

CEDAR: tools for event generator tuning

Andy Buckley

Institute for Particle Physics Phenomenology, Durham University, UK.

E-mail: andy.buckley@durham.ac.uk

I describe the work of the CEDAR collaboration in developing tools for tuning and validating Monte Carlo event generator programs. The core CEDAR task is to interface the Durham HepData database of experimental measurements to event generator validation tools such as the UCL JetWeb system — this has necessitated the migration of HepData to a new relational database system and a Java-based interaction model. The “number crunching” part of JetWeb is also being upgraded, from the FORTRAN HZTool library to the new C++ Rivet system and a generator interfacing layer named RivetGun. Finally, I describe how Rivet is already being used as a central part of a new generator tuning system, and summarise two other CEDAR activities, HepML and HepForge.

*XI International Workshop on Advanced Computing and Analysis Techniques in Physics Research
April 23-27 2007
Amsterdam, the Netherlands*

1. Introduction

Monte Carlo event generators are an essential tool for particle physics, simulating aspects of collider events ranging from the parton-level signal process to cascades of QCD and QED radiation in both initial and final states, non-perturbative hadronisation, underlying event physics and specific particle decays. Event generators provide experimentalists and phenomenologists with samples of fully exclusive events drawn from physical distributions, and are therefore central to the design of both detector hardware and data analysis strategies: this is more true than ever for LHC physics.

However, event generators are not fully predictive: various phenomenological parameters must be tuned to experimental data to bootstrap a general purpose generator before physically meaningful predictions can be obtained. Such parameters include the parton density functions (PDFs) of the colliding beam particles, parton shower cutoffs and evolution variables, the running of α_s , choice of Λ_{QCD} , and a variety of hadronisation parameters which strongly depend on the hadronisation model being applied. Observable distributions calculated from a generator's output depend on these parameters in an extremely non-trivial way. This leads to the generator tuning problem — how to choose the parameter sets that give the best fits to experimental data, given that we have no ready parameterisation of the output, and an exhaustive exploration of the many-dimensional parameter space is out of the question?

In this paper, I will describe the work done by the CEDAR collaboration [1, 2] towards addressing this problem for the LHC. The main theme of this work is the development of tools to generate and efficiently analyse simulated events, and the validation of generator tunings against experimental data using these tools.

2. HepData

HepData [3] is a database of experimental particle physics data, which has been maintained at Durham University since the 1970s. HepData's contents are not the raw experimental data, but the data as presented in the plots and data tables of peer-reviewed experimental papers. HepData contains data from a wide range of collider and fixed target experiments, covering many initial states and \sqrt{s} : it is therefore an excellent reference point for the distributions of observables that event generators should be tuned against.

2.1 The legacy database

Since its inception, HepData has been based on a hierarchical database management system (HDBMS). In the intervening years, the database world has firmly centred its attention on more flexible relational database systems (RDBMS), and today a wide range of high quality RDBMS systems are available for free, with strong support for the SQL query language and networked availability. By contrast, the hierarchical systems have evolved little or not at all, and even simple operations like query changes or schema updates are substantial tasks involving writing FORTRAN routines.

These shortcomings of the legacy database meant that, with the incentive of placing HepData as a data service at the centre of projects like CEDAR, the decision was taken to upgrade HepData to use a relational database system [4, 5].

2.2 The new database

HepData's data model is intrinsically hierarchical: a given data point is located within the hierarchy *paper* \rightarrow *dataset* \rightarrow *axis* \rightarrow *point* and there is little point in comparing data points from different distributions. Fortunately, a relational model is flexible enough that implementing a hierarchical structure is easy.

In practice the RDBMS used is MySQL [6], but due to the ANSI-standardised SQL query language this is easily swappable for almost any other modern RDBMS. A major design principle has been the decoupling of database access and implementation details, such as which parts of a data record are stored in which fields and tables, from the semantic model of how various aspects of the data are related to each other. Hence, the semantics of the data are reflected in a Java "object model", which makes no reference to database implementation, and the Hibernate [7] object-RDBMS persistency system is used as a layer to isolate "client" code from the database implementation details, using Java mechanisms such as reflection and annotations. With Hibernate, the persistency between the persistent database storage and the ephemeral in-memory object representation is governed by configuration files and Java metadata, and so much tedious and fragile database access coding is avoided.

