H5PartROOT—A VISUALIZATION AND POST-PROCESSING TOOL FOR ACCELERATOR SIMULATIONS

Thomas Schietinger
Paul Scherrer Institut, CH-5232 Villigen PSI, Switzerland

Abstract

Modern particle tracking codes with their parallel processing capabilities generate data files of the order of 100 Gigabytes. Thus they make very high demands on file formats and post-processing software. H5PartROOT is a versatile and powerful tool addressing this issue. Based on ROOT, CERN's object-oriented data analysis framework developed for the requirements of the LHC era, and the HDF5 hierarchical data format, supplemented by an accelerator-specific interface called H5Part, H5PartROOT combines the statistical and graphical capabilities of ROOT with the versatility and performance of the HDF5 technology suite to meet the needs of the accelerator community. Providing the user with both a graphical user interface (data browser) and a shared library to be used in an interactive or batch ROOT session, H5PartROOT passes on the full power of ROOT without presupposing any knowledge about the intricacies of either ROOT or C++.

INTRODUCTION AND MOTIVATION

Three-dimensional particle simulations (e.g., OPAL [1]) follow the trajectories of a large number (up to 10^9 and more) of macro-particles through space as they are influenced by external (electro-magnetic and or gravitational) and internal (space charge) fields. The result of such a simulation is typically stored as a sequence of time steps. Each time step contains some quantities (often scalars or 3-vectors for the three spatial dimensions) describing properties of the macro-particle ensemble (bunch) as a whole (e.g., centroid position, mean particle energy etc.) and, if detailed analysis of the bunch evolution is desired, a full dump of the macro-particle phase space (i.e., positions and momenta of all macro-particles in the simulation).

With the increasing size of datasets produced by such simulations, swift post-processing becomes an issue of paramount importance. To address the problem we created a tool based on ROOT, the data analysis software used by CERN and its user community to analyze the vast amounts of data produced by the Large Hadron Collider, and HDF5, an extremely versatile and powerful data file format enjoying growing popularity throughout the scientific community. The considerable power and flexibility of both HDF5 and ROOT come at the prize of rather complex user interfaces. To spare the user the learning curves of these packages as much as possible, we built a ROOT application which allows fast extraction of statistical data and generation of publication-quality plots with just a few mouse clicks. Since it makes use of the H5Part interface to HDF5, the application is called *H5PartROOT*.

BUILDING BLOCKS

HDF5

HDF5 ("Hierarchical Data Format 5") is a highly sophisticated, "self-describing" data storage format. Originally created by NCSA, it is now supported by the HDF group [2]. The HDF5 technology suite allows the management of extremely large and complex data collections. Its versatile data model can represent very complex data objects and a wide variety of meta-data. The file format is completely portable and puts no limits on the number or size of data objects, making it an ideal format for large accelerator simulations. The HDF5 software library provides various high-level interfaces (C, C++, Fortran 90, Java) and runs on almost every computing platform, from laptops to massively parallel systems. Furthermore HDF5 comes with built-in performance features that optimize access time and storage space as well as a whole set of tools and applications for managing, manipulating, viewing, and analyzing the data.

H5Part

H5Part is a thin software layer on top of HDF5 to facilitate I/O for the simulation of particle accelerators (or any other multi-particle system that evolves in time) [3]. Designed as a portable, high-performance parallel data interface to HDF5 [4], it constrains HDF5's very general data format to a subset useful for three-dimensional particle accelerator simulations, i.e., it knows about time steps, phase space variables etc. H5Part is co-developed by LBNL and PSI.

ROOT

ROOT is an object-oriented data-processing framework developed at CERN for the requirements of the LHC era, i.e., to handle complex datasets of sizes measured in Terabytes [5]. Within the last decade, it has become the data analysis and visualization tool of choice in high-energy physics around the world.

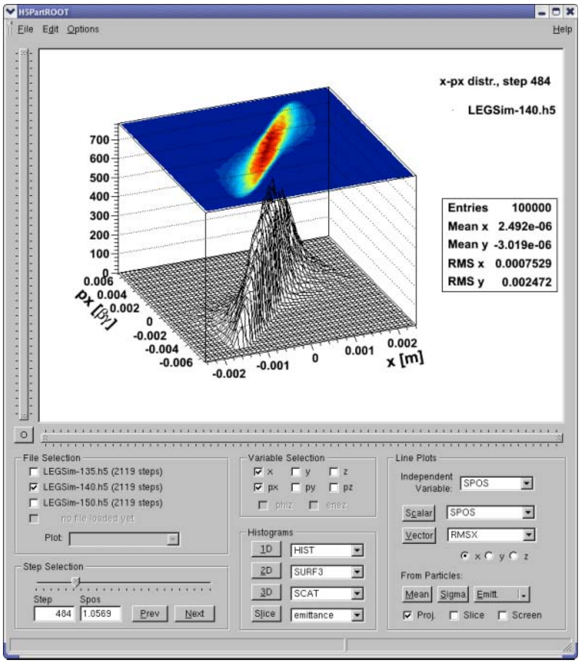

Figure 1: H5PartROOT main GUI showing an example phase space plot.

