World Scientific
www.worldscientific.com

SixTrack V and runtime environment

R. De Maria*

*Beam Department (BE-ABP-HSS),
CERN, 1211, Geneva 23, Switzerland
riccardo.de.maria@cern.ch*

J. Andersson, V. K. Berglyd Olsen, L. Field,
M. Giovannozzi, P. D. Hermes, N. Høimyr, S. Kostoglou,
G. Iadarola, E. Mcintosh, A. Mereghetti, J. Molson,
D. Pellegrini, T. Persson and M. Schwinzerl

CERN, 1211, Geneva 23, Switzerland

E. H. Maclean

CERN, 1211, Geneva 23, Switzerland

University of Malta, Msida, MSD 2080, Malta

K. N. Sjobak

CERN, 1211, Geneva 23, Switzerland

University of Oslo, Boks 1072 Blindern, 0316, Oslo, Norway

I. Zacharov

EPFL, Rte de la Sorge, 1015, Lausanne, Switzerland

S. Singh

*Indian Institute of Technology Madras,
IIT P.O., Chennai 600 036, India*

SixTrack is a single-particle tracking code for high-energy circular accelerators routinely used at CERN for the Large Hadron Collider (LHC), its luminosity upgrade (HL-LHC), the Future Circular Collider (FCC) and the Super Proton Synchrotron (SPS) simulations. The code is based on a 6D symplectic tracking engine, which is optimized for long-term tracking simulations and delivers fully reproducible results on several platforms. It also includes multiple scattering engines for beam–matter interaction studies,

*Corresponding author.

as well as facilities to run the integrated simulations with external particle matter inter-action codes. These features differentiate SixTrack from general-purpose, optics-design software. The code recently underwent a major restructuring to merge the advanced features into a single branch, such as multiple ion species, interface with external codes and high-performance input/output. This restructuring also removed a large number of compilation flags, instead enabling/disabling the functionality with runtime options. In the process, the code was moved from Fortran 77 to Fortran 2018 standard, also allowing and achieving a better modularization. Physics models (beam–beam effects, Radio-Frequency (RF) multipoles, current carrying wires, solenoid and electron lenses) and methods (symplecticity check) have also been reviewed and refined to offer more accurate results. The SixDesk runtime environment allows the user to manage the large batches of simulations required for accurate predictions of the dynamic aperture. SixDesk supports several cluster environments available at CERN as well as submitting jobs to the LHC@Home volunteering computing project, which enables volunteers contributing with their hardware to CERN simulation. SixTrackLib is a new library aimed at provid-ing a portable and flexible tracking engine for single- and multi-particle problems using the models and formalism of SixTrack. The library is able to run in CPUs as well as graphical processing units (GPUs). This contribution presents the status of the code, summarizes the main existing features and provides details about the main development lines SixTrack, SixDesk and SixTrackLib.

Keywords: Beam dynamics; interaction particles matter.

1. Introduction

SixTrack[1] (available in http://cern.ch/sixtrack) is a 6D single-particle symplectic tracking code able to compute the trajectories of individual relativistic charged particles in circular accelerators for studying dynamic aperture (DA) or evaluating the performance of beam-intercepting devices like collimators.[2,3] It can compute linear and nonlinear optics functions, time-dependent effects and extract indicators of chaos from tracking data. SixTrack implements scattering routines and aperture calculations to compute "loss maps" i.e. leakage from collimators as a function of longitudinal position along the ring and collimation efficiency.[2,3]

Different from a general-purpose code like MAD-X,[4] SixTrack is optimized for speed and numerical reproducibility. It can be also linked with the BOINC library[5] to use the volunteering computing project LHC@Home.[6] SixTrack studies, such as estimation of DA of large storage rings like the Large Hadron Collider (LHC) or the Future Circular Collider (FCC), require massive computing resources, since they consist of scans over large parameter spaces for probing nonlinear beam dynamics over long periods.

The SixDesk runtime environment manages SixTrack simulations from input generation, job queue management (using HTCondor[a], LSF[b] or BOINC), to collecting and post-processing results.

SixTrackLib is a new library built from scratch in C with the main aim of offering a portable tracking engine for other codes and offloading SixTrack simulation to graphical processing units (GPUs).

---

[a]High Throughput Condor (HTCondor), https://research.cs.wisc.edu/htcondor.
[b]IBM spectrum LSF.

2. Main Features of SixTrack

SixTrack tracks an ensemble of particles defined by a set of coordinates through several beam-line elements, using symplectic maps,[7–9] or scattering elements.

2.1. *Coordinates*

The set of coordinates is larger than the minimum needed to describe the motion.

