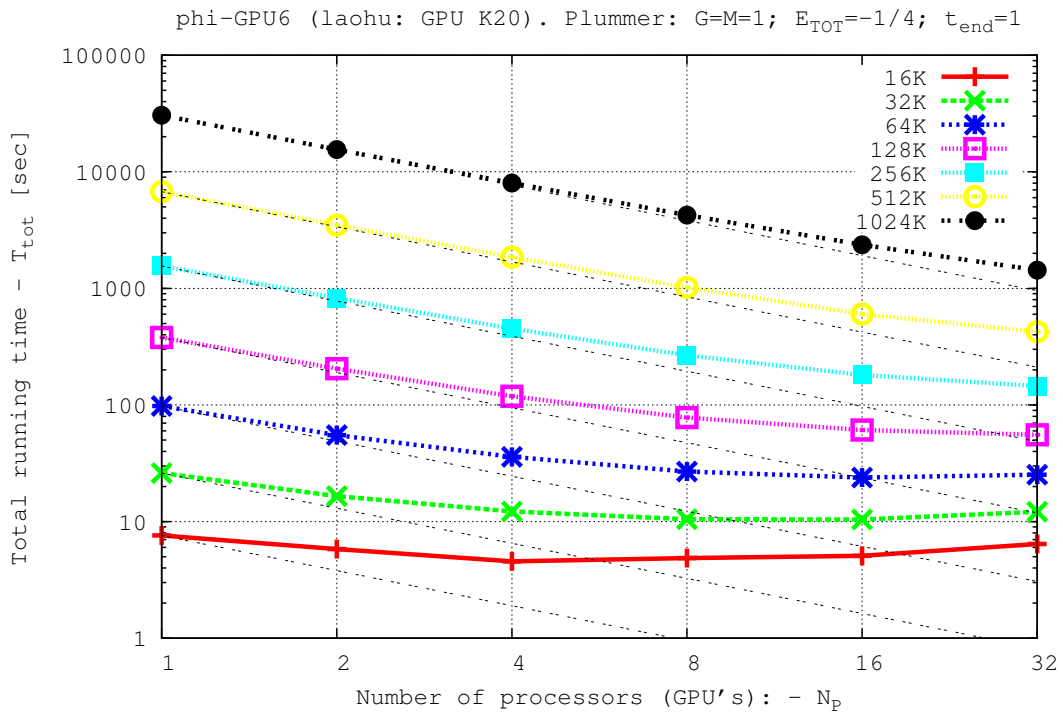
Competitive astronomical and astrophysical research requires access to competitive computing facilities. Theoretical numerical modeling of astrophysical objects, their composition, radiation, and dynamical evolution has become a third basic method of astrophysical research, besides observation and pure theory [21]. Numerical modeling allows one to compare theory with observational data in unprecedented detail, and it also provides theoretical insight into physical processes at work in complex systems [22]. Similarly, data processing of astrophysical observations comprises the use of complex software pipeline to bring raw data into a form digestible for

---

<sup>1</sup>[ftp://ftp.mao.kiev.ua/pub/berczik/phi-GPU/](http://ftp.mao.kiev.ua/pub/berczik/phi-GPU/)



**Figure 1.** Speed performance with mixed (fp32 + fp64) precision of the  $\varphi$ -GPU 6<sup>th</sup> order scheme on the K20 GPU cards. The lines with different symbols presents the different particle numbers.



**Figure 2.** Total wall clock time of 1 time unit integration with the  $\varphi$ -GPU 6<sup>th</sup> order scheme on the K20 GPU cards. The lines with different symbols presents the different particle numbers.

observational astronomers and ready for exchange and publication; these are, e.g., mathematical transformations like Fourier analysis of time series or spatial structures, complex template analysis or huge matrix-vector operations. Here fast access to and transmission of data, require supercomputing capacities.

We are undergoing a new revolution on parallel processor technologies, especially with regard to the Graphic Processing Units. GPU's have become widely used nowadays to accelerate a broad range of applications, including computational physics and astrophysics, image/video processing, engineering simulations, quantum chemistry<sup>2</sup>. Graphics processing units (GPU's) are rapidly emerging as a powerful and cost-effective platform for high performance parallel computing. Recent GPU's, such as the NVIDIA Fermi C2050 and Kepler K20 Computing Processors offer correspondingly 448 and 2496 processor cores with the extremely fast on-chip-memory, as compared to only 6 or 8 cores on a standard Intel or AMD CPU. In this paper we present the set of large scale N-body benchmarks of our astrophysical applications on a large Kepler and Fermi based GPU clusters.