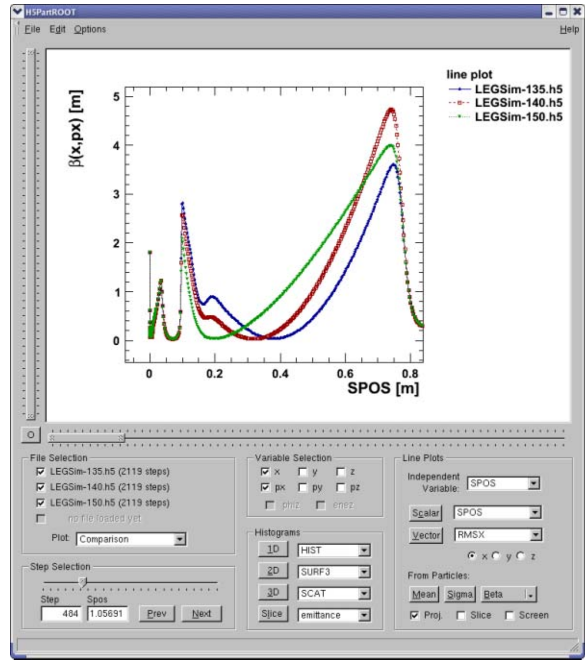

Figure 2: Example of a line plot produced with the H5PartROOT GUI (comparison of beta functions).

Written in C++, ROOT can be used either as an extensive class library to be linked to a C++ main program, or via its C++ interpreter (CINT) as an interactive command-line tool. The ROOT libraries include classes for plotting, fitting and other statistical analysis, as well as graphical user interfaces. ROOT is maintained by a core development team at CERN supported by numerous users around the world.

DESIGN OVERVIEW

The driving design philosophy behind H5PartROOT is the desire to pass on to the user the full power of ROOT without bothering him or her with the (sometimes rather cumbersome) intricacies of the ROOT (or generally C++) syntax. Since this is best achieved by a user-friendly graphical user interface (GUI) a main data browsing window is at the center of the H5PartROOT design. The GUI is implemented as a class, `TH5MainFrame`,[1] which processes all user requests via buttons or pull-down menus and displays information in a graphics window ("canvas") embedded in the GUI frame. A second class, `TH5Dataset`, provides the link to a given HDF5 data file, and returns data processed as plots or statistical quantities as requested by the client. Figure 1 shows a screen-shot of H5PartROOT's main GUI with an example phase-space distribution. By dragging the mouse, the user can rotate the plot or zoom in along the axes (functionalities inherited from ROOT).

While an intuitive GUI is convenient for rapidly browsing one or several datasets, it quickly reaches its limitations when it comes to producing elaborate plots combining information from different sources (e.g., measurement data). For this reason, H5PartROOT can also be used as a shared library within ROOT, which gives the user full access to all `TH5Dataset` methods from within the ROOT environment.

[1] Like ROOT itself, H5PartROOT uses the Taligent coding convention.

FEATURES

Besides simple data representations such as histograms and scatter plots of phase-space distributions, or line plots of time-step attributes, H5PartROOT provides a number of features relevant to particle accelerator physics:

- *Emittance:* computation of the rms emittance $\varepsilon(\vec{x},\vec{x}') = \sqrt{\langle x^2\rangle\langle x'^2\rangle - \langle xx'\rangle}$ for any two phase space variables x, x' from the particle distribution.

- *Twiss parameters and phase space ellipses:* Computation of the rms Twiss parameters $\alpha(\vec{x},\vec{x}')$, $\beta(\vec{x},\vec{x}')$ from the particle distribution, graphical representation as a phase space ellipse.

- *Bunch slicing:* Evaluation of the above quantities or graphs for longitudinal slices (subsets of the particle distribution as a function of z) of the bunch (see Fig. 3). Both length and relative position of the slice within the bunch can be set by the user. Gauss-weighted slices are also supported.

- *Screen projection:* Evaluation of emittance and Twiss parameters for fixed longitudinal position s (as opposed to fixed time).

- *Bunch clipping:* Particle distributions or bunch properties after particles in sparsely populated regions of the bunch have been removed. This feature is typically used to suppress the effect of bunch halos.

