

Interactive Methods for Exploring Particle Simulation Data

Christopher S. Co[†], Alex Friedman^{‡§}, David P. Grote^{‡§}, Jean-Luc Vay[§], E. Wes Bethel[§], and Kenneth I. Joy[†]

Abstract

In this work, we visualize high-dimensional particle simulation data using a suite of scatterplot-based visualizations coupled with interactive selection tools. We use traditional 2D and 3D projection scatterplots as well as a novel oriented-disk rendering style to convey various information about the data. Interactive selection tools allow physicists to manually classify “interesting” sets of particles that are highlighted across multiple, linked views of the data. The power of our application is the ability to correspond new visual representations of the simulation data with traditional, well understood visualizations. This approach supports the interactive exploration of the high-dimensional space while promoting discovery of new particle behavior.

1. Introduction

The study of multiparticle dynamics spans a wide variety of applications, from galaxy simulations to particle accelerator design. In particular, *heavy-ion fusion* (HIF) is a topic of intense current scientific study for it is a potential source of energy that is low-cost, widely available, and environmentally friendly [HIF, Fri02, CBH*03]. Particle accelerators themselves are expensive facilities that cost hundreds of millions of dollars to design, construct, and then operate over a period of decades. Prior to construction of such a facility, simulations are developed to produce the best possible design. Accelerator simulations use discrete-particle models that are carried out in *phase space*, where each particle is characterized by position (x, y, z) and momentum (p_x, p_y, p_z) . The data produced by accelerator simulations are large, time-varying, and multidimensional. The simulation data contain representations of complex physical phenomena. Compounding matters is the fact that the physical phenomena hidden in the data often have never been observed before. As researchers undertake data analysis, they are not sure exactly what features in the data are meaningful, nor are they sure exactly how to quantitatively express “regions of interest” or “hotspots” in the simulation data to aid and accelerate

discovery. The need for effective visualization capabilities to aid in data understanding motivates our work.

The optimal design of particle accelerators remains a difficult problem, primarily because of the complex nature of the forces involved. Two primary physical forces contribute to the shape and behavior of the beam. In most particle accelerators, beams of charged particles are controlled using externally applied electromagnetic fields at intervals along the length of the accelerator. The second force is the electromagnetic field induced by the charged particles themselves—the beam generates its own electromagnetic field. Understanding the interplay between the two electromagnetic field contributors is critical to controlling the particle beam, and requires a self-consistent field description, since the latter factor is a function of particle motion.

During data analysis, scientists are interested in understanding a number of issues. For example, based upon prior experience, they know that most of the beam particles are “well behaved,” but a small number of particles drift away from the core of the beam. The particles that diverge from the core form what is known as a “halo,” and the halo particles pose a potential problem for the stability of the beam. The halo is the low density region of particles located “far” from the central core of the particle beam [MQR99]. Particles in the halo region are more likely to collide with accelerator walls. When they do so, electrons and other debris can be emitted that can lead to degradation of the beam quality. Also, high energy particles can activate the wall material, rendering it radioactive. In extreme conditions, the halo particles may potentially cause a rupture to the accelerator vac-

[†] Institute for Data Analysis and Visualization (IDAV), University of California, Davis, {co,joy}@cs.ucdavis.edu

[‡] Lawrence Livermore National Laboratory, Livermore, CA, {af,grote1}@llnl.gov

[§] Lawrence Berkeley National Laboratory, Berkeley, CA, {jl-vay,ewbethel}@lbl.gov

uum vessel, thereby causing a loss of vacuum and potential damage to surrounding structures.

Accelerator physicists seek to understand the causal factors that produce halos so they can be suppressed. They also seek to understand the factors that result in a “well behaved” beam: placement and strength of controlling applied electromagnetic fields, cross-sectional shape of the containment device, etc. Their primary tool for scientific inquiry is simulation and the subsequent data analysis.