Additional variables (see Table 1) are used to store the energy-related quantities used in the tracking maps that are updated only on energy changes, which do not occur very frequently in synchrotrons in the absence of radiation effects, to save computational time. Maps for dipoles and quadrupoles with nonzero length also reuse the energy-dependent factors of the first- and second-order polynomials of the map that are recalculated at each energy change. Furthermore, different ion species, such as debris from interaction with matter, can be also tracked at the same time[10] by correctly taking into account their relativistic speed as well as transverse magnetic rigidity. Variables used internally in tracking are not canonical, however, once they are converted to canonical form, the maps are symplectic. Differently from other codes such as MAD-X or PTC, SixTrack uses

$$\sigma = s - \beta_0 ct$$

as the longitudinal coordinate during tracking to avoid rounding errors associated to the relativistic β when updating σ in drifts[c].

Table 1. Coordinates of the particles tracked by SixTrack. Extra-dependent coordinates are stored to save expensive computations. Variables m_0, β_0, P_0 refer to the reference particle. SixTrack maps are symplectic when re-expressed in canonical coordinates.

x	Horizontal position
$x_p = \frac{P_x}{P}$	Horizontal direction cosine
y	Vertical position
$y_p = \frac{P_y}{P}$	Vertical direction cosine
$\sigma = s - \beta_0 ct$	Negative time delay with respect to the reference particle
$\delta = \frac{P\frac{m_0}{m} - P_0}{P_0}$	Relative momentum deviation with respect to the reference particle
E	Energy of the particle
P	Momentum of the particle
m	Mass of the particle
$\chi = \frac{q}{q_0}\frac{m_0}{m}$	Relative mass to charge ratio with respect to the reference particle
$r_v = \frac{\beta}{\beta_0}$	Velocity ratio with respect to the reference particle
$r_1 = \frac{1}{1+\delta}$	Quantity used in tracking maps
$r_2 = \frac{\chi}{1+\delta}$	Quantity used in tracking maps
$r_3 = \frac{1-\chi}{1+\delta}$	Quantity used in tracking maps
$r_4 = \frac{\delta}{1+\delta}$	Quantity used in tracking maps

[c]$\Delta\sigma/l = 1 - \frac{\beta_0}{\beta}(1 + \frac{x_p^2+y_p^2}{2})$ compared to $\Delta ct/l = \frac{1}{\beta_0} - \frac{1}{\beta}(1 + \frac{x_p^2+y_p^2}{2})$ in MAD-X or PTC.

Table 2. Beam-line elements implemented in SixTrack.

Expanded drift	Drift resulting from an approximated Hamiltonian to second order
Exact drift[11]	Drift using the exact equation of motion
Single thin multipole	Infinitely thin magnetic multipole
Thin multiple block	Combination of magnetic multipole components up to the 20th order
Thick dipole–quadrupole	Combined magnetic dipole and quadrupole of a finite length
Thin solenoid	Infinitely thin solenoid
Accelerating cavity	Infinitely thin accelerating cavity
RF-multipoles[12]	Multipolar kick with longitudinal position sinusoidal dependence
4D Beam–beam	Transverse-only beam–beam interaction
6D beam–beam[13]	Full 6D beam–beam interaction with crossing angle, hourglass effects
Wire[14]	Infinitely thin wire carrying current
Hollow electron lens[15, 16]	Interaction with hollow electron distribution

2.2. *Beam-line elements*

Table 2 shows the different types of beam-line elements implemented in Six-Track. Thin multipoles are used in conjunction with the MAKETHIN and SIXTRACK commands in MAD-X to implement symplectic integrators for elements of finite length. Thin multipoles include the effect of the curvature, when present, up to the second order. The tracking maps have been recently reviewed and benchmarked against MAD-X and its optics module for consistency[d].

2.3. *Scattering*

SixTrack is capable of simulating the basic scattering processes undergone by an ultra-relativistic proton in the multi-TeV range when passing through matter.[2, 3] The simulated processes range from ionization energy loss and multiple Coulomb scattering to point-wise interactions like Coulomb, elastic and inelastic events, including single diffractive scattering. Compound materials of interest for the low-impedance upgrade of the LHC collimators are implemented via averaged nuclear and atomic properties.[17] Other scattering models can be imported and made available in the SixTrack executable, such as that of Merlin[18,e] and Geant4.[19, 20]

A new module (SCATTER) is under development to offer a general framework for simulating scattering events in SixTrack. Currently, it supports beam scattering against a target specified as an area density distribution at a thin marker inserted into the lattice. Internally, the scattering module supports elastic scattering through Monte Carlo sampling of experimental data from the *TOTal cross-section, Elastic scattering and diffraction dissociation at the LHC Measurement* (TOTEM[21]). Alternatively, scattering events can be generated on the fly by Pythia8,[22] in which case elastic and diffractive processes are supported.