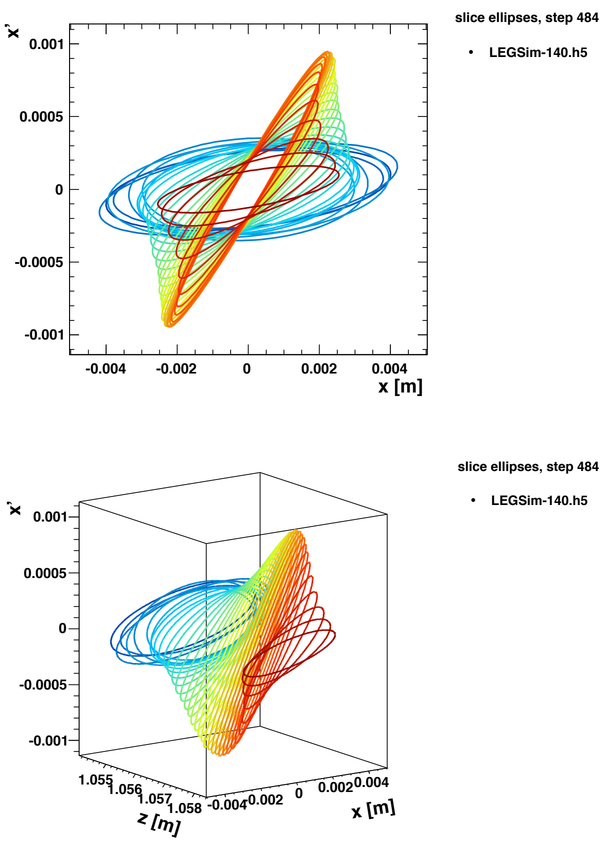

Figure 3: Twiss phase space ellipses for bunch slices in 2D (top) and 3D (bottom) representation.

USAGE

The Graphical User Interface

The main advantage of the GUI consists in the ease with which a user can analyze a dataset and produce publication-quality plots by clicking just a few buttons—without any knowledge of neither ROOT nor C++. The GUI has fields *File Selection*, *Step Selection* and *Variable Selection*, to tell H5PartROOT what to plot and *Histograms* and *Line Plots* to specify how to visualize the data. More sophisticated settings (choice of binning, log scale etc.) are available via an options menu.

Once a plot is done, its embellishments can easily be changed via context menus. Thanks to ROOT's many graphics interfaces, it can then be printed to almost any commonly known graphics format. A particularly useful print format is *C++ Macro*, which will produce the exact C++ code necessary to reproduce the plot inside a ROOT session. This gives the user the possibility to save a result without fixing the exact plot style, which may have to be adapted later to the style of a larger document.

The Shared Library

To take full advantage of all ROOT features, however, it is advisable to use H5PartROOT as a shared library. Depending of the complexity of the task at hand this library may be

- loaded into an interactive ROOT session,
- used inside a ROOT macro, or
- linked to a compiled ROOT executable.

Figure 4 shows an example plot generated by the use of the H5PartROOT shared library.

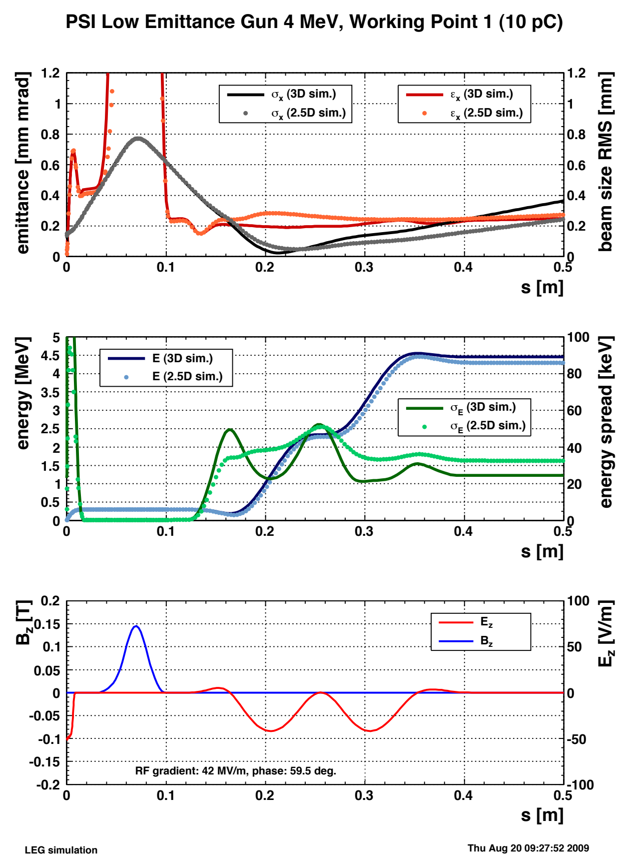

Figure 4: Code comparison plot produced by a ROOT macro after loading the H5PartROOT shared library.