Historically, the scatterplot has been the staple method for visualizing the particle data generated by high energy physics simulations and experiments. Scientists create many “small multiples” [Tuf90] that depict different 2D projections of phase space. The different projections cover the gamut of permutations: x vs. phase- x , y vs. phase- y , and so forth. Each different scatterplot provides a depiction of spatial partial distribution in some projection. Scientists then review the many projections in search of distributions characteristic of known phenomena. For instance, an s-shaped distribution in a plot of an (x, p_x) projection of the phase space indicates the accumulated action of non-ideal, anharmonic forces that do not vary linearly as functions of the transverse coordinates. With practice, physicists are able to construct a mental model that is the integration of all 2D projections. While powerful individually, each particular type of phase space projection only offers a limited representation of the complex behavior exhibited by the simulation data.

While the shape of the beam is important to physicists, what is more important—and the scientific question we help to answer with the work presented in this paper—is how a given beam shape or configuration came to be in the first place, and how it will evolve over time. The general workflow we have refined in conjunction with physicists is to “classify” simulation particles that are part of a halo population at some point in time, then observe their behavior at earlier and later points in time. This approach is the inverse to classification based upon data value ranges. In this case, the scientist can quickly visually identify a halo region due to some characteristic shape or through intuition, then interactively classify the particles as being part of a halo region. Our workflow model lets them then track the evolution of the halo particles forward and backward in time. The intent is to facilitate rapid discovery of the conditions that led to the formation of a beam halo in the first place. The techniques we describe in this paper—designed and implemented as part of an accelerator modeling project—are designed to facilitate such discoveries.

Our work improves the speed and accuracy of visual data analysis of particle data generated by accelerator modeling simulations or experiments. We describe several straightforward techniques that combine into a powerful and easy-to-use application. Our application, PPaint, provides the ability to generate small multiples, which is the same familiar presentation that the high energy physics community has come

to know and understand. We add the ability to interactively select particles, and then have the selected particles appear in each of the multiple linked views. Interaction with PPaint as well as implementation details are provided in Section 4. The views may be a combination of 2D or 3D projections, depending upon the user’s preference. This metaphor, while simple, is powerful since a collection of particles in a 2D phase space plot may turn out to occupy a scattered region in 3-space. We add a novel particle visualization method as an option to enhance visual comprehension of complex, high-dimensional data. We describe this novel rendering approach along with other aspects of rendering in PPaint in Section 3. Once particle halos are selected, the particles comprising a halo population may then be tracked forward or backward in time to facilitate understanding of the phenomena that led to the formation of the halo. The enhancements our visualization tool provides help streamline the process of forming a mental model of the data.

2. Related Work

Perhaps the best-known visualization tool in the high energy physics community is the “Physics Analysis Workstation” software, or PAW [PAW]. PAW is an analysis and visualization toolkit designed to work with data generated by the GEANT4 simulation code. GEANT4 has been traditionally used to model particle detectors [A*03]. PAW provides a rich collection of charting and plotting capabilities that have evolved over time to be especially well-suited to the needs of the high energy physics community. PAW and tools like it are not well-suited to the demands posed by today’s large datasets. It does not provide the ability to spatially select data points, nor to track them over time and display multiples in a mixture of modalities. These shortcomings have provided substantial motivation for our work.

Standalone methods for 3D visualization of particle data have included volume rendering methods. One such project used volume rendering to show halo particles in high-resolution accelerator data [MQR99]. Recently, a hybrid particle-volume rendering method has proven effective in showing both the overall shape of the beam using direct volume rendering, but retaining the detail of individual particles using geometry [WMQR02, WMM02, MSW*02]. In both methods, interactively specifying a transfer function plays a critical role in allowing the scientist to explore the volume rendered phase space projections.