On top of this object model, the primary mode of access to the database will remain the Web interface. This is also written in terms of Java objects, executed in the Apache Tomcat [8] servlet engine. More decoupling is desirable here, and we intend to use the Tapestry [9] templating system to isolate the data content from the Web presentation according to the model-view-controller (MVC) design idiom, but as yet there has been little effort expended on the new Web interface.

While the Web interface, when completed, will be the primary method of searching and browsing the database, our design means that the object model can also be imported from, and exported to, various other representations. Datasets will be representable as graphics (in a variety of formats), plain text, the AIDA [10] XML and ROOT data formats and HepML [11]. The last of these is the canonical file format representation of HepData records and mechanisms for data import and export in HepML format are built into HepData, using the Castor [12] object-XML marshalling system. Another form of HepML is used by JetWeb and will be mentioned later.

2.3 Database migration

Migrating from the legacy hierarchical database to a new relational database has proven to be a substantial task. In large part the difficulties have been due to the relative unstructured form of the legacy database, which has no strong type system: data entries are simply text strings. This means that the structure and integrity of the data has been subject to the whims of various data submitters over a long time period, tempered by the diligence of the database managers, and there are many cases where a data record which appears attractively formatted when presented as a Web page turns out to be hard to fit into a more rigidly defined data model. To attempt to decouple the target relational database design from the vagaries of the legacy system, the migration of data has become a multi-stage process.

Legacy DBMS to flat files The first stage is to use a mixture of FORTRAN and Perl scripts to massage the hierarchical data records into a set of tab-delimited plain text files, each of which contains all the information for one aspect of all the papers in the database. As might be expected, the data

point values and errors files are extremely large. This procedure needs only be done infrequently, if the legacy database changes substantially, and so an immediate decoupling is achieved.

Flat files to HepML The second stage is to transform these “flat files” into a series of HepML files, one for each paper, using a Python script. In practice, we have found that the efficiency of the HepML builder can be greatly increased by splitting each of the flat files so that there is one file per legacy paper, and then sorting the content of these files. This reduces the number of scans required through large files, and allows caching to be made use of during the building of the XML object model. There is not an exact match between the papers from the legacy database and those which are available at the end of the HepML-building procedure because, due to technical limitations, long papers often had to be split into several parts in the legacy system: the “logical” papers which result from the HepML builder are the preferable form.

HepML to RDBMS The final stage of the HepData migration is the reading of HepML files into the database. While the previous two migration steps are only of use in the migration process, importing records from HepML is a core feature of the HepData system. The mechanism used is an object persistency framework, Castor, of which we are only using the XML part. Like Hibernate, Castor describes how the object model will be stored in a persistent form, but in this case the persistency medium is an XML string rather than database tables. As some features of the HepML representation are not strictly hierarchical, the JDOM [13] Java XML processing library is used as an input filter when importing records from HepML files.

3. JetWeb

The second main component of CEDAR is JetWeb [14] an application for validating the performance of various generator tunings. JetWeb combines a system for running a variety of event generator programs, a database of distributions calculated from simulated events and a Web interface written in Java and run on a Tomcat Java server. JetWeb’s development was motivated by the need to avoid misleading tunings, where one distribution is fitted at the expense of unseen others: accordingly, tunings considered by JetWeb will be compared to as many distributions as possible.

Consistently generating data for such a large number of distributions requires both a good understanding of the physics models involved and a lot of computational power: JetWeb helps here by providing a relatively user-friendly way of configuring the models, by archiving generated data in such a way that extra statistics can be requested through the Web interface, and providing a browsable archive of stored results. JetWeb shows the overall χ^2 for a chosen model against all distributions, as well as the fit quality to individual plots, so the overall quality of a given tuning can be readily assessed.