---

[d]J. Andersson, "Comparison of tracking codes.", in CERN openlab summer students lighting talks 2018, http://cds.cern.ch/record/2634289.
[e]Available at https://github.com/Merlin-Collaboration/Merlin.

2.4. *Optics calculations*

SixTrack contains matrix code for 5D optics calculation and a 6D tracking engine using Truncated Power Series Algebra (TPSA[23–25]) library for 6D optics calculation. The 6D tracking engine uses the following canonical variables:

$$\left(x,\ p_x = \frac{P_x}{P_0}\right), \quad \left(y,\ p_y = \frac{P_y}{P_0}\right), \quad \left(\zeta = \frac{\beta}{\beta_0}\sigma,\ \delta = \frac{P - P_0}{P_0}\right)$$

as conjugate canonical variables in 6D optics calculations which use explicitly symplectic maps. A symplecticity check is routinely performed on the one-turn map.

Coupled Twiss parameters (using the Mais–Ripken formalism[26]) can be extracted along the lattice. The optics parameters are optionally used to define strong beam sizes in beam–beam elements. The 6D optics module has been recently improved by removing some unnecessary ultra-relativistic approximations which introduced small symplectic errors.

2.5. *Dynamic effects*

A general functionality for dynamically changing simulation settings on a turn-by-turn basis has been implemented.[27] This allows setting magnet strengths including multipoles, radio-frequency (RF) amplitude and phase, reference energy and beam–beam element as a function of turn number. This can be useful for a number of different studies, e.g. magnet snap-back in the LHC,[28] HL-LHC crab cavity failure scenarios,[29–31] studies of beam losses during energy ramp[32] and hollow electron-lens modulation.[16] The settings can be internally computed as a function of turn number, or loaded from a file. These functions are specified using a flexible language that allows combining functions to achieve the required effect. The architecture of the functionality makes it easy to add support for new elements or new functions.

2.6. *Online aperture checking*

SixTrack is capable of performing a check during tracking (online) of whether tracked particles still fall within the mechanical acceptance of the machine or not. The check is activated by the presence of aperture markers in the accelerator lattice. All aperture types of relevance for LHC, HL-LHC and FCC can be handled, namely: RECTELLIPSE (i.e. the aperture identified by the intersection of a rectangle with an ellipse), RACETRACK (i.e. the aperture identified by a rectangle with quadrants of ellipses at the corners) and OCTAGON (i.e. an octagon with sides at multiples of 45°). The code is capable of handling horizontal and vertical offsets of the aperture markers, as well as rotations about the longitudinal axis. The implementation comes with a smooth, linear interpolation between all supported aperture markers, such that the aperture is very well defined at any location along the ring.

A smooth linear interpolation of the aperture model of the machine is essential to the detailed identification of the actual loss point tracked particles hitting the mechanical aperture of the machine (back-tracking algorithm). The loss location is

Fig. 1.   Longitudinal pattern (left frame) and transverse position (right frame) of losses in the inner triplet upstream of the IR5 experimental area for HL-LHC.[33] Losses get concentrated on the upper-left quadrant (right frame) because of the local position of the closed orbit. Simulation parameters: optics v1.3, $\beta^* = 15$ cm, full nonlinear model of the HL-LHC (no beam–beam collisions), no collimation system.

identified comparing particle trajectories with the aperture model of the machine by means of a bi-section method; the user can set the precision of the identification of the longitudinal position. For the time being, the feature is available only for thin-lens tracking, where back-tracking is comfortably performed along a drift; extension to all other thick-lens linear elements is in plan. As an example, Fig. 1 shows the longitudinal pattern and transverse position of losses in the inner triplet upstream of one experimental area for HL-LHC;[33] losses coming from DA in the absence of collimation system.

2.7. *Electron lenses*

Electron lenses can be deployed to deplete over-populated beam tails in a controlled manner; as such, they are currently being investigated in the framework of the HL-LHC project.[33] The module is in under development. Recent changes include: porting to ion tracking, relevant for benchmarking recent measurements in RHIC[34] at Brookhaven National Laboratory; possibility to deploy measured radial profiles of electron beams; compatibility to the BOINC platform for volunteering computing, relevant for evaluating long-term effects on DA. Ongoing development is focused on implementing 2D maps of measured electron profiles by means of Chebyshev polynomials.[35]

2.8. *Post-processing*

Long-term tracking with SixTrack is used extensively at CERN for studying the DA, with a typical study consisting of up to $\sim 2 \times 10^6$ individual tracking simulations over $10^5 - 10^6$ turns (see Fig. 2 for an example).

Fig. 2. Survival plot of a fine phase-space scan for the LHC. The simulations was divided in task of 60 particles pairs covering the phase space in 2σ and $1.5°$ steps.