With regard to user interaction methods applied to visualization research, others have successfully employed multiple linked views and painting-style user interfaces. Becker and Cleveland [BC87] used brushes to label and delete data points across multiple linked scatterplots. Martin and Ward implemented high-dimensional brushing tools that operated in data space [MW95]. Doleisch et al. [DGH03] promote the visualization of essential features in exploring high-dimensional data. To accomplish this, they designed an in-

teractive system where features can be identified by the user through “brushing” and other techniques. They note that such feature extraction is often preferred over automatic and semi-automatic methods, since autonomous approaches still require information from the user as to what features are “interesting.” Incorporating user intervention is essential since scientists possess intuitive understandings of the data that are difficult for autonomous methods to mimic. This was also a major motivating factor in the development of the Protein-Shop program [KMC*03], which is an interactive protein manipulation package for computational biology. Many successful applications, such as decision tree visualization and classification applications have benefited from the use of a painting interface [TM03]. Tzeng et al. [TLM03] also used a painting interface and the assistance of a neural network to perform classification in high-dimensional domains.

In our work, we make use of user-driven classification via a painting interface, as in the work of Teoh and Ma [TM03]. The classified particles are highlighted across multiple views, similar in spirit to the approach taken by Doleisch et al. [DGH03]. This approach leverages the visualization power of each particular rendering style used in each view. Such classification techniques combined with novel rendering techniques and temporal tracking facilitate the creation of a mental model of the data by guiding the user to a deeper understanding of information hidden inside the data.

3. Rendering

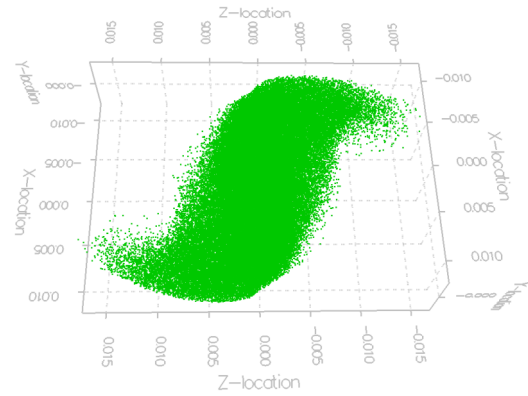
3.1. 2D Scatterplots

We use standard 2D projections of phase space to create the small multiples familiar to physicists. In each scatterplot, one data component provides the ordinate and another provides the abscissa. A point plotted at the resulting coordinate represents the presence of a particle at that point in phase space. For example, the x -component of position provides an ordinate while the x -component of the momentum (known as p_x) provides the abscissa. Together, these combine to produce the 2D Cartesian coordinate (x, p_x) . This style of visualizing the particle data, at a single simulation step or as a movie spanning many simulation steps, is well understood and serves as a reference when evaluating newer visualization techniques applied to the same particle data.

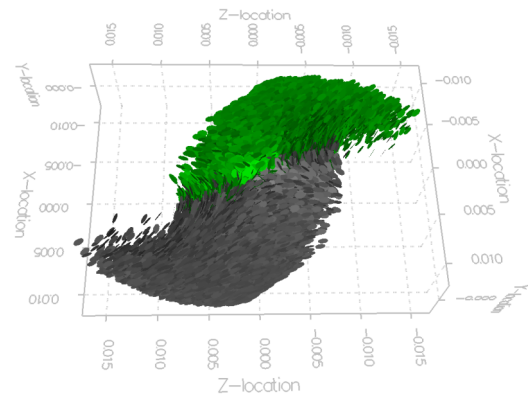
3.2. 3D Scatterplots

A variety of 3D plotting styles are utilized in our application to reveal information in the particle data that would otherwise be hidden from view if using traditional 2D scatterplots.

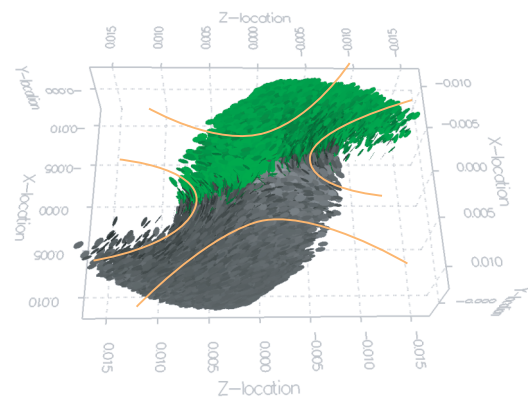
3D Phase Space Projections – using three dimensions rather than two offers the potential to convey more information, particularly when combining interactive 3D transformations with stereo. In terms of understanding depth relationships and 3D structure, use of a static 3D perspec-