## 2 Codes description

The application codes which we use for benchmarking here are direct N-body simulation codes for astrophysics, both using a high order Hermite integration scheme and hierarchical block time steps. One is called  $\varphi$ -GPU [2, 1], it has been developed from our earlier published versions  $\varphi$ -GRAPE<sup>3</sup> (which used GRAPE hardware as a accelerator [3]), the other is a legacy code named NBODY6++GPU [4, 5], which has only recently been accelerated by GPU. Both use MPI (Message Passing Interface), and on each node use many cores of the special GPU hardware to compute gravitational forces between particles. Parts of the codes were developed in cooperation with Keigo Nitadori (RIKEN Japan) and Tsuyoshi Hamada (Nagasaki Univ. Japan).  $\varphi$ -GPU [2, 1] is written in C++ with MPI and CUDA-C. NBODY6++GPU is written in Fortran with CUDA-C extensions [6]. We present benchmarks on different GPU accelerated hardware, including Fermi and Kepler architectures.

## 3 Large scale astrophysical results

We report results obtained with our parallel direct N-body code on the CAS GPU clusters at our institute NAOC/CAS and at the IPE/CAS. We are also planning to run benchmarks on the Oak Ridge Titan supercomputer (using its Kepler K20 GPU's), but at the time of writing we cannot foresee whether they will be ready for the conference.

The NAOC cluster (85 nodes) `1aohu` is now, after a recent upgrade, running with the new Kepler K20 GPU's (each with 2496 cores), and has in total roughly  $\sim 170k$  computing GPU cores. The IPE/CAS system `mole-8.5` is a much larger system with 362 nodes (each with 6 Fermi C2050 accelerators) and has in total almost  $\sim 1M$  (973k) GPU cores. On the newest Kepler K20 GPU's we reach almost 1.5 Tflop/s sustained speed per one GPU using the 6<sup>th</sup> order  $\varphi$ -GPU code (see **Figure 1.** and **Figure 2.**). On the Fermi C2050 cards on the `mole-8.5` system we get the sustained speed of around 550 Gflop/s per one GPU with the same  $\varphi$ -GPU code (see **Figure 5.** and **Figure 6.**). These numbers are roughly close to the 50% of the peak performance of such a GPU cards, which shows a very good utilization of the GPU hardware with our high order direct N-body codes. Note that for the largest simulation conducted for this system used 1536 Fermi C2050 accelerators each with 448 computing cores totaling almost 700k GPU cores of the `mole-8.5` system.


After our extensive scaling measurement on the `mole-8.5` system, which we have presented in our earlier paper [2], we make a detail performance analysis of our  $\varphi$ -GPU code and derive the "semi-analytic" scaling formula for the performance (see also the equations (17), (18) and (19) in the paper [2]):

$$P \approx \frac{\gamma N^{2+x}}{\frac{\alpha N^{2+x}}{N_{GPU}} + \beta(N^{0.33+x}\tau_{lat} + N^{1+x})\log(N_{GPU})} \quad (1)$$