Tracking data are post-processed during the study and summary files, containing the main results of the simulation for each initial condition, are returned back to the user. In particular, tracking summary files for each initial condition identify particle loss/survival, final surviving turn numbers and the inferred particle amplitudes.

Particle's invariants are calculated for each initial condition and based on the average invariant over a user-defined range of turns. An initial estimate of the invariant is obtained by assuming no coupling between the planes of motion, via the usual relation for the Courant–Snyder ellipse, e.g. for the horizontal plane. Alternatively, an estimate of the decoupled single-particle emittance for the three oscillation modes can be calculated from the eigenvectors of the motion ($\bar{v}$), which may be constructed from the one-turn map, see for example Ref. 26. Various parameters relevant to the nonlinear motion, such as smear and detuning, are also evaluated.

In addition to quantities relevant to particle survival, estimates of the long-term stability are obtained through a Lyapunov-like analysis performed by examining the phase-space separation of initially close by particle pairs. In particular, the angular separation in phase for the three oscillation modes

$$\frac{1}{\pi}\sqrt{\frac{\Delta\phi_1^2 + \Delta\phi_2^2 + \Delta\phi_3^2}{N}},$$

where $N = (1, 2, 3)$ for (2D,4D,6D) motion, respectively, is considered. Linear fits to the logarithm of the separation as function of the logarithm of the turn number identify the maximum separation rate between each particle pair. This quantity is returned, along with the maximum separation in phase, to provide approximate indicators of the onset of chaotic motion in place of the far more computationally intensive Lyapunov exponents.[36,37]

Summary files for the outcome of each initial condition are collectively post-processed by the user using external tools, in order to identify minimum boundaries

Fig. 3.   Frequency map using the NAFF method for LHC in the presence of a tune modulation in the triplets around the experimental areas with frequency of 550 Hz and relative amplitude of 10^{-7}.

in the (σ_x, σ_y) space for particle survival over the tracked number of turns, as well as to study the evolution of DA as a function of the turn number.

2.9. *Frequency analysis*

A collection of routines for frequency analysis has been linked in SixTrack, namely PLATO[38] and a C++ implementation of the Numerical Analysis of Fundamental Frequency (NAFF) algorithm have been developed.[39] These algorithms allow for a more refined, compared to plain Fast-Fourier Transformations (FFT), tune determination with a much faster convergence, i.e. requiring a shorter number of turns. By comparing the tune determination at different time intervals, diffusive frequency maps can be computed.[40] With the resolution of the frequency map, many resonance lines become visible (example shown in Fig. 3).

2.10. *Input and output*

Initial conditions can be given in amplitude steps or taken from an external file. A dump module offers multiple ways to extract tracking data both in terms of type observable (physical coordinates, canonical coordinates, normalized coordinates, averages and first-order distribution momenta) in a selection of turns and observation points. Data are written in text (i.e. ASCII) and, in a few cases, a binary option is also available. Support for output of simulation data to Hierarchical Data Format (HDF5[f]) files and ROOT[41] using the XrootD protocol is also currently being developed.

[f] Available at https://www.hdfgroup.org/HDF5/.

Furthermore, it is planned to develop a new way to generate the distribution that is used as initial conditions for tracking. This will provide the functionality to create matched or mismatched distributions in both physical and normalized coordinates.

2.11. *Interfaces to external programs*

Collimation studies can also be performed by running SixTrack coupled[42] to FLUKA,[43, 44] a specialized simulation tool for describing the interaction (i.e. scattering) of particles with matter. In this configuration, the two codes exchange particles at runtime, with the aim of combining the refined tracking through the accelerator lattice, performed by SixTrack, with the detailed scattering models, implemented in FLUKA, when the beam reaches intercepting devices. Coupling the two software tools in this manner efficiently combines the strengths of each program while, by virtue of using the same FLUKA geometries as for subsequent energy-deposition calculations run with FLUKA, also allows for an excellent level of consistency in the results.

Additionally, a more generalized interface "BDEX" (Beam Distribution EXchange) for interfacing external codes is also included, enabling for example tracking of multiple bunches or coupling to cavity simulation codes. Here, the exact protocol can be implemented as a plug-in to SixTrack.[20]

2.12. *Building and testing*

A CMake-based build and test system has recently been added.[20] This greatly simplified the maintenance of the dependencies between the various build options, as well as the setup for building on the large range of supported platforms.

The testing framework CTest is also provided as part of CMake. For SixTrack, this is used to verify that the executables are still providing the expected output after the code has been modified, and to track the changes to the output. Furthermore, it is used for checking that the results from versions compiled for different platforms are in agreement, which is vital for BOINC. The main benefit of using CTest lies in automating the execution of tests and the generation of pass/fail summary output. Full tests run automatically every night and a subset of tests runs for each proposed code contribution (i.e. pull-request in git jargon).