(a)



(b)



(c)

Figure 1: 3D scatterplots as rendered by our program. (a) A standard xyz-plot of time step 60 of a 276 time step linear accelerator simulation is shown. (b) The same data set is shown rendered using our novel oriented-disk approach. To a certain extent, the influence of the (c) electromagnetic field on the particles can be observed in this style of rendering by visually piecing together the facets of neighboring disks.

tive view offers little more than a static 2D view. However, adding interactive transformation to the 3D view has been shown to provide a dramatic and measurable increase in comprehension of 3D depth relations and 3D structure [WF96].

Disk Rendering – while the interactive 3D perspective views are better than their static 2D or 3D counterparts, there is still the need to understand the nature of the phase space vector field in conjunction with 3D shape. To help in that regard, we have employed a specialized glyph to depict the phase space vector field. The glyph consists of an oriented and colored disk for each particle. The location of the disk is given by the location (x, y, z) of the particle, while the orientation of the disk is given by the normal vector parallel to the momentum vector $\langle p_x, p_y, p_z \rangle$. To further aid in distinguishing vector field polarity, we use different material properties for front- and back-facing disks. A neutral gray depicts back-facing disks, while a fully saturated color is applied to front-facing disks. Lighting aids in further depicting the orientation of the disks. These disks are essentially surface elements, or *surfels*, as described by Pfister et al. [PvBZG00] in the context of surface representations. This style of rendering is extremely useful in determining the direction of each particle’s momentum vector and enhances the scientist’s understanding of the overall path of individual particles as well as the entire particle beam. Figure 1 shows the disk rendering style juxtaposed with a standard 3D xyz -plot.

3.3. The Use of Color

In our system, color is used in two manners. First, color is used as a label to indicate selection membership. In other words, all the red particles belong to one group selected by the user, while all the green particles belong to a different group. Groups of particles may be made “invisible” to reduce visual clutter. This programmatic feature is illustrated in Figures 2, 4, and 5. The second use of color is to indicate a measurement. Particles are colorized according to a user-defined colormap driven by additional stored or derived quantities associated with each particle, such as the magnitude of the momentum vector. Figure 3 illustrates a particle distribution colormapped according to the magnitude of each particle’s momentum vector. The rest of the figures demonstrate the use of color to highlight particle selection classes.

3.4. Animations

Our system includes an animation feature, such that the simulation data can be examined over time to obtain a higher-level notion of the particle beam path. As particles are selected and isolated, they can also be visually tracked over time by playing back each time step in order. The scientist can pause, rewind and jump to any loaded time step in the sequence to see the path of particles of interest.

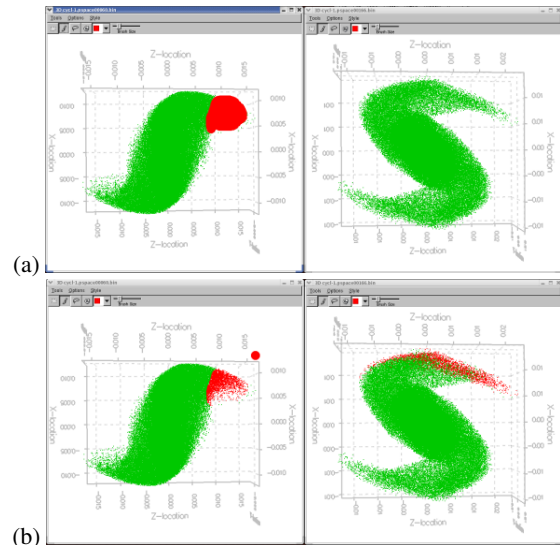


Figure 2: Demonstration of manual particle selection through the use of a painting tool. (a) Particles belonging to the “spiral arm” of an early time step (left) are painted on. (b) These particles are marked red and can be seen in a later time step (right) of the simulation. The red circle represents the paint brush used for selection.