JetWeb was initially developed at UCL, based on analyses using the HZTool [15–17] library and various versions of the Herwig [18] and Pythia [19] FORTRAN event generators. The reference data in this version was transcribed from a variety of sources, including HepData. Extension of JetWeb’s prototype generator interface to deal with generators other than Pythia and Herwig proved difficult, and hence CEDAR’s work on JetWeb has centred on improving the way that generators are modelled, adding mechanisms for combining generator runs, and separating run parameters from model parameters[4]. The Web user interface has also been considerably enhanced, and a version of

HepML for describing generator and run configurations has been developed and incorporated into the JetWeb system. The other way in which JetWeb’s modernisation shows is that event generation now uses Grid resources and authentication rather than local batch farms.

Another JetWeb change planned under CEDAR is the linking of JetWeb to HepData— a “single-sourcing” approach which should result in more robust data–MC comparisons and easier adding of new analyses. Since the HepData migration process has been more lengthy than anticipated, it is only recently that JetWeb has begun reading some data directly from the HepData database (via the HepData Java object model). Before JetWeb’s internal database of experimental reference data can be eliminated, some extra datasets must be added to HepData and the HepData migration must be essentially complete so that data entries are reliably retrievable.

JetWeb is a relatively sophisticated tool for steering event generators based on “high level” event type requests, and for comparing generated data to reference plots, but it does not actually interface to the various generator codes or analyse the generated events directly. These rôles are filled by a pair of native applications for event analysis and steering — in the existing incarnations of JetWeb the FORTRAN programs HZTool [15] and HZSteer [20] are used, but CEDAR has developed C++ replacements for these, titled Rivet [21, 22] and RivetGun [23] respectively.

4. Rivet and RivetGun

Rivet is a C++ replacement for the FORTRAN HZTool library, initially developed for the HERA experiments *H1* and *ZEUS* (hence the “HZ”). HZTool as a library has two rôles: firstly, to provide a collection of physics utility routines for calculating commonly used quantities, such as implementations of jet definitions; and second, to collect a set of analyses which use these routines and produce histograms comparable with experimental results. As a result of its HERA legacy, the majority of HZTool analyses are from DIS and photoproduction experiments.

Initially, HZTool analyses included code to specifically configure particular generators. However, this approach scales badly, and a concerted effort was made in 2005 to decouple HZTool routines from generator specifics. The result was a steering package, HZSteer [20], which contains (almost) all of the generator-specific code, and the current version of HZTool is purely concerned with the physics analyses of event records, and not where any given event came from. The current version of JetWeb uses HZTool and HZSteer to generate and analyse events.

Even as HZSteer was being split off from HZTool, it was clear that time was running short for FORTRAN-based analysis systems. The rising prominence of C++-based generators, such as Herwig++ [25, 26], SHERPA [27] and Pythia 7/8 [29, 30], was evident, and FORTRAN does not have the level of sophistication as an application framework language to steer these generators¹. A not unimportant secondary point is that the success of a system like HZTool relies on the support of the community in providing new analyses to keep pace with the appearance of new data: the *de facto* language used by LHC-era experimentalists is C++, and an analysis system written in any other language is less likely to be embraced. The result is a new, C++-based analysis system, Rivet, to replace HZTool, and a generator steering system, RivetGun, to replace HZSteer.

¹Indeed, technical concerns with how C++ encodes symbol names mean that steering C++ from anything other than C++ is troublesome.

4.1 Rivet

Rivet [21, 22] (an acronym for “Robust Independent Validation of Experiment and Theory”) is a generator-agnostic analysis framework. Guiding design principles include the implementation in object oriented C++ and compatibility with existing standard data formats such as HepMC [24] and AIDA [10]. Rivet has no dependence on generator-specific features and sees data only via HepMC event records either supplied from file or a generator steering package such as RivetGun (see below). This makes it easier to incorporate new Monte Carlo generators into a Rivet-based validation system than is currently the case with HZTool.

The Rivet analysis system is based on a concept of “event projections”, which project a simulated event into a lower-dimensional quantity such as scalar or tensor event shape variables. Projections can be nested and their results are automatically cached to eliminate duplicate computations, using C++ runtime type information (RTTI) and comparison operators between projection classes. The infrastructure has been designed to place as little burden as possible on the authors of projection and analysis classes, which should be concerned almost entirely with the analysis algorithm.