2.13. *Performance*

Thanks to the recent re-factoring, the internal particles arrays are fully dynamic, therefore the number of particles that can be tracked in parallel is limited by the system memory and not by a build-time flag. A machine model like the LHC, using about 18 k elements, including 4.6 k high-order multipole blocks among them, needs about 220 μs per particle per turn on a single CPU core at 3.4 GHz. Typical studies require in the order of $10^9 - 10^{12}$ particle turns or even more for parameter scans. For this reason, SixTrack is often used in conjunction with high-performance computing (HPC) facilities described in the following section.

3. SixDesk Runtime Environment

The SixDesk environment[45] is the simulation framework used to manage and control the large amount of information necessary for, and produced by, SixTrack studies of DA. It supports CERN batch systems[g] as well as the BOINC platform for volunteering computing available at the LHC@Home project.[6] The SixDeskDB post-processing tool collects data from SixDesk, performs post-processing analysis and prepares reports and plots. It also offers a Python API for interactive analysis. Similarly to the SixTrack code, the SixDesk environment and SixDeskDB are continuously updated, extending the coverage of the studies and keeping the environment up to date with the latest developments in the CERN IT infrastructure.

3.1. *LHC@Home and the CERN batch system*

Volunteer computing has been used successfully at CERN since 2004 with the LHC@Home project; it has provided additional computing power for CPU-intensive applications with small data sets, as well as an outreach channel for CERN activities. LHC@Home started off with SixTrack, which had been successively ported from mainframe to supercomputer to emulator farms and PCs. In order to run on the largest number of volunteer computers, SixTrack is compiled for the most common operating systems, architectures and CPU instruction sets.

In terms of computing power provided by the volunteers to LHC@home, the system is capable of handling 1×10^5 tasks on average, with peaks of 3.5×10^5 tasks simultaneously running on 2.4×10^4 hosts observed during intense SixTrack simulation campaigns (see Fig. 4). Every SixTrack task is run twice to eliminate random host errors and minimize the impact of a failing host. The LHC@Home capacity available for SixTrack can be compared to the average of 2.5×10^5 running tasks on 1.4×10^5 processor cores in the CERN computer centre, which is fully loaded with tasks from analysis and reconstruction of collisions recorded by LHC experiments, and has limited spare capacity for beam dynamics simulations.

The CERN batch system is presently managed by means of the High Throughput Condor (HTCondor) package. Contrary to BOINC, most suitable for a steady stream of work units, the CERN batch system provides users with a responsive computing resource. Also, contrary to LHC@Home, no redundancy is implemented during task submission since the code runs in a controlled environment, although very rarely hardware errors do appear in the results.

3.2. *Developments*

One of the main development lines of SixDesk is porting the collimation studies to the BOINC platform for volunteer computing. This entailed a thorough revision of the collimation part, currently ongoing, to make results numerically stable and

---

[g]CERN Batch service, http://information-technology.web.cern.ch/services/batch.

Fig. 4. Summary of tasks and users during the last two years (upper frame: 2017; lower frame: 2018 to date) of the SixTrack application in the LHC@Home project. The number of users tends to increase, with the number of tasks absorbing most of the time all pending tasks.

reproducible across platforms. The possibility to interrupt and restart the computation (check-point/restart capability), which is essential to run on BOINC, is being added as well.

Other lines of development include: the addition of new parameters for dedicated scans of DA, the possibility of running chains of jobs in BOINC, for simulating extended periods of beam time in the ring; a pre-filtering stage of submission to the CERN batch system prior to submission to BOINC, to avoid short tasks, with consequent inefficient use of volunteer resources, like bandwidth and time.

4. Main Features of SixTrackLib

In the context of single-particle simulations, tracking requires no interaction between the calculations carried out for each particle. The memory requirement for

representing each particle is about 160 bytes. The machine description can, over a single turn, be considered constant, with different elements and sections of the ring requiring different amounts of local resources. SixTrack presents itself as an ideal candidate for a parallel implementation because computations are strongly CPU bound with inherent parallelism and resource requirements do not scale with the number of parallel processes.