4. Interaction and Implementation

Classification in our program is performed manually with a familiar painting interface. Box selection, lasso selection, and paint brush selection tools provide intuitive methods for selecting particles of interest. As the user selects a set of particles of interest, the selected particles are color-labeled across all of the particle data sets loaded. Manual classification in this particular application is a key feature, since it is through the interaction process that a deeper understanding of particle behavior is obtained. Figure 2 shows a screen capture of the painting selection tool being used to classify particles belonging to the “spiral arm” of a particle distribution.

The application workflow and methods we present here offer a very usable system that facilitates deeper understanding of complex scientific data. Classification of interesting data is performed via user-defined selections interactively, allowing a physicist to explore spatial properties of the particle data. Our disk rendering technique allows information about particle momentum to be conveyed in addition to the spatial distribution of the data. Motion parallax across multiple linked views accelerates understanding of 3D depth relationships and shape. Temporal browsing allows a scientist to understand the evolution of interesting particles. All of these features support the formation of a mental model that is necessary for developing a deeper understanding of the particle beam data.

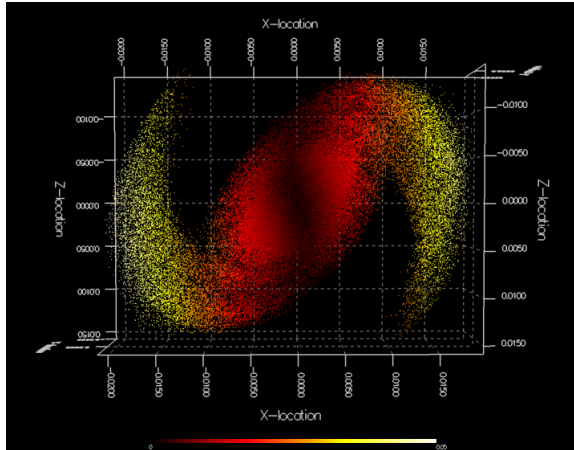


Figure 3: Example showing a xyz-plot of time step 214 of a 276 time step linear accelerator simulation run. Points are colorized according to the magnitude of their momentum vector.

In addition to the spatial information, additional per particle information is required to make user interaction possible. We keep an adjacent array of flags that specifies the membership of each particle to a particular highlight group. In our implementation, we use eight bit flags, thus allowing 255 distinct classifications, reserving zero to indicate membership in the class of non-selected particles. Particle IDs are associated with each particle in a given time step. Since corresponding particles in different time steps of the data may have different relative location in their respective arrays, a mapping from particle ID to array location must be computed. We accomplish this with the use of a hash table. We note that these particle IDs are often maintained as part of the simulation and can be directly imported *as is* into our system.

5. Results

Figure 1 illustrates our novel oriented-disk rendering method. The orientation of each particle is clearly conveyed using lighting combined with two-sided material surface properties. Since the user visually clusters disks together, it is relatively easy to understand the behavior of particle groups, which in turns provides inferences about the influence of the electromagnetic field. For example, Figures 1b and 1c show the influence of the electromagnetic field on clusters of particles. Animations of the simulation are enhanced by this type of glyph, since motion “toward” and “away from” the user’s viewpoint are easily communicated by shading and color.

One stage in the path toward building a full scale driver for heavy-ion fusion is an accelerator experiment that examines most of the design issues of a driver at full scale pa-

rameters. Such an *integrated beam experiment* (IBX) differs from a full-scale fusion driver in that it only has a few beams and does not accelerate to the full energy, and thus is far less costly. This type of reduced experiment can be done since most of the issues occur, or can be studied, at low energies. The data shown in Figure 4 is from 3D simulations of a design of such an experiment.

The typical “hoe”-shaped structure that can be seen in the polar plot is a result of the particles at the transverse edge of the beam (particles colorized red and blue) being over-focused, indicating non-linear focusing/defocusing forces. These non-linearities are due to both inherent geometric aberrations in the injector design and numerical errors in the simulation (the source region was not finely enough resolved). These particles (red and blue) initially form a higher density rim on the beam and later some are expelled from the core of the beam into a halo. The selection tools allow these overfocused particles to be visually tracked, leading to a better understanding of tolerances to such errors.