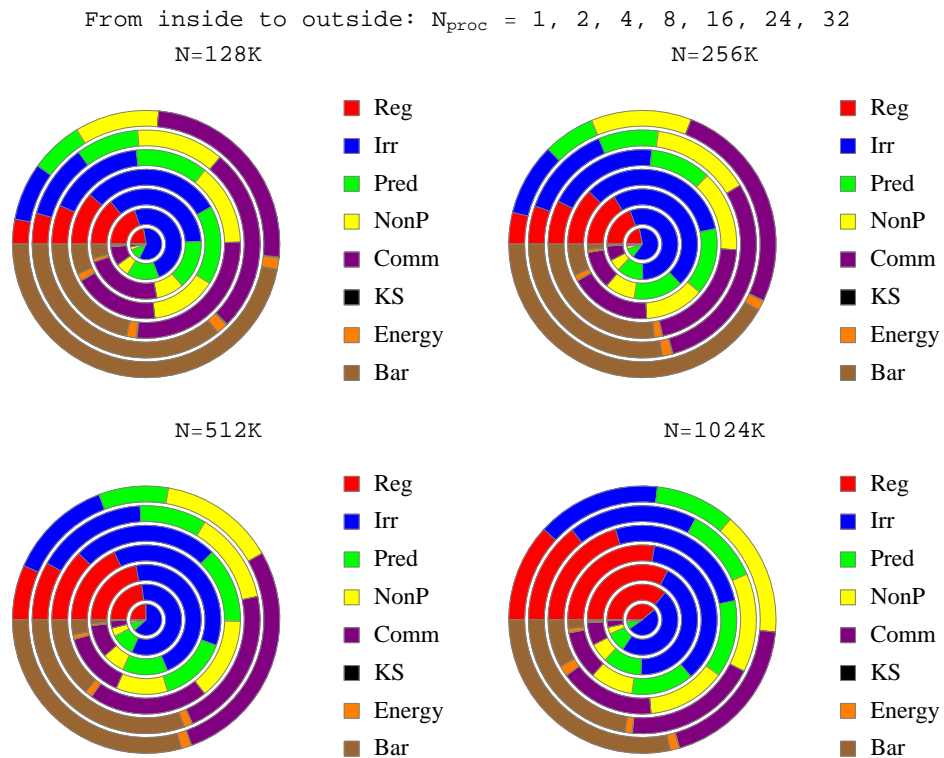
The formula (based on the experimental data fitting) quite well predict the code speed for different combination of particle and processor numbers (see the **Figure 4.** in the paper [2]). Based on this and using the extrapolation we predict our  $\varphi$ -GPU code performance for the much larger system (like the Oak Ridge Titan supercomputer system with  $\sim 20k$  GPU accelerator cards).

<sup>2</sup><http://gpgpu.org/>

<sup>3</sup><ftp://ftp.mao.kiev.ua/pub/berczik/phi-GRAPE/>



**Figure 3.** Total wall clock time of 1 time unit integration with the NBODY6++GPU on the K20 GPU cards. The lines with different symbols presents the different particle numbers.



**Figure 4.** Distribution of the total wall clock time of 1 time unit integration with the NBODY6++GPU on the K20 GPU cards. The different pie charts presents the code timing data for the different particle numbers.

In the **Figure 6.** we present this extrapolation exercise. We first simply extend the scale of the  $N=6M$  particle simulation for the larger GPU accelerator numbers up to 20k. As we can see even for this very high particle number the maximum performance which we can reach with the current GPU and network hardware is around  $\sim 350 - 400$  Tflop/s (see the line with  $N=6M$  label). We can have a significantly better speedup factor if we were to use the roughly 4 times faster network (see the line with LANx4 label). In this case we can reach above the 1 Pflop/s speed. If we could simultaneously speedup the network 10 times and also speedup the GPU hardware with 100 times we can reach sub Exaflop/s speed for  $N=300M$  particles. Such a very large particle number is needed to describe the full galactic bulge system, consisting of a few  $10^8 M_{\odot}$ , in star by star way.

The second set of the tests were done on the **laohu** cluster using the newly developed GPU enabled NBODY6++GPU code. This code is more complex including many binaries, stellar evolution and an Ahmad-Cohen neighbor scheme, and will be used for high resolution star cluster modeling. The results are shown in the plots **Figure 3.** and **Figure 4.** The new NBODY6++GPU code also shows a very good GPU utilization performance data and quite good network communication scaling. Right now, we are still working on the NBODY6++GPU code optimization, so, these results should be considered as "preliminary". After the next round of network communication optimization, we hope to get a speedup factor of  $\sim 2$  especially for larger particle and processor numbers.

As examples of the science applications done with these codes we will present new simulations for binary black holes to reach relativistic coalescence after galaxy mergers [11, 12, 13]. The simulations show unexpectedly short times to reach full coalescence of the black holes, compared to the typical galaxy merging time (less than a Gyr). In another paper we show how the dynamics of binary black holes before the final merger gives rise to significant gravitational wave emission in the pulsar timing band, due to their high eccentricities [14, 15].

## 4 Conclusions

Our massively parallel codes ( $\varphi$ -GPU and NBODY6++GPU), which use MPI parallelization as well as acceleration by many GPU's, scale well on large numbers of cores. They both run very well with no sign of saturation e.g. by communication on the new Kepler K20 GPU accelerator, reaching almost 1.5 Tflop/s per GPU with 2496 cores. These codes are currently being used for astrophysical research on galactic nuclei, requiring large particle resolution. With realistic technical improvements of GPU hardware and network speed we expect to reach approximately 0.2 Exaflop/s speed for  $N=300M$  particles.