Introducing parallelism into a mature code-base-like SixTrack, even from such a favorable starting point, is challenging. It entails a high level of complexity due to competing paradigms and concepts of parallel computing. In particular, a fast-changing technological landscape in combination with a diverse, multi-vendor and long-tailed selection of hardware available via initiatives like LHC@Home,[6] as well as the realities of limited development resources are the main decision-making factors. These and other constraints motivated the design and ongoing development of SixTrackLib[h], a clean-room implementation of the core functionality of SixTrack as a stand-alone library, allowing users to off-load the task of particle tracking onto supported HPC resources. The library provides (a) a representation for a set of particles, (b) a set of beam elements such as drifts, magnetic multipoles (e.g. dipoles, quadrupoles, sextupoles, etc.), RF cavities, 4D and 6D beam–beam elements, etc. (c) a set of maps describing the tracking of particles over the beam elements, (d) a beam monitor lattice element, allowing to represent the functionality of beam instrumentation devices such as beam-position monitors, beam-profile monitors, etc., (e) a range of contexts representing different parallelism and high-performance paradigms (e.g. auto-vectorized CPU code, OpenCL,[46] CUDA[47]) and the associated computing nodes within these paradigms, (f) a dedicated generic buffer for managing collections of particles, beam elements and for transfer data to and from and the computing nodes and (g) a set of high-level APIs in C, C++ and Python.

The design of the library aims to completely separate the parts containing the physics (i.e. the particles, the beam elements and the tracking maps) in a header-only module usable from C-like languages (e.g. C, C++, CUDA, OpenCL, etc.) from the platform specific implementation pertinent to different computing environments and paradigms. Expressing the components in the abstracted header-only physics module allows the reuse of the modeled physics across all supported computing paradigms and to easily tap into the functionality of the tracking library from concrete and already existing applications and frameworks. Another benefit of this approach is the ability to extend the physics (for example by adding additional or alternative implementations for beam elements) without the need to address different computing backends (or, vice versa, being able to add support for additional HPC environments without having to revisit the existing collection of physics-related code-paths). Automatically generated code to aid users in such scenarios and to improve the maintainability of the code by reducing the library's footprint in terms of lines-of-code is currently under development.

---

[h]SixTrackLib source code repository, http://github.com/SixTrack/SixTrackLib.

Fig. 5.   Absolute numerical difference of SixTrackLib and pysixtrack simulations of the LHC lattice over 100 turns.

SixTrackLib is developed in concert with pysixtrack[i], a reference implementation of the SixTrackLib physics using pure Python. It focuses on a straight forward and easy to understand version of the modeled physics, numerical reproducibility and numerical precision (by virtue of allowing the use of extended precision numerical types present in Python) rather than performance or scalability. By comparing results between pysixtrack, SixTrackLib and SixTrack, the feasibility of the chosen design, the numerical accuracy and the performance can be evaluated and quality assured as part of the test-driven development efforts.

Consider for example a simulation on a lattice representing the LHC over a sequence of 100 turns while comparing the results generated via SixTrackLib against pysixtrack (Fig. 5) and SixTrack (Fig. 6). The presented magnitude of the absolute numerical difference in each of the six degrees of freedom is within admissible limits to prove the feasibility of the chosen design of both SixTrackLib and pysixtrack and lays the ground-work for further investigations and optimizations into this aspect.

Concerning an evaluation of the performance improvements achievable by using the parallel hardware capabilities of CPUs and GPUs via SixTrackLib and given the inherent parallelism of the problem outlined above, the main limiting factors for achieving high scalability and numerical throughputs are expected to be (a) the finite availability of resources such as registers and high-bandwidth/low-latency memory on computing nodes and (b) the ability to compensate for any occurring

---

[i]pysixtrack source code repository, https://github.com/rdemaria/pysixtrack.

Fig. 6.   Absolute numerical difference of SixTrackLib and SixTrack simulations of the LHC lattice over 100 turns.

latencies by having enough parallel tasks scheduled to prevent computing units from stalling or idling.

In order to study the impact of these effects, it is illustrative to simulate the same lattice using different transfer maps resulting in the same physics but requiring different amounts of hardware resources (e.g. registers). The by far most demanding transfer maps present in SixTrackLib as of this writing are the 4D and 6D beam–beam interaction implementations. By comparing the performance of simulating a lattice without beam–beam interactions once with a streamlined set of maps online during the simulation and once with passive (i.e. unused) support for beam–beam interactions available, the effect can be illustrated and its impact in real-world scenarios estimated.

Consider to this end the simulation of $1 \leq N \leq 10^7$ particles on a lattice representing the LHC without beam–beam interactions using SixTrackLib, as presented in Fig. 7. Evaluating the impact of the beam–beam maps using the OpenCL backend is especially interesting as it allows a direct comparison of the performance implications on both CPU and GPU configurations across vendors. The presented results encompass both CPUs (Intel Xeon E5-2630 20 × 2.2 GHz hyperthreads, AMD Ryzen 1950X, 16 × 3.40 GHz hyper-threads) and a range of both high-end GPUs (NVIDIA Tesla V100 PCIe 16GB GPU) as well as consumer-grade graphical processors (NVIDIA GTX 1050Ti 4GB, AMD RX560 4GB). The results demonstrate parallel speed-ups approaching (for large N) factors of 10^1 to and exceeding 10^2.

Fig. 7. Benchmarking results of a LHC study using SixTrackLib on different target hardware. Increasing the complexity of the parallel code by enabling beam–beam elements (BB), but not using them, leads to decreased performance on GPUs, but not on CPU-based systems.

The impact of passively enabling the beam–beam interactions (labeled using BB in Fig. 7) in the parallel tracking code leads, unsurprisingly, to virtually no change on CPU-based systems; these systems are not as sensitive to register pressure and resource starvation as graphical processor-based environments. For the GPU-based configurations, the increased complexity and pressure on resources impairs the performance on all studied systems, although the degree of the impact and the economics of keeping complex maps on-line even if the features are not actively used are differing considerably even within the context of this limited set of considered hardware configurations. The results motivate studying ways to split the monolithic parallel code into smaller specialized blocks and to provide APIs and facilities to execute these blocks in sequence, thereby trading in synchronization and dispatching overheads for a potentially better utilization of hardware resources.

5. Conclusions

The SixTrack tracking code is the main code used to simulate long-term stability, collimation cleaning and machine failure scenarios in the LHC, super proton synchroton (SPS) and FCC due to its unique features of speed and integration with HPC resources. It comes with a fully developed running environment to easily perform the massive numerical simulations that include scans on beam and ring parameters and the option of using different computing resources, from standard batch services to volunteer computing.

In spite of its maturity, SixTrack is still in an intense development phase. On the short time scale, it is planned to merge all the code and combine all the features developed within the framework of the studies of the LHC collimation system into a single code-base. On a longer time scale, the main lines of development include tighter integration of existing features, interoperability with other codes and deployment on new architectures such as GPUs.

Acknowledgments

References

1. F. Schmidt *et al.*, Sixtrack version 4.2.16 single particle tracking code treating transverse motion with synchrotron oscillations in a symplectic manner, CERN/SL/94-56 (2012).
2. R. Bruce, *et al.*, Status of sixtrack with collimation, in *Proc. Tracking for Collimation Workshop*, CERN, Geneva, Switzerland, 2018, pp. 1–10.
3. A. Mereghetti *et al.*, Sixtrack for cleaning studies: 2017 updates, in *Proc. 8th Int. Particle Accelerator Conf.*, Copenhagen, Denmark, May 2017, pp. THPAB046.
4. L. Deniau *et al.*, Upgrade of MAD-X for HL-LHC project and FCC studies, in *Proc. ICAP'18*, Key West, FL, USA, Oct 2018, paper TUPAF01.
5. D. P. Anderson, BOINC: A system for public-resource computing and storage, in *Proc. 5th IEEE/ACM Int. Workshop on Grid Computing*, Pittsburgh, USA, 2004, pp. 4–10.
6. J. Barranco *et al.*, *Open Eng.* **7**, 378 (2017).
7. G. Ripken and F. Schmidt, A symplectic six-dimensional thin-lens formalism for tracking, DESY 95-63, CERN/SL/95-12(AP) (1995).
8. K. Heinemann, G. Ripken and F. Schmidt, Construction of nonlinear symplectic six-dimensional thin-lens maps by exponentiation, DESY 95-189 (1995).
9. D. P. Barber *et al.*, A non-linear canonical formalism for the coupled synchro-betatron motion of protons with arbitrary energy, DESY 87-36 (1987).
10. P. D. Hermes, Heavy-ion collimation at the Large Hadron Collider — Simulations and measurements, Ph. D. thesis, University of Münster, Münster, Germany (2016).
11. M. Fjellstrom, Particle tracking in circular accelerators using the exact Hamiltonian in sixtrack, CERN-THESIS-2013-248 (2013).
12. A. Latina and R. De Maria, RF multipole implementation, CERN-ATS-Note-2012-088 TECH (2012).
13. L. H. A. Leunissen *et al.*, *Phys. Rev. ST Accel. Beams* **3**, 124002 (2000).
14. A. Patapenka *et al.*, Simulation in support of wire beam-beam compensation experiment at the LHC, in *Proc. 2016 North America Particle Accelerator Conf.*, Chicago, IL, USA, Oct 2016, pp. TUPOB17.
15. V. Previtali *et al.*, Numerical simulations of a hollow electron lens as a scraping device for the LHC, in *Proc. 4th Int. Particle Accelerator Conf.*, Shanghai, China, May 2013, pp. MOPWO044.
16. M. Fitterer *et al.*, Implementation of hollow electron lenses in sixtrack and first simulation results for the HL-LHC, in *Proc. 8th Int. Particle Accelerator Conf.*, Copenhagen, Denmark, May 2017, pp. THPAB041.
17. E. Quaranta, Investigation of collimator materials for the high luminosity Large Hadron Collider, Ph. D. thesis, Politecnico di Milano, Milan, Italy (2017).
18. S. Tygier *et al.*, Simulating novel collimation schemes for high-luminosity LHC with Merlin++, in *Proc. 10th Int. Particle Accelerator Conf.*, Melbourne, Australia, May 2019, pp. MOPRB060.
19. J. Allison *et al.*, *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.* **835**, 186 (2016).

20. K. Sjobak *et al.*, New features of the 2017 sixtrack release, in *Proc. 8th Int. Particle Accelerator Conf.*, Copenhagen, Denmark, May 2017, pp. THPAB047.
21. TOTEM Collab. (G. Anelli *et al.*), *J. Instrum.* **3**, S08007 (2008), doi:10.1088/1748-0221/3/08/S08007.
22. T. Sjstrand *et al.*, *Comput. Phys. Commun.* **178**, 852 (2008).
23. M. Berz, *Nucl. Instrum. Methods A* **258**, 431 (1987).
24. M. Berz, *Part. Accel.* **24**, 109 (1989).
25. M. Berz, *Modern Map Methods in Particle Beam Physics* (Academic Press, 1999).
26. H. Mais and G. Ripken, Theory of coupled synchro-betatron oscillations, DESY-M-82-05, DESY (1982).
27. K. Sjobak *et al.*, Dynamic simulations in sixtrack, in *Proc. Tracking for Collimation Workshop*, CERN, Geneva, Switzerland (2018), pp. 123–134.
28. L. Bottura, L. Walckiers and R. Wolf, Field errors decay and snap-back in LHC model dipoles, CERN-LHC-Project-Report-55 (1996).
29. A. Santamaría García *et al.*, Limits on failure scenarios for crab cavities in the HL-LHC, in *Proc. 6th Int. Particle Accelerator Conf.*, Richmond, USA, May 2015, pp. THPF095.
30. K. Sjobak *et al.*, Time scale of crab cavity failures relevant for high luminosity LHC, in *Proc. 7th Int. Particle Accelerator Conf.*, Busan, Korea, May 2016, pp. THPOY043.
31. A. Santamaría García *et al.*, Machine protection from fast crab cavity failures in the high luminosity LHC, in *Proc. 7th Int. Particle Accelerator Conf.*, Busan, Korea, May 2016, pp. TUPMW025.
32. S. J. Wretborn *et al.*, Study of off-momentum losses at the start of the ramp in the Large Hadron Collider, CERN-ACC-NOTE-2017-0065 (2017).
33. G. Apollinari *et al.*, High-luminosity Large Hadron Collider (HL-LHC): Technical design report v. 0.1, CERN-2017-007-M, CERN, Geneva, Switzerland.
34. D. Mirarchi *et al.*, Hollow electron-lenses assisted collimation and plans for the LHC, in *Proc. 61st ICFA Advanced Beam Dynamics Workshop on High-Intensity and High-Brightness Madron Beams*, Daejeon, Korea, June 2018, pp. TUP1WE02.
35. G. Stancari, Calculation of the transverse kicks generated by the bends of a hollow electron lens, FERMILAB-FN-0972-APC, Fermilab, Batavia, Illinois, USA (2014).
36. F. Schmidt, F. Willeke and F. Zimmermann, Comparison of methods to determine long-term stability in proton storage rings, CERN-SL-91-14-AP (1991).
37. M. Böoge and F. Schmidt, Estimates for longterm stability for the LHC, LHC Project Report 114 (1997).
38. M. Giovannozzi *et al.*, *Nucl. Instrum. Methods A* **388**, 1 (1996).
39. S. Kostoglou *et al.*, Development of computational tools for noise studies in the LHC, in *Proc. 8th Int. Particle Accelerator Conf.*, Copenhagen, Denmark, May 2017, pp. THPAB044.
40. J. Laskar, C. Froeschle and C. Celletti, *Phys. D* **56**, 253 (1992).
41. R. Brun and F. Rademakers, *Nucl. Instrum. Methods Phys. Res. A* **389**, 81 (1997).
42. E. Skordis *et al.*, FLUKA coupling to sixtrack, in *Proc. Tracking for Collimation Workshop*, CERN, Geneva, Switzerland (2018), pp. 17–26.
43. G. Battistoni *et al.*, *Ann. Nucl. Energy* **82**, 10 (2015).
44. A. Ferrari, P. R. Sala, A. Fassò and J. Ranft, FLUKA: A multi-particle transport code, CERN-2005-10, INFN/TC_05/11, SLAC-R-773 (2005).
45. E. McIntosh and R. De Maria, The sixdesk run environment for sixtrack, CERN-ATS-Note-2012-089 TECH (2012).
46. J. Stone *et al.*, *Comput. Sci. Eng.* **12**, 3 (2010).
47. J. Nickolls *et al.*, *ACM Queue* **6**, 2 (2008).