Sometimes, a beam injector is required that can produce a high current beam with low transverse temperature. Typical injectors use a single, monolithic source of ions. Because of various scaling laws, though, the required emitting area increases at a high power of the current. A way around this poor scaling is to use a large number of small sources, a hundred or more, and merge the resulting beamlets. An experiment testing the concept is in the process of being designed and built. The data from Figure 5 is from 2D, transverse slice, simulations of the design. As the many beamlets merge and mix, short wavelength density waves are launched on the combined beam and travel across it. This process can knock some particles out of the core of the beam into a small halo. A dramatic case can be seen in Figure 5a where four small jets of particles leave the core of the beam. Figures 5b, 5c, and 5d show the particles at earlier time, with the particles that form the later jets selected, and the core of the beam made invisible. Note how each group remains in a relatively tight formation, even though heavy particle mixing is occurring, but is spread over multiple beamlets. This is information which would be difficult to extract without the PPaint program and that can provide a much deeper understanding of the subtle processes involved in the merging.

6. Conclusion

The work we have described in this paper offers a highly practical approach to visual analysis and understanding of particle data generated by accelerator modeling simulations. Our work was conducted in close collaboration with discipline scientists resulting in an application that offers them the ability to more quickly explore and understand simulation results. Using our novel oriented-disk rendering method, physicists are able to quickly visually distinguish between particles that are moving towards or away from the viewer as well as understand shape and flow information of




Figure 4: Correlating points across different views of the same data. An integrated beam experiment (IBX) data set represents a beam after it has undergone significant phase space distortion. (a) An Rv_Rv_Θ -plot of a polar coordinate version of the data is plotted in Cartesian coordinates. R denotes the radial coordinate, v_R the velocity of R , and v_Θ the velocity of the angular coordinate Θ . A typical “hoe”-shaped distribution is observed. (b) The corresponding xp_xp_y -plot of a Cartesian coordinate version of the data is plotted. (c) Attention is brought to the interesting shape representing the effects of strong anharmonic forces in the xp_xp_y -plot by making selections other than the red selection invisible.

particles undergoing influence by multiple electromagnetic fields. Brushing to create user-defined classifications of particles allow physicists to explore a multitude of projections and representations of the high-dimensional data, thereby promoting scientific discovery. While few of the features in our application are patently new, the combination and usability of features results in a highly efficacious system. PPaint is currently being used in day-to-day operations as part of an accelerator modeling and design project, and has proven effective as a visual data analysis tool.

We envision several avenues for enhancement in our system as well as many important further applications of our work. Additional rendering styles, such as star coordinates [Kan01] and parallel coordinates [Ins85], offer the possibility of increasing the effectiveness of PPaint. We are currently developing and evaluating the effectiveness of dimension-reduction schemes, such as principal components analysis (PCA) [Jol86], to improve our visualization tool. Adding more visual and contextual cues, such as the geometry representing the particle accelerator, will aid in data understanding. From an application point of view, we see great promise for the direct coupling of PPaint to the particle simulation code itself, such that particle data can be visualized as it is being generated, and user interactions through our system can steer subsequent simulator computations.

Acknowledgments

This work was supported by the National Science Foundation under contracts ACR 9982251 and ACR 0222909, through the National Partnership for Advanced Computing Infrastructure (NPACI); the Lawrence Livermore National

Laboratory under contract B523818, by Lawrence Berkeley National Laboratory (LBNL), and the Director, Office of Science, of the U. S. Department of Energy under Contract No. DE-AC03-76SF00098. We thank Andreas Adelmann of LBNL for the use of his time-varying beam dynamics simulation data. We thank the members of the Visualization and Graphics Group of the Institute for Data Analysis and Visualization (IDAV) at UC Davis.