## 5 Acknowledgments

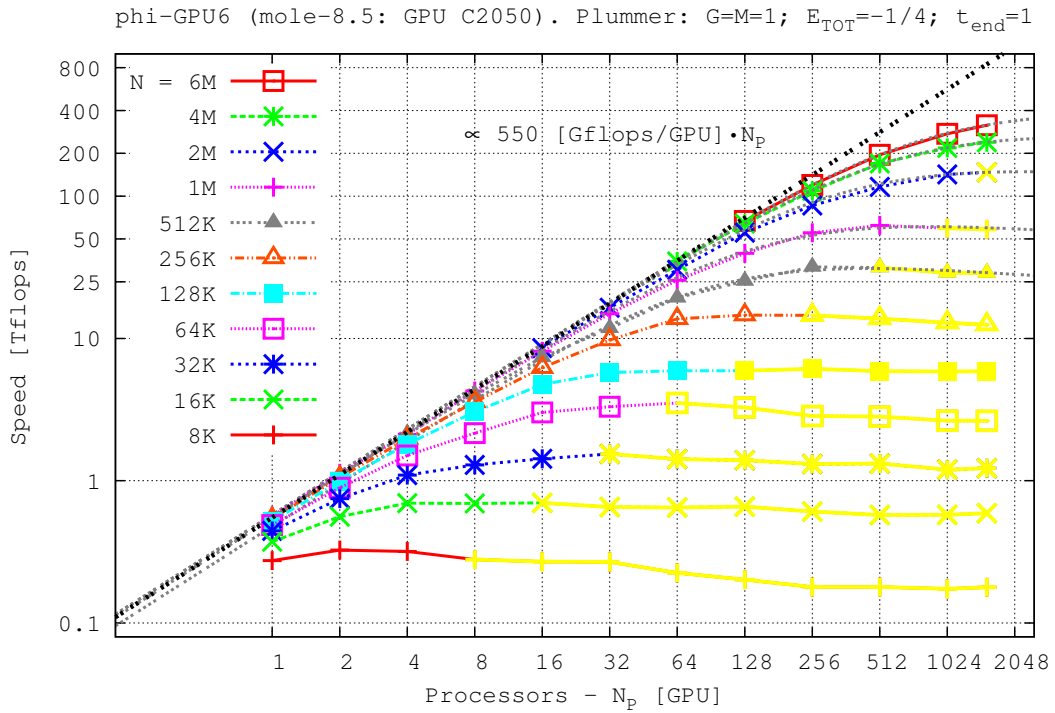
We acknowledge support by Chinese Academy of Sciences through the Silk Road Project at NAOC, through the Chinese Academy of Sciences Visiting Professorship for Senior International Scientists, Grant Number 2009S1 - 5 (RS), and through the "Qianren" special foreign experts program of China.

The special GPU accelerated supercomputer **laohu** at the Center of Information and Computing at National Astronomical Observatories, Chinese Academy of Sciences, funded by Ministry of Finance of People's Republic of China under the grant *ZDYZ2008 - 2*. We also used smaller GPU clusters **titan**, **hydra** and **kepler**, funded under the grants I/80041-043 and I/84678/84680 of the Volkswagen Foundation and grants 823.219-439/30 and /36 of the Ministry of Science, Research and the Arts of Baden-Württemberg, Germany. Some code development was also done on the Milky Way supercomputer, funded by the Deutsche Forschungsgemeinschaft (DFG) through Collaborative Research Center (SFB 881) "The Milky Way System" (subproject Z2), hosted and co-funded by the Jülich Supercomputing Center (JSC).