Sets of standard projections and analyses are included with the Rivet package, and this collection will grow with subsequent releases. Analysis data is accumulated using the AIDA interfaces, and exported primarily in the AIDA XML histogram format. If ROOT is present on the build system, ROOT format files can also be exported, allowing use of Rivet for n-tuple based analyses as well as the primary design purpose of semi-automated event generator validation. To complement the generated analysis data, HepData-generated AIDA records for each bundled analysis are included in the Rivet package and can be used to define the binnings of generated data observables: this improves the robustness of analysis implementations and allows easy data-theory comparisons without requiring network access to HepData.

At the time of writing, the stable version of Rivet is 0.9, available from the Rivet development website [22]. This first version includes 5 analyses — two from Tevatron Run 2, one from LEP and two from HERA — as well as the library of projections which currently includes e^+e^- event shapes, DIS kinematical boosts, the DØ “improved legacy cone” and k_{\perp} jet algorithms via KtJet [35, 36], a variety of final state projections including particle vetoing, and several others. With the main Rivet design now stable, we intend for the next release to have much more substantial libraries of both analyses and projections, such that Rivet can entirely replace HZTool.

4.2 RivetGun

Rivet is primarily a code library for use by generator steering packages, although it also includes a command line tool, `rivet`, which can read in HepMC “ASCII” event files. The main tool for running Rivet is the RivetGun generator steering program. RivetGun is written in C++ and provides a uniform programmatic and command line interface to running event generators, with common generator configuration features such as setting initial states, named control parameters and random seeds possible through the most general level of the interface. Using runtime dynamic library loading, even different versions of the same generator can be used from the same executable, which would not be possible with compile time library linking.

The RivetGun C++ class structure uses inheritance to define a base class, `Generator` on which these operations are declared — each specific generator then implements the methods de-

clared in the interface in its own way, be that by mapping common blocks, calling FORTRAN library routines or using the C++ class methods of the target generator. With this approach, switching between generators with the same initial state configuration is trivial, although knowledge of generator-specific parameters is still needed for any proper study. Formally, the generator interfaces are part of a library called AGILe (“A Generator Interface Library”), since we wish to keep open the possibility of using the interface without using Rivet at all.

RivetGun currently provides generator interfaces for FORTRAN Herwig [18] and Pythia [19], plus enhanced versions of those generators using the AlpGen [31], Charybdis [32] and Jimmy [33,34] auxiliary generators. Preliminary bindings to the Herwig++ [25,26], SHERPA [27,28] and Pythia 8 [30] generators are also available. RivetGun has not yet been formally released, but the development code is sufficiently usable that it is by far the easiest way to generate distributions using Rivet.

Since RivetGun is intended to be run from within JetWeb, the generator configuration will eventually be able to be set from HepML generator description files, as well as the current methods of command line arguments and/or simple key–value parameter files.

5. Tuning vs. validation

So far we have addressed efforts towards CEDAR’s central goal, which is the *validation* of existing event generator tunings. An obvious criticism of this is that nowhere in the framework is there a procedure for finding a better or, ideally, optimal tuning. Let us consider the general problem before describing a particular, CEDAR-centric solution.

5.1 The tuning problem

Naïvely, one might expect that an event generator can be optimally tuned by either grid-scanning the parameter space, evaluating a goodness of fit (GoF) measure against reference data and choosing the best point. Anyone with experience in sampling problems will be aware of the fallacies at work here: at the root of the difficulties is the exponential scaling, $\mathcal{O}(A^n)$, of computational requirements with the dimensionality n of the parameter space. This makes comprehensive scanning unrealistic for typical tuning problems with $n \sim \mathcal{O}(10)$. Even adaptive grid-scanning, where the grid is non-uniform, or adapts to the local GoF, is subject to the exponential scaling.

Sampling specialists may suggest a Markov chain Monte Carlo (MCMC) approach to this problem. MC sampling works because the scaling is independent of n , although there are only rules of thumb available for estimating sampler convergence rates and the choice of proposal distribution and use of gradient information can have very significant effects. However, the typical burn-in for an MCMC sampler is likely to be in the hundreds or thousands of iterations even with a well-chosen MCMC setup: this is fine for easily evaluated functions, but the “function” we are attempting to minimise (the GoF for, say 100,000 events generated with the proposed parameter vector) renders this approach on the edge of current computational feasibility.