The submitted manuscript has been authored by a contractor of the U.S. Government under contract No. DE-AC03-76SF00098. Accordingly, the U.S. Government retains a nonexclusive royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

References

- [A*03] AGOSTINELLI S., ET AL.: Geant4: A simulation toolkit. *Nucl. Instrum. Meth. A506* (2003), 250–303. 2
- [BC87] BECKER R. A., CLEVELAND W. S.: Brushing scatterplots. *Technometrics* 29, 2 (1987), 127–142. 2
- [CBH*03] CELATA C. M., BIENIOSEK F. M., HENESTROZA E., KWAN J., LEE E. P., LOGAN G., PROST L., SEIDL P. A., VAY J.-L., WALDRON W. L., YU S. S., BARNARD J. J., CALLAHAN D. A., COHEN R., FRIEDMAN A., GROTE D. P., LUND S. M., MOLVICK A., SHARP W. M., WESTENSKOW G., DAVIDSON R. C., EFTHIMION P., GILSON E., GRISHAM L. R., KAGANOVICH I., QIN H., STARTSEV E. A., BERNAL S., CUI Y., FELDMAN D., GODLOVE T. F., HABER I., HARRIS J., KISHEK R. A., LI H., O’SHEA P. G.,

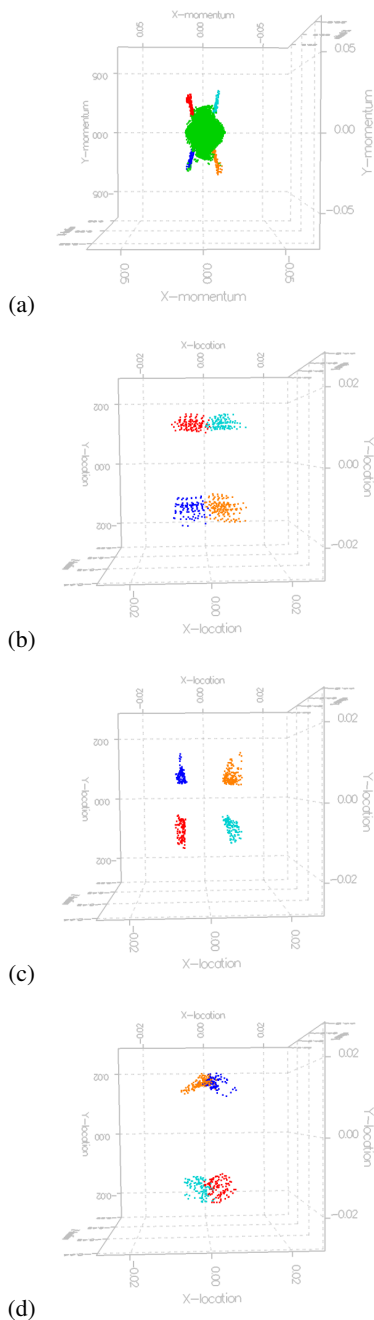


Figure 5: Tracking particle motion over time in a 500 time step beam injector simulation. (a) The $p_x p_y p_z$ -plot of the full distribution in time step 500, at the end of the simulation, where four jets of particles are clearly visible and have been selected in four distinct groups, each with its own color. Figures (b), (c), and (d) show xyz -plots of time steps 90, 370, and 500, respectively, of the four selected groups, the core of the beam being rendered invisible. Note that each group remains in a tight formation even though heavy particle mixing is occurring, but is spread over multiple beamlets.

QUINN B., REISER M., VALFELLS A., WALTER M., ZOU Y., ROSE D. V., WELCH D. R.: Progress in heavy ion fusion research. *Physics of Plasmas* 10, 5 (May 2003), 2064–2070. 1

[DGH03] DOLEISCH H., GASSER M., HAUSER H.: Interactive Feature Specification for Focus+Context Visualization of Complex Simulation Data. In *Proceedings of the symposium on Data visualisation 2003* (2003), Hansen C. D., Bonneau G.-P., Hahmann S., (Eds.), Eurographics Association, pp. 239–248. 2, 3

[Fri02] FRIEDMAN A.: Overview of Heavy Ion Fusion Accelerator Research in the U.S. *Proceedings of the 10th Workshop on Advanced Accelerator Concepts AIP Conf. Proc. 647* (June 22–28 2002), 106–120. 1