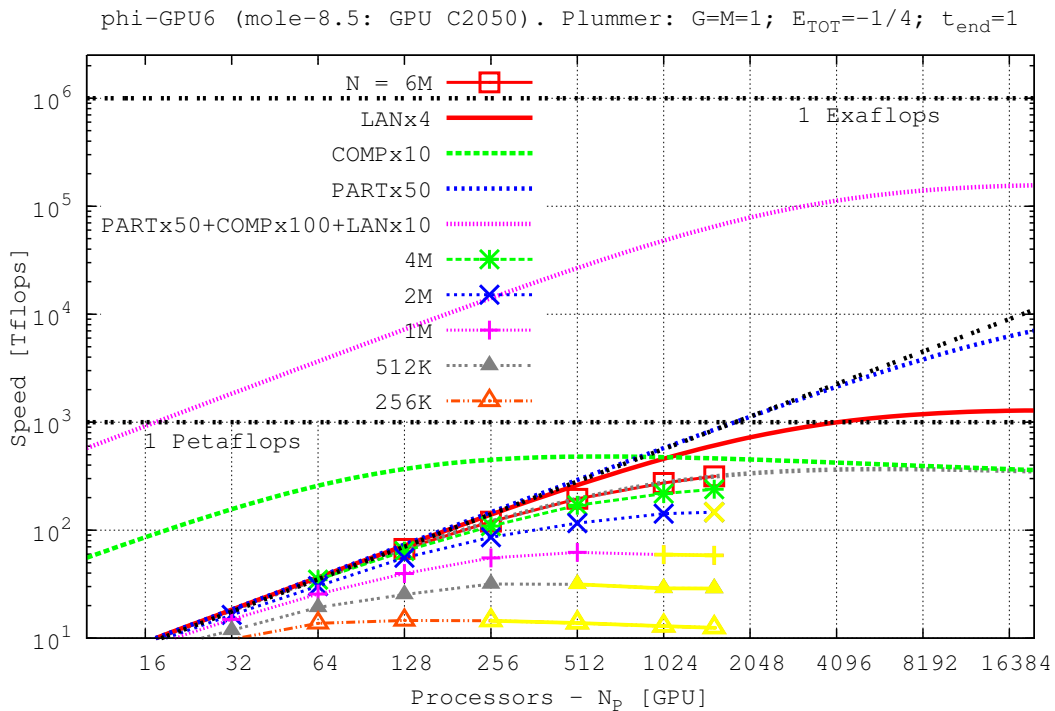
PB acknowledges the special support by the NASU under the Main Astronomical Observatory GRID/GPU computing cluster project.

## References

- [1] Spurzem R., Berczik P., Zhong S., Nitadori K., Hamada T., Berentzen I., Veles A.: Supermassive Black Hole Binaries in High Performance Massively Parallel Direct N-body Simulations on Large GPU Clusters. *Advances in Computational Astrophysics: Methods, Tools, and Outcome. Astronomical Society of the Pacific Conference Series*, Edited by R. Capuzzo-Dolcetta, M. Limongi, and A. Tornambe, 453: 223-233, 2012.



**Figure 5.** Speed performance with mixed (fp32 + fp64) precision of the  $\varphi$ -GPU 6<sup>th</sup> order scheme on the C2050 GPU cards. The lines with different symbols presents the different particle numbers.



**Figure 6.** Speed performance prognosis with mixed (fp32 + fp64) precision of the  $\varphi$ -GPU 6<sup>th</sup> order scheme on the C2050 GPU cards. The lines with different symbols presents the different particle numbers. The dotted horizontal lines shows the 1 Petaflops & 1 Exaflops levels.