There is a second reason to be wary of such approaches, though: a global GoF measure is not sensitive enough for a non-exhaustive search such as MCMC to find a global optimum. There are simply too many ways for parts of distributions to fit better or worse than others and the combinatorics generate a parameter space vastly dominated by mediocre tunings — what is needed is

a method which is more sensitive to the dependence of each element of each distribution on the tuning parameters. Such a method was used by the LEP Delphi experiment [37, 38], and it is this approach which we will now describe.

5.2 Professor — tuning with bin-by-bin interpolation

The Delphi approach to event generator tuning was to fit a function to the generator output on a bin-by-bin basis and then to minimise the goodness of fit in all bins simultaneously, using the interpolating functions. While previous approaches fitted a linear function, Delphi fitted a second order polynomial since that is the first order at which inter-parameter correlations are taken account of [38].

Delphi’s FORTRAN code implementing this algorithm was called Professor, and so is its continuation, although the implementation language is now Python combined with the C++ Rivet and RivetGun systems. Professor is not a CEDAR project — it is a collaborative effort between the Durham IPPP and Dresden MPI — but its aims are so closely connected to CEDAR that it seems prudent to mention it here.

The procedure implemented by Professor is as follows:

1. Define a hypercube in the n -dimensional parameter space, by specifying sample ranges in each parameter.
2. Generate N random parameter n -vectors in the hypercube. There is no upper limit on the number of samples — indeed, the more the better as might be expected, — but there is a minimum number, given later.
3. Run RivetGun and Rivet using the sampled parameter values: this will produce N Rivet output AIDA files, each of which, say, describes B bins. The number of events generated in each run should be sufficient to reduce statistical error to a near-negligible level.
4. For each bin, b , fit a polynomial function to the N generated values, using a singular value decomposition (SVD) [39]. The SVD allows calculation of a “pseudoinverse” — a matrix inverse for non-square matrices [40, 41] whose use in overconstrained systems performs a least-squares fit [42]. The function to be fitted is the general second-order polynomial in n dimensions,

$$f_b(\mathbf{p}) = \alpha_b + \sum_i \beta_{b,i} \mathbf{p}'_i + \sum_{i,j \geq i} \gamma_{b,ij} \mathbf{p}'_i \mathbf{p}'_j. \quad (5.1)$$

where $\mathbf{p}' = \mathbf{p} - \mathbf{p}_0$, the shift of the parameters from their nominal/central values \mathbf{p}_0 and α_b , $\beta_{b,i}$ & $\gamma_{b,ij}$ are the sets of polynomial coefficients to be determined for each bin b . Coefficient counting reveals that there must be $N \geq 1 + (n^2 + 3n)/2$ parameter space samples for there to be an inverse. In principle, \mathbf{p}_0 itself can be considered as a set of parameters to be fitted, which would add another n to the minimum N .

5. Use reference data from HepData to compute a goodness of fit function for each bin, such as error-weighted square deviation, $\phi_b(\mathbf{p}) = (f_b(\mathbf{p}) - r_b)^2 / E_b^2$ where r_b and E_b are the experimental value and uncertainty respectively.

6. Analytically or numerically minimise the individual ϕ_b functions and the corresponding global GoF figure, the ubiquitous χ^2 , defined² as $\chi^2(\mathbf{p}) = \sum_b \phi_b(\mathbf{p})$. Flag any significant deviations of the bin-by-bin minima from the global minimum.
7. Generate a final MC event sample using the interpolation-optimal parameters and compare with the prediction.

Professor is currently in active development, testing the method with toy models and low-dimensional samplings using Rivet and RivetGun. The first uses of it will be on relatively simple cases such as re-implementing the original Delphi optimisation, and then proceeding to more complex tunings for the LHC where data from various generators must be combined and extrapolated. There is a great opportunity for statistical sophistication in this area, including bin–bin correlations, various GoF measures [43] and weighting of particular distributions and bins.