[HIF] Heavy-Ion Fusion Virtual National Laboratory. URL: <http://hif.lbl.gov>. 1

[Ins85] INSELBERG A.: The plane with parallel coordinates. *The Visual Computer* 1, 2 (Oct. 1985), 69–92. 6

[Jol86] JOLLIFFE I.: *Principal Component Analysis*. Springer-Verlag, New York, NY, 1986. 6

[Kan01] KANDOGAN E.: Visualizing Multi-dimensional Clusters, Trends, and Outliers using Star Coordinates. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-01)* (New York, Aug. 26–29 2001), Provost F., Srikant R., (Eds.), ACM Press, pp. 107–116. 6

[KMC*03] KREYLOS O., MAX N., CRIVELLI S. N., BETHEL W., HAMANN B.: Interactive protein manipulation. IEEE. URL: <http://vis.lbl.gov/Research/ProteinShop/index.html>. 3

[MQR99] MCCORMICK P. S., QIANG J., RYNE R. D.: Visualization viewpoints: Visualizing high-resolution accelerator physics. *IEEE Computer Graphics and Applications* 19, 5 (Sept./Oct. 1999), 11–13. 1, 2

[MSW*02] MA K.-L., SCHUSSMAN G., WILSON B., KO K., QIANG J., RYNE R.: Advanced Visualization Technology for Terascale Particle Accelerator Simulations. In *SC'2002 Conference CD* (Baltimore, MD, Nov. 2002), IEEE/ACM SIGARCH. 2

[MW95] MARTIN A. R., WARD M. O.: High dimensional brushing for interactive exploration of multivariate data. In *VIS '95: Proceedings of the 6th conference on Visualization '95* (1995), IEEE Computer Society, p. 271. 2

[PAW] Physics Analysis Workstation (PAW). URL: <http://paw.web.cern.ch/paw/>. 2

[PvBZG00] PFISTER H., VAN BAAR J., ZWICKER M., GROSS M.: Surfels: Surface elements as rendering primitives. In *Proceedings of the Computer Graphics Conference 2000 (SIGGRAPH-00)* (New York, July 23–28 2000), Hoffmeyer S., (Ed.), ACM Press, pp. 335–342. 4

- [TLM03] TZENG F.-Y., LUM E., MA K.-L.: A Novel Interface for Higher-Dimensional Classification of Volume Data. In *Proceedings of IEEE Visualization 2003* (Seattle, WA, Oct.19–24 2003), Turk G., Wijk J., Moorhead R., (Eds.), IEEE, pp. 505–512. 3
- [TM03] TEOH S. T., MA K.-L.: PaintingClass: Interactive Construction, Visualization and Exploration of Decision Trees. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD-03)* (New York, Aug. 24–27 2003), Domingos P., Faloutsos C., Senator T., Kargupta H., Getoor L., (Eds.), ACM Press, pp. 667–672. 3
- [Tuf90] TUFTE E. R.: *Envisioning Information*. Graphics Press, Cheshire, Connecticut, 1990. 2
- [WF96] WARE C., FRANCK G.: Evaluating stereo and motion cues for visualizing information nets in three dimensions. *ACM Transactions on Graphics* 15, 2 (Apr. 1996), 121–140. 4
- [WMM02] WILSON B., MA K.-L., MCCORMICK P. S.: A Hardware-Assisted Hybrid Rendering Technique for Interactive Volume Visualization. In *Proceedings of the 2002 IEEE Symposium on Volume Visualization and Graphics (VOLVIS-02)* (Piscataway, NJ, Oct. 28–29 2002), Spencer S. N., (Ed.), IEEE, pp. 123–130. 2
- [WMQR02] WILSON B., MA K.-L., QIANG J., RYNE R.: Interactive Visualization of Particle Beams for Accelerator Design. In *Workshop on High Performance Computing in Particle Accelerator Science and Technology* (2002), 2002 International Conference on Computational Science. 2