Black Holes in Galactic Nuclei simulated with large GPU clusters in CAS

Rainer Spurzem, Peter Berczik
National Astronomical Observatories of China, Chinese Academy of Sciences (NAOC/CAS)
20A Datun China Rd., Chaoyang District, Beijing 100012,

Supermassive Black Holes in Galactic Nuclei

Many, if not all galaxies harbour supermassive black holes. If galaxies merge, which is quite common in the process of hierarchical structure formation in the universe, their black holes sink to the centre of the merger remnant and form a tight binary. Depending on initial conditions and time supermassive black hole binaries are prominent gravitational wave sources, if they ultimately come close together and coalesce. 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 (Kupi et al. 2006, Berentzen et al. 2009).

Kupi, G., Amaro-Seoane, P., Spurzem, R. , Dynamics of compact object clusters: a post- Newtonian study, 2006, Mon. Not. Royal Astron. Soc. 371, 45.
Berentzen, I., Preto, M., Berczik, P., Spurzem, R., Binary Black Hole Merger in Galactic Nuclei: Post-Newtonian Simulations, 2009, The Astrophysical Journal 695, 455

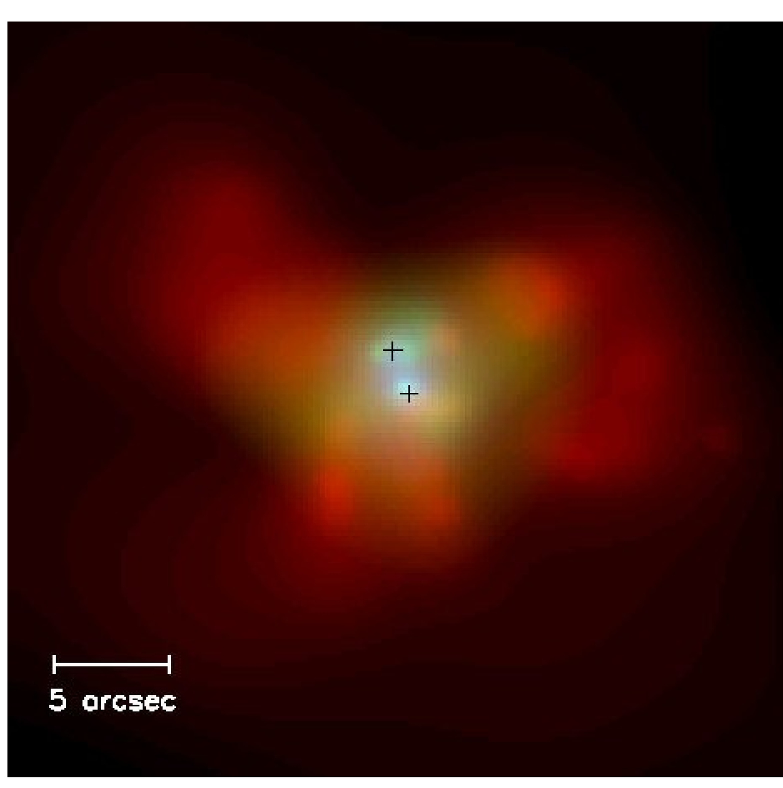

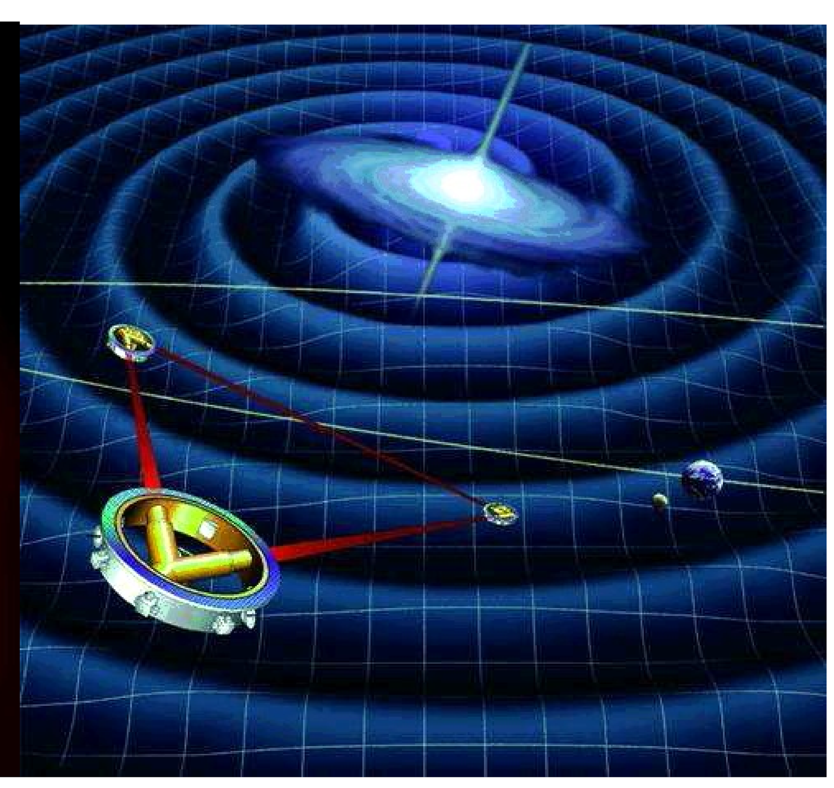

Figures on Top:

Left: X-Ray observation of two nuclei of a galaxy after merging, both bright spots are interpreted as hot gas around a supermassive black hole, which cannot be directly resolved. The distance between both black holes is 3000 light years (Picture from S. Komossa, Max-Planck-Institute for Extraterrestrial Physics in Garching near Munich, Germany).

Right: Artist's Impression of the space-based LISA laser interferometer satellites, designed to detect gravitational waves from massive black hole coalescences in the entire universe (Picture: ESA). Our team is linked to the German LISA consortium for the research on astrophysical gravity wave generation by binary black holes in galactic nuclei.

Gravitational Wave from Black Holes in Galactic Centres

-

Supermassive black hole binaries (SMBHBs) are products of galaxy mergers, and are important sources of gravitational wave (GW) radiation For this reason, it will be important to detect electromagnetic counterparts of GW radiation sources. In this work we focus on stars being disrupted by tidal forces in the vicinity of supermassive black holes, they generate an X-ray flare, which is observable. We are also interested in supermassive black holes which are kicked out of galactic nuclei due to recoil after a gravitational radiation induced merger.