- [2] Berczik P., Nitadori K., Zhong S., Spurzem R., Hamada T., Wang X., Berentzen I., Veles A., Ge W.: High performance massively parallel direct N-body simulations on large GPU clusters. *International conference on High Performance Computing*, Kyiv, Ukraine, October 8-10, 2011, p. 8-18.  
[http://www.hpc-ua.org/sites/default/files/proceedings-2011/1.1\(8\).pdf](http://www.hpc-ua.org/sites/default/files/proceedings-2011/1.1(8).pdf)
- [3] Harfst S., Gualandris A., Merritt D., Spurzem R., Portegies Zwart S., Berczik P.: Performance analysis of direct N-body algorithms on special-purpose supercomputers. *New Astronomy*, 12: 357-377, 2012.
- [4] Aarseth S.J.: From NBODY1 to NBODY6: The Growth of an Industry. *Publications of the Astronomical Society of the Pacific*, 111: 1333-1346, 1999.
- [5] Spurzem R.: Direct N-body Simulations. *Journal of Computational and Applied Mathematics*, 109: 407-432, 1999.
- [6] Nitadori K., Aarseth S.J.: Accelerating NBODY6 with graphics processing units. *Monthly Notices of the Royal Astronomical Society*, 424: 545-552, 2012.
- [7] Li S., Liu F.K., Berczik P., Chen X., Spurzem R.: Interaction of Recoiling Supermassive Black Holes with Stars in Galactic Nuclei. *The Astrophysical Journal*, 748: 65 (14 pp), 2012.
- [8] Fiestas J., Porth O., Berczik P., Spurzem R.: Evolution of growing black holes in axisymmetric galaxy cores. *Monthly Notices of the Royal Astronomical Society*, 419: 57-69, 2012.
- [9] Just A., Khan F.M., Berczik P., Ernst A., Spurzem R.: Dynamical friction of massive objects in galactic centres. *Monthly Notices of the Royal Astronomical Society*, 411: 653-674, 2011.
- [10] Merritt D., Berczik P., Laun F.: Brownian Motion of Black Holes in Dense Nuclei. *The Astronomical Journal*, 133: 553-563, 2007.
- [11] Khan F.M., Holley-Bockelmann K., Berczik P., Just A.: Supermassive Black Hole Binary Evolution in Axisymmetric Galaxies: The Final Parsec Problem is Not a Problem. *The Astrophysical Journal*, 773: 100 (6 pp), 2013.
- [12] Khan F.M., Berentzen I., Berczik P., Just A., Mayer L., Nitadori K., Callegari S.: Formation and Hardening of Supermassive Black Hole Binaries in Minor Mergers of Disk Galaxies. *The Astrophysical Journal*, 756: 30 (10 pp), 2012.
- [13] Khan F.M., Preto M., Berczik P., Berentzen I., Just A., Spurzem R.: Mergers of Unequal-mass Galaxies: Supermassive Black Hole Binary Evolution and Structure of Merger Remnants. *The Astrophysical Journal*, 749: 147 (14 pp), 2012.
- [14] Preto M., Berentzen I., Berczik P., Spurzem R.: Fast Coalescence of Massive Black Hole Binaries from Mergers of Galactic Nuclei: Implications for Low-frequency Gravitational-wave Astrophysics. *The Astrophysical Journal Letters*, 732: L26 (6 pp), 2011.
- [15] Berentzen I., Preto M., Berczik P., Merritt D., Spurzem R.: Binary Black Hole Merger in Galactic Nuclei: Post-Newtonian Simulations. *The Astrophysical Journal*, 695: 455-468, 2009.
- [16] Preto M., Berentzen I., Berczik P., Merritt D., Spurzem R.: Merger of massive black holes using N-body simulations with post-Newtonian corrections. *Journal of Physics: Conference Series*, 154: 012049 (6 pp), 2009.
- [17] Berentzen I., Preto M., Berczik P., Merritt D., Spurzem R.: Post-Newtonian simulations of super-massive black hole binaries in galactic nuclei. *Astronomische Nachrichten*, 329: 904-907, 2008.
- [18] Berczik P., Merritt D., Spurzem R., Bischof H.-P.: Efficient Merger of Binary Supermassive Black Holes in Nonaxisymmetric Galaxies. *The Astrophysical Journal*, 642: L21-L24, 2006.
- [19] Berczik P., Merritt D., Spurzem R.: Long-Term Evolution of Massive Black Hole Binaries. II. Binary Evolution in Low-Density Galaxies. *The Astrophysical Journal*, 633: 680-687, 2005.
- [20] Just A., Yurin D., Makukov M., Berczik P., Omarov C., Spurzem R., Vilkoviskij E.Y.: Enhanced Accretion Rates of Stars on Supermassive Black Holes by Star-Disk Interactions in Galactic Nuclei. *The Astrophysical Journal*, 758: 51 (10 pp), 2012.
- [21] Ernst A., Just A., Berczik P., Olczak C.: Simulations of the Hyades. *Astronomy & Astrophysics*, 536: A64 (12 pp), 2011.

- [22] Ernst A. Just A. Berczik P. Petrov M.I.: Calibration of radii and masses of open clusters with a simulation. *Astronomy & Astrophysics*, 524: A62 (9 pp), 2010.
- [23] Kharchenko N.V., Berczik P., Petrov M.I., Piskunov A.E., Röser S., Schilbach E., Scholz R.-D.: Shape parameters of Galactic open clusters. *Astronomy & Astrophysics*, 495: 807-818, 2009.
- [24] Just A., Berczik P., Petrov M.I., Ernst A.: Quantitative analysis of clumps in the tidal tails of star clusters. *Monthly Notices of the Royal Astronomical Society*, 392: 969-981, 2009.