6. HepForge

As a spin-off from our own development requirements, CEDAR now provides a free online collaborative development facility, HepForge [44, 45], for HEP projects which aim to provide useful, well-engineered tools to the community.

HepForge currently offers feature-enhanced Web hosting, hosting and HTTP access to the Subversion code management system (a modern replacement for CVS), an integrated bug tracker and wiki system strongly integrated with Subversion and mailing lists for developer contact, project announcements and discussion. All the CEDAR projects and about 20 others are hosted with HepForge and it has proven a popular alternative to CERN’s Savannah system, particularly for small phenomenology collaborations.

The long-term plan for HepForge is that it will provide search facilities for a wide range of HEP computational tools, but at present we are consolidating our developer support and improving the existing system.

7. Summary

CEDAR is providing a wide variety of computational tools, which are the wide foundation on which a systematic and global validation and tuning of HEP event generators is to be built. The largest-scale components of CEDAR are HepData and JetWeb— established resources which have been substantially re-designed and upgraded by CEDAR. HepData’s migration from a legacy hierarchical database to a much more rigorously structured relational database and Java object model and persistency system has been a substantial task and is now approaching its final stages. JetWeb has been internally restructured a great deal, but the most obvious consequence of the CEDAR upgrades is the forthcoming use of HepData as a source of reference data. The HepML XML formats provide the glue between these systems, and a file persistency format.

At a finer-grained level, the Rivet and RivetGun systems are the C++ replacements for HZTool and HZSteer being created by CEDAR. The first official release of Rivet has recently taken place

²We’re being somewhat sloppy about the definition of the error in this χ^2 definition: strictly it should be the “theory error” to avoid a biased distribution.

and work is now proceeding on adding new analyses to it and preparing RivetGun for its first stable release. At Rivet's core is the concept of event projections and generator independence — with these and a design aim to make it as easy as possible to write algorithm-focused analysis code, Rivet is an excellent framework for LHC validation and tuning studies and users are encouraged to try it out.

Finally, mention was made of the Professor event generator tuning effort and of the HepForge development environment. These will respectively build on the foundation provided by CEDAR and continue to provide facilities for the development of HEP computational tools.

8. Acknowledgements

The CEDAR team would like to thank the UK Science and Technology Funding Council (STFC) for their generous support. CEDAR's work has also been supported in part by a Marie Curie Research Training Network of the European Community's Sixth Framework Programme under contract number MRTN-CT-2006-035606.

References

- [1] J. M. Butterworth *et al.*, hep-ph/0412139, presented at CHEP'04, Interlaken, September 2004
- [2] CEDAR Web site: <http://www.cedar.ac.uk>
- [3] Legacy HepData Web interface: <http://durpdg.dur.ac.uk/hepdata/>
- [4] A. Buckley *et al.*, "HepData and JetWeb" in *Proceedings of CHEP'06*, arXiv:hep-ph/0605048v1
- [5] HepData development: <http://projects.hepforge.org/hepdata/>
- [6] MySQL: <http://www.mysql.com>
- [7] Hibernate: <http://www.hibernate.org>
- [8] Apache Tomcat: <http://tomcat.apache.org>
- [9] Tapestry: <http://tapestry.apache.org>
- [10] Abstract Interfaces for Data Analysis (AIDA), <http://aida.freehep.org/>
- [11] HepML: <http://projects.hepforge.org/hepml/>
- [12] Castor: <http://www.castor.org>
- [13] JDOM: <http://www.jdom.org>
- [14] J. M. Butterworth and S. Butterworth, *Comput. Phys. Commun.* **153** (2003) 164, <http://jetweb.cedar.ac.uk/>
- [15] J. Bromley *et al.*, "HZTOOL: A package for Monte Carlo-data comparison at HERA (version 1.0)," in *Workshop on Future Physics at HERA, 1996*
- [16] J. M. Butterworth, H. Jung, V. Lendermann, B. M. Waugh, in *HERA and the LHC: A Workshop on the implications of HERA for LHC physics*, Proceedings, Part B., hep-ph/0601013
- [17] HZTool: <http://hepforge.cedar.ac.uk/hztool/>
- [18] G. Corcella *et al.*, "HERWIG 6.5 release note," arXiv:hep-ph/0210213