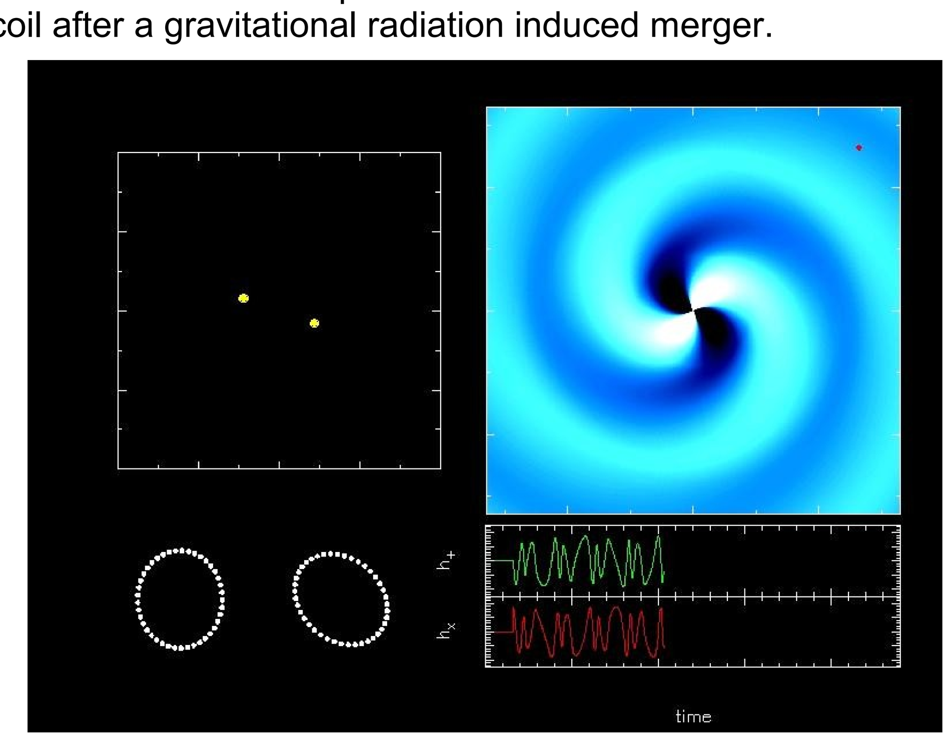

Snapshot from Animation of gravitational wave emission due to one of our simulated binary black holes. Animation picture thanks to Ingo Berentzen, Univ. of Heidelberg.

Top left: position of black holes
Top right: gravitational wave pattern in coordinate plane
Bottom left: polarisations of gravitational waves
Bottom right: amplitudes of the two polarisations as a function of time at hypothetic earth observer's point (black dot in upper right picture).

Other teams, interested students and postdocs from NAOC, or from other Chinese institution are invited to test their own, develop new and use their codes on the NAOC/CAS GPU supercomputer. International Collaborations are also possible. Please contact spurzem@bao.ac.cn and look at the Silk Road Project Webpage http://silk0.bao.ac.cn/silkroad

Software φGPU

The application code which we use for benchmarking here is a direct N-body simulation code for astrophysics, using a high order Hermite integration scheme and hierarchical block time steps (the code supports time integration of particle orbits with 4 , 6 , and 8 order schemes). The code is called φGPU, it has been developed from our earlier published versions φGRAPE (using GRAPE hardware instead of GPU, Harfst et al. 2007). It is parallelised using MPI (Message Passing Interface), and on each node using many cores of the special GPU hardware to compute gravitational forces between particles. The code was developed in cooperation with Keigo Nitadori (RIKEN Japan) and Tsuyoshi Hamada (Nagasaki Univ. Japan). It is written in C++ with CUDA, and it is based on Nitadori & Makino (2008) earlier CPU serial code (yebisu).

MPI parallelization was done in the so-called j-particle parallelization mode as in the earlier φGRAPE code (Harfst et al. 2007). The particles are divided equally between the working nodes (GPU memory) and in each node we calculate only the fractional forces for the active i-particles at the current timestep. Due to the hierarchical block time step scheme the number Nact of active particles (due for a new force computation at a given time level) is usually small compared to the total particle number N, but its actual value can vary from 1... N. This distinguishes the code from a simple shared time step code, where at every time step all particles are integrated. The full forces from all the particles acting on the active particles we get after using the global MPI_SUM communication routines.

We use native GPU support and direct code access to the GPU with only CUDA. Recently we use CUDA 2.2. Multi GPU support is achieved through MPI parallelization; each MPI process uses only a single GPU, but we can start two MPI processes per node (to use effectively the dual CPU's and GPU's in the NAOC cluster) and in this case each MPI process uses its own GPU inside the node. Communication always (even for the processes inside one node) works via MPI. We do not use any of the possible OMP (multi-thread) features of recent gcc 4.x compilers inside one node.

References:
Harfst, S., Gualandris, A., Merritt, D., Spurzem, R., Portegies Zwart, S., Berczik, P., Performance analysis of direct N-body algorithms on special-purpose supercomputers, 2007, New Astronomy 12, 357
Nitadori, K., Makino, J., Sixth- and eighth-order Hermite integrator for N-body simulations, 2008, New Astronomy 13, 498

Performance on GPU Supercomputer at NAOC

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. The NAOC cluster is running with 85 nodes, each with two NVIDIA Tesla C1070 GPUs.

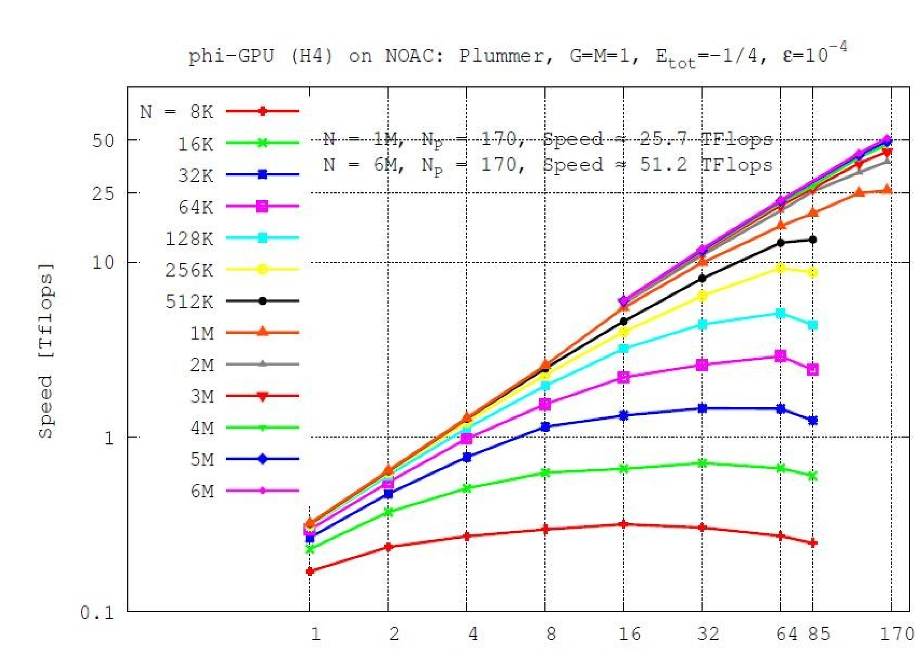

Figure 2 on Top: Sustained Speed in Teraflop/s of astrophysical N-body Simulation for different particle numbers, as a function of number of GPU multiprocessors used. Near ideal speed-up is reached for large N, approaching the diagonal.

Figure 1 on Top: NAOC GPU Cluster.

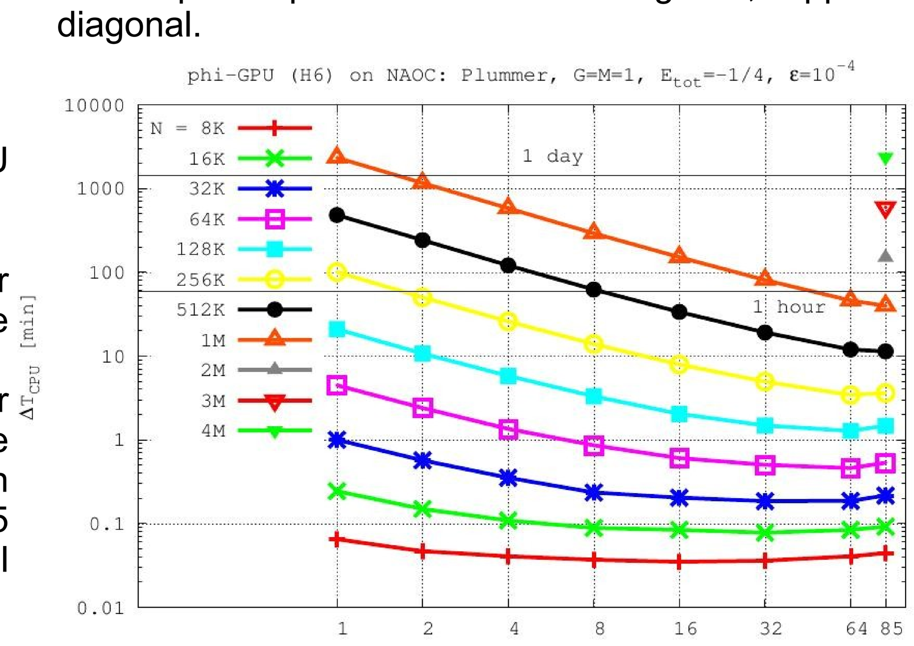

Figure 3 Right: Strong Speed-Up for Varying particle number, for the same
N-body simulation as in Figure 2. For too small particle number the simulation gets communication limited. At one million bodies and 85 GPU multi-processors we are still computationally dominated.