

# Towards enhanced databases for High Energy Physics

---

**Andrea Ceccarelli**

*Dipartimento di Matematica e Informatica, Università degli Studi di Firenze,*  
E-mail: [andrea.ceccarelli@unifi.it](mailto:andrea.ceccarelli@unifi.it)

**Andrea Cioni**

*Dipartimento di Matematica e Informatica, Università degli Studi di Firenze,*  
E-mail: [andrea.cioni@stud.unifi.it](mailto:andrea.cioni@stud.unifi.it)

**Maria Vittoria Garzelli\***

*Dipartimento di Fisica e Astronomia, Università degli Studi di Firenze