- [19] T. Sjöstrand, S. Mrenna and P. Skands, *JHEP* **0605** (2006) 026, arXiv:hep-ph/0603175
- [20] HZSteer: <http://hepforge.cedar.ac.uk/hzsteer/>
- [21] B. M. Waugh *et al.*, “HZTool and Rivet” in *Proceedings of CHEP’06*, arXiv:hep-ph/0605034v1
- [22] Rivet: <http://projects.hepforge.org/rivet/>
- [23] RivetGun: <http://projects.hepforge.org/rivetgun/>
- [24] M. Dobbs and J. B. Hansen, *Comput. Phys. Commun.* **134** (2001) 41, <https://savannah.cern.ch/projects/hepmc/>
- [25] S. Gieseke *et al.*, “Herwig++ 2.0 release note,” arXiv:hep-ph/0609306
- [26] Herwig++: <http://projects.hepforge.org/herwig/>
- [27] T. Gleisberg, S. Hoche, F. Krauss, A. Schalicke, S. Schumann and J. C. Winter, “SHERPA 1.alpha, a proof-of-concept version,” *JHEP* **0402** (2004) 056, arXiv:hep-ph/0311263
- [28] <http://projects.hepforge.org/sherpa/>
- [29] L. Lönnblad, “ThePEG, PYTHIA7 and ARIADNE,” prepared for *12th International Workshop on Deep Inelastic Scattering (DIS 2004)*
- [30] Pythia 8: <http://www.thep.lu.se/~torbjorn/pythiaaux/future.html>
- [31] M. L. Mangano, M. Moretti, F. Piccinini, R. Pittau and A. D. Polosa, *JHEP* **0307** (2003) 001, arXiv:hep-ph/0206293
- [32] C. M. Harris, P. Richardson and B. R. Webber, *JHEP* **0308** (2003) 033, arXiv:hep-ph/0307305
- [33] J. M. Butterworth, J. R. Forshaw and M. H. Seymour, “Multiparton interactions in photoproduction at HERA,” *Z. Phys. C* **72** (1996) 637, arXiv:hep-ph/9601371
- [34] Jimmy: <http://hepforge.cedar.ac.uk/jimmy/>
- [35] J. M. Butterworth, J. P. Couchman, B. E. Cox and B. M. Waugh, “KtJet: A C++ implementation of the k_{\perp} clustering algorithm,” *Comput. Phys. Commun.* **153** (2003) 85, arXiv:hep-ph/0210022
- [36] KtJet: <http://hepforge.cedar.ac.uk/ktjet/>
- [37] P. Abreu *et al.*, DELPHI Collaboration, *Z. Phys. C* **73** (1996) 11.
- [38] K. Hamacher and M. Weierstall, “The Next Round of Hadronic Generator Tuning Heavily Based on Identified Particle Data,” arXiv:hep-ex/9511011
- [39] H. Abdi, *Singular Value Decomposition (SVD) and Generalized Singular Value Decomposition (GSVD)*, in N. J. Salkind (Ed.), *Encyclopedia of Measurement and Statistics*, Thousand Oaks, 2007
- [40] E. H. Moore, “On the reciprocal of the general algebraic matrix,” in *Bulletin of the American Mathematical Society* **26** (1920) 394–395
- [41] R. Penrose, “A generalized inverse for matrices,” in *Proceedings of the Cambridge Philosophical Society* **51** (1955) 406–413
- [42] R. Penrose, “On best approximate solution of linear matrix equations,” in *Proceedings of the Cambridge Philosophical Society* **52** (1956) 17–19
- [43] S. Donadio, S. Guatelli, B. Mascialino, M. G. Pia, A. Pfeiffer, A. Ribon and P. Viarengo, “A statistical toolkit for data analysis,” *Nucl. Phys. Proc. Suppl.* **150** (2006) 50
- [44] A. Buckley *et al.*, “HepForge: A lightweight development environment for HEP software” in *Proceedings of CHEP’06*, arXiv:hep-ph/0605048v1
- [45] HepForge: <http://www.hepforge.org>