In interactive sessions or macros the library is loaded with the command

```
gSystem->Load("<libpath>/libh5root.so");
```

Once this is done, the user has access to the H5PartROOT classes with all their methods:

```
TH5Style::SetStyle()
```

— *set the H5PartROOT plotting style*

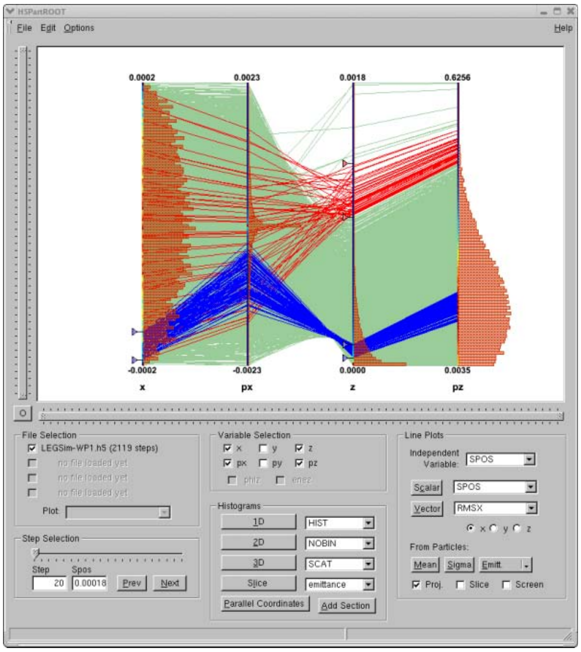

Figure 5: Prototype version of H5PartROOT featuring parallel coordinates. The plot shows a phase space distribution (x, p_x, z, p_z) with one selection in z (red) and another one in x and z (blue).

```
TH5Dataset data("example.h5");
```
— *load data from file example.h5*
```
data.Histo2d("x","y");
```
— *plot x and y (for default step 0, with default range)*

INSTALLATION

The H5PartROOT source code and instructions on how to install and run the application can be downloaded from our web site [6]. Usage of the GNU build system (autotools) ensures straight-forward installation across different platforms.

FUTURE PLANS

Plans to improve and extend H5PartROOT include some more particle selection features in the GUI, better 3D plotting functionality (mainly by integration of the OpenGL graphics engine [7]) as well as new visualization concepts. An example are parallel coordinates, which have recently been added to ROOT [8]. Figure 5 shows a prototype version of the H5PartROOT data browser with parallel coordinates.

Further down the line we plan to adapt H5PartROOT to visualize slice-based simulations such as Homdyn [9]. Another possible development would be the parallelization of the analysis code given sufficient user interest.

CONCLUSION

ROOT's excellent data analysis and visualization capabilities, developed for the high-energy physics community, can also be harnessed in the context of computational accelerator physics. In conjunction with HDF5 (and H5Part) it provides a powerful yet elegant solution to the post-processing needs of large-scale particle simulations. With H5PartROOT, we have presented our implementation of this promising approach.

ACKNOWLEDGMENTS

It is a pleasure to acknowledge substantial contributions to H5PartROOT by Andreas Adelmann, Bo Liu and Yves Ineichen, as well as valuable feedback by many other users.

REFERENCES

[1] A. Adelmann et al., "The Object Oriented Parallel Accelerator Library (OPAL), Design, Implementation & Application", these proceedings.

[2] The HDF5 Group, `http://www.hdfgroup.org/HDF5/`.

[3] H5Part: a Portable High Performance Parallel Data Interface to HDF5, `http://vis.lbl.gov/Research/AcceleratorSAPP/`; code available from `https://codeforge.lbl.gov/projects/h5part/`.

[4] A. Adelmann et al., in Proceedings of Particle Accelerator Conference (PAC), Albuquerque, NM, USA, 2007, p. 3396.

[5] ROOT: an object-oriented data analysis framework, `http://root.cern.ch`.

[6] H5PartROOT: a ROOT Based Graphical User Interface to H5Part, `http://amas.web.psi.ch/tools/H5PartROOT`.

[7] M. Tadel, J. Phys.: Conf. Ser. 119 (2008) 042028.

[8] O. Couet, J. Phys.: Conf. Ser. 119 (2008) 042007.

[9] M. Ferrario et al., SLAC-PUB-8400 (2000).