Physics Simulations on the DESY HPC platform

I.Agapov, G.Geloni (European XFEL GmbH)
J.Grebenyuk, Y.Kemp, A.M.de la Ossa, J.Osterhoff, B.Lewendel, F.Schlünzen, S.Sternberger (DESY Hamburg)

High performance or massively parallel computing is gaining momentum at DESY, driven by the recent accelerator and detector developments, and new experiments like single particle imaging or crystallography at Free Electron Laser facilities like FLASH, LCLS or - in the near future - at the European XFEL. A small-to-medium scale HPC platform was launched in 2011 at DESY Hamburg and is now being expanded and optimized for specific simulation and data processing tasks. The main usage so far has been for studies of novel particle acceleration techniques and for simulations required for preparation for the startup of the European XFEL facility. Later on the platform should also be able to handle challenging high-rate data processing tasks. To this end, specific applications are currently being bechmarked, to test performance and stability of cluster/parallel filesystems. Here we report on the hardware platform status, experience with physics simulations, and software developments to facilitate optimal usage of the cluster.

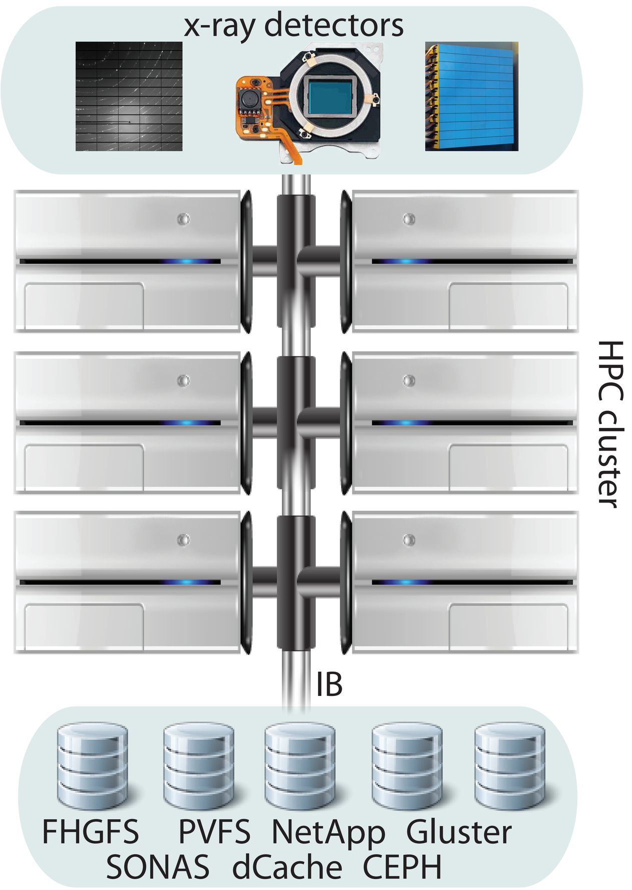

Hardware platform

The DESY HPC Cluster currently consists of 64 HP ProLiant compute blades, each node equipped with 2x Intel Xeon CPU E5345 @ 2.33 GHz and 16 GByte RAM, totalling 512 cores and 1.024 GByte memory. A Mellanox MT25418 DDR InfiniBand is utilized for inter-node communication, each compute blade directly connected to the central Voltaire ISR 2012 DDR InfiniBand switch. The cluster will later on be scaled up with 16 nodes each equipped with 64-cores using AMD Interlagos technology and 192 GByte of RAM, totalling 1.024 cores and 3.072 GByte of memory. All nodes will be interconnected with QDR InfiniBand. Physics simulations such as plasma-wakefield interaction studies or free-electron laser studies pose relaxed requirements on the file system performance. Processing experimental data however presents a challenge. Current x-ray detectors operate at 10-100Hz generating data streams of up to 500GByte/s. After initial quality assessment locally buffered data need to be transferred to fast, high capacity storage for full analysis. Each of the roughly 50 experimental stations at DESY produces similar data streams, which need to be processed in parallel. Various files systems are being investigated for suitability to cope with such data streams, and parallel processing using MPI-I/O and pHDF5. Filesystems (and protocols) under testing are for example FhGFS, Sonas, NetApp, CEPH or PVFS. In addition dCache (NFS4.1) will be added as a fast, scalable storage and archival backend. Plots on the right show initial meta-data and detector performance tests.

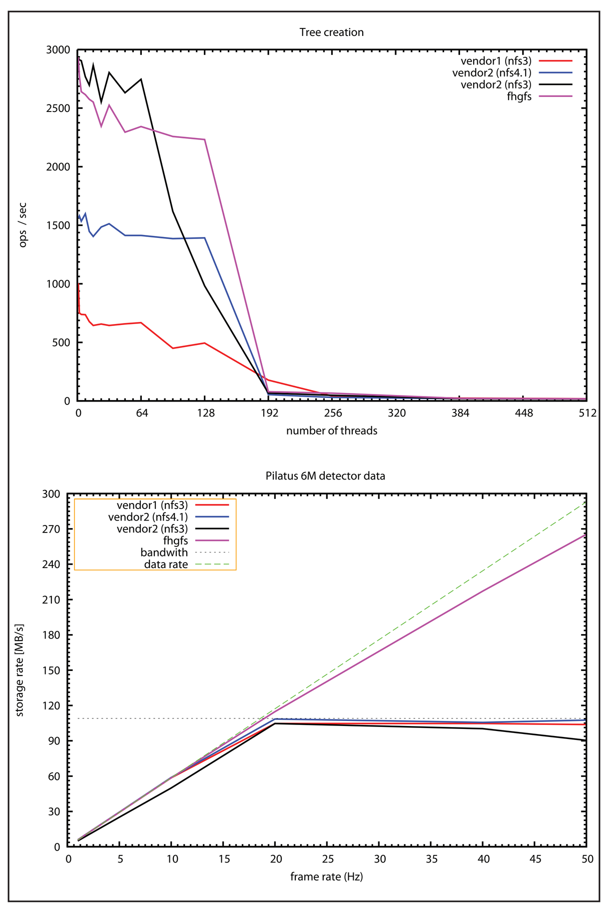

Simulations of plasma-wakefield acceleration

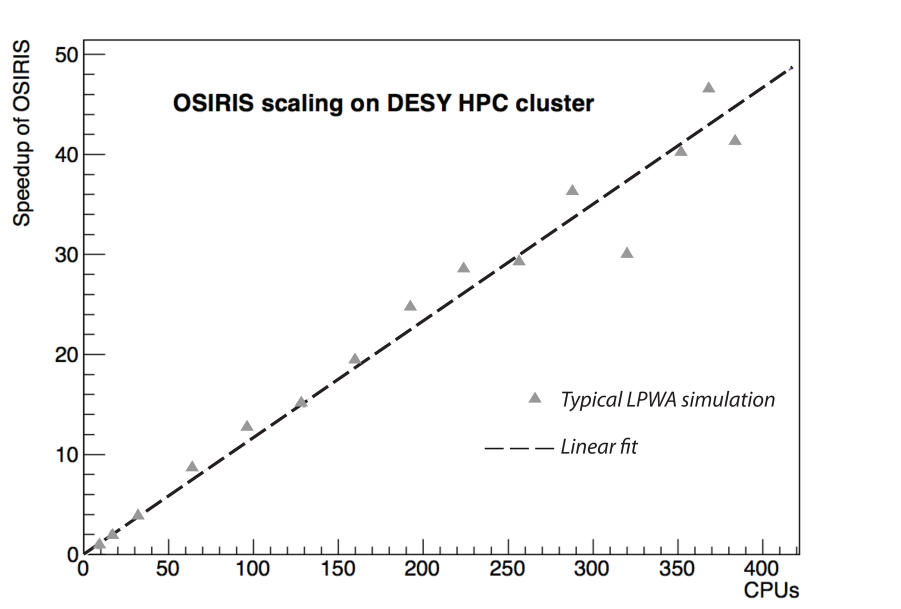

Figure 2. Scalability of OSIRIS parallel efficiencies for a typical LWFA simulation on the DESY HPC cluster

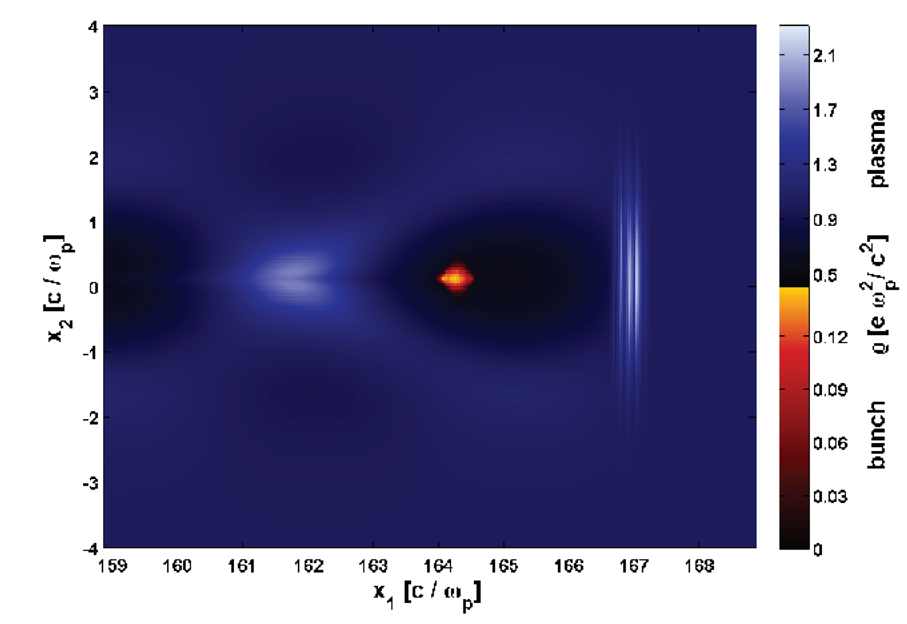

Figure 3. Electron bunch externally injected into a laser-induced plasma wake, simulated with the PIC code OSIRIS.

Laser-wakefield acceleration (LWFA) has experienced growing scientific interest and fast development during the last decade. Short and highly intense laser pulses focused into a gas target ionise the gas and may excite large amplitude, co-propagating plasma waves that support extreme electric fields (>10 GV/m) which can be used to accelerate charged particles on very short distances. One of the key issues of the LWFA is control over the process of self-injection of electrons into the accelerating phase of the wake. Our studies concentrate on controlled external injection of pre-accelerated tailored bunches from a conventional accelerator, which open numerous opportunities for more fundamental understanding of the underlying processes, such as probing the wakefield and studying the bunch phase-space evolution.

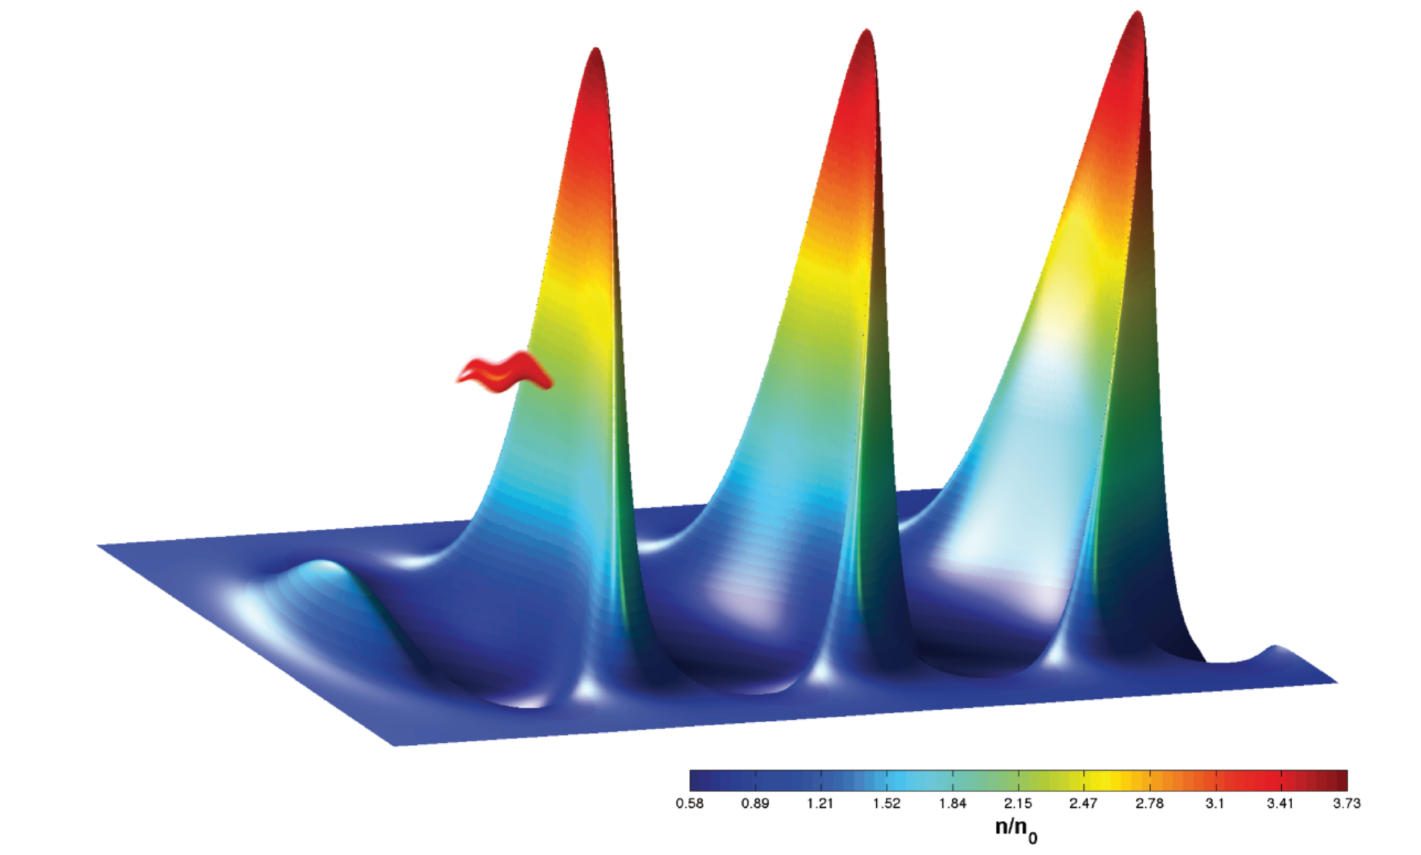

Figure 1. Schematic representation of a wake in electron density generated by high-intensity laser, with an externally injected electron bunch. High field gradients in the wake can be used for electron acceleration on short distances

A deeper theoretical insight into laser-plasma interaction is gained by particle-in-cell (PIC) simulations. The main PIC iteration loop consists of four steps: particle current deposition on a grid, Maxwell equations solving to obtain the fields, interpolation of the fields back to particle positions, and finally particle push according to the Lorentz force. PIC code OSIRIS, developed at Instituto Superior Tecnico (IST, Lisbon, Pt) and University of California (UCLA, LA, USA), is a fully explicit, fully relativistic, object-oriented code which takes full advantage of high performance computing. OSIRIS incorporates a dynamic loadbalancing algorithm that adjusts the partition of the simulation dynamically as the simulation progresses, ensuring optimal computational load per core. Scaling tests of OSIRIS at the HPC cluster at DESY showed high parallel scalability efficiencies (See Figure 1). OSIRIS LWFA simulations were performed to prepare external injection experiments at DESY facilities (see Figure 2).

Radiation simulations for European XFEL

European XFEL is the world-wide leading facility for research with photons currently under construction in the Hamburg area. Design, commissioning and operation of x-ray beamlines requires extensive numerical studies of FEL as well as spontaneous synchrotron radiation properties. FEL simulations are performed by means of the PIC code GENESIS based on the MPI architecture. Spontaneous radiation simulations are based on Green's function approach which results in evaluating certain quickly oscillating integrals. Taking the effects of emittance, energy spread and quantum diffusion into account leads to evaluating such integrals for a multidimensional parameter space which presents a numerical challenge. Using a modified MPI-based version of the code SRW on the HPC cluster allows to perform most of such calculation nearly real time.The numerical studies in questions often require running a certain chain of simulation codes and post-processing steps repeatedly. To facilitate this a generic software framework is being developed which features common I/O libraries, a workflow engine, and various common post-processing routines.

This framework allows interoperability of various simulation codes (by means of common graphical user interfaces, common input and output files), transparent open data formats (hdf5-based for simulation output, xml-based for geometry and problem description), and exposing simple common interfaces to high-performance computing paradigms (Cluster, multi-core, GPU) so that the users are not aware of the technical details (such as command syntax, local directory structures etc.) required to use specific hardware platforms.

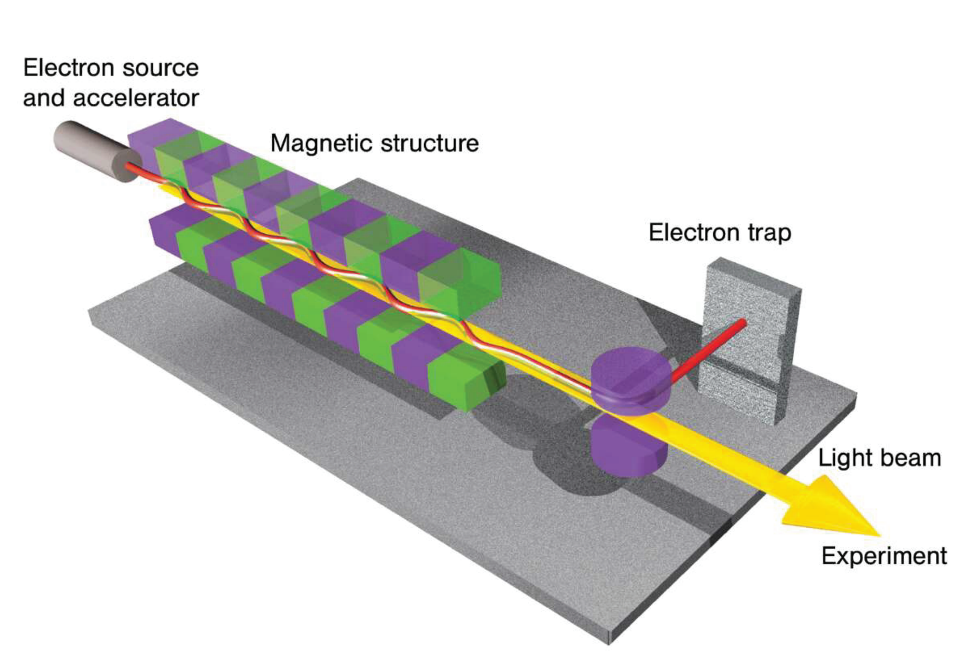

Figure 4. Sketch of a Free Electron Laser layout

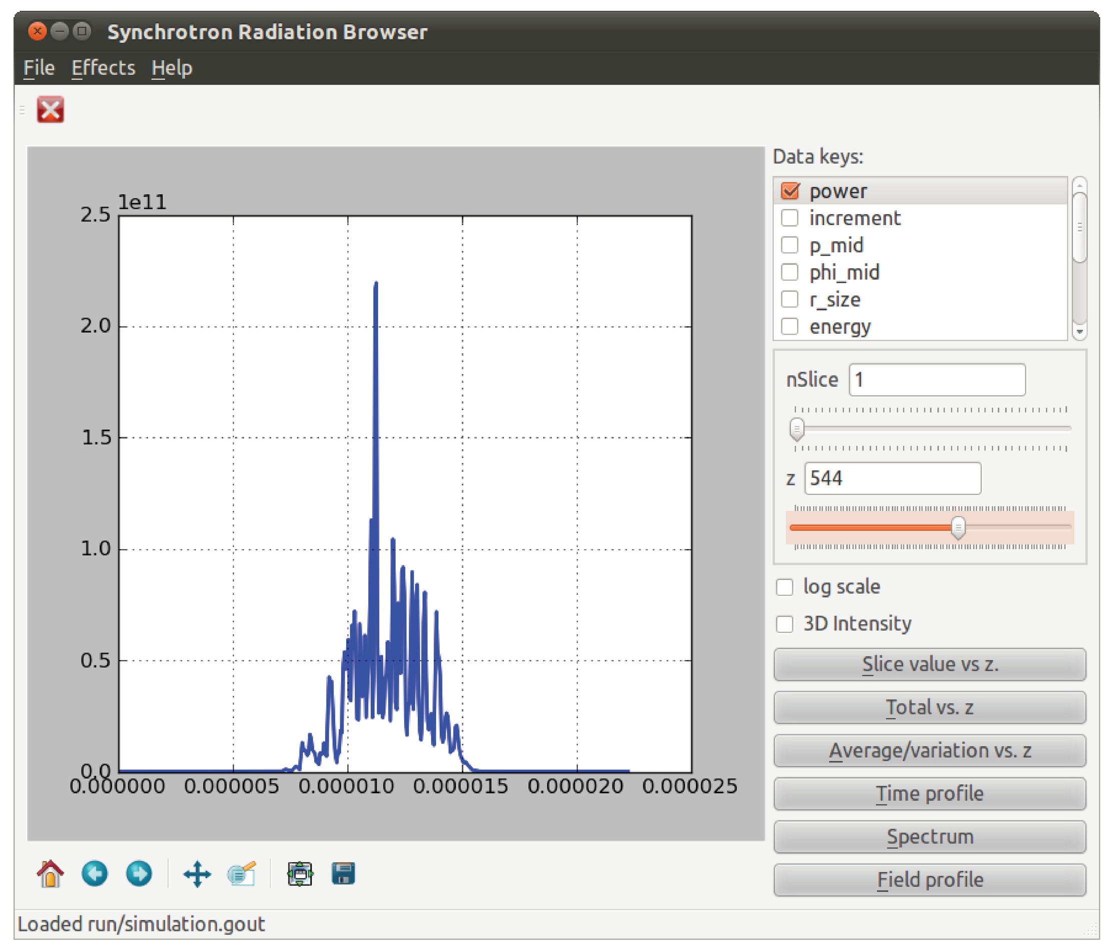

Figure 6. Simulated FEL radiation time profile for European XFEL undulators

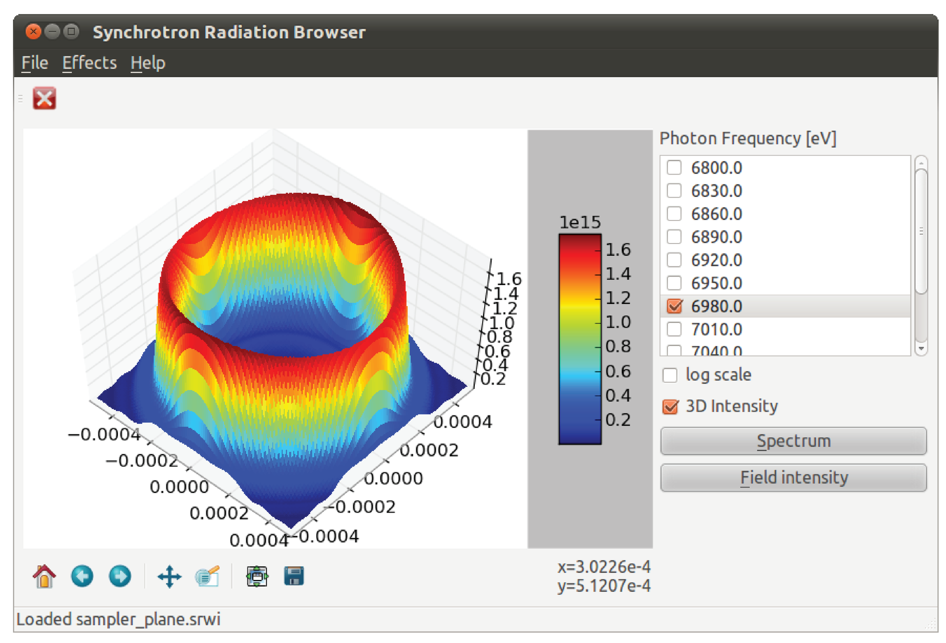

Figure 5. Simulated spontaneous radiation intensity profile for European XFEL undulators

This work is partially supported by the European Commission under the 7th Framework Programme Grant Agreement 283745 (CRISP) and 283556 (PaNdata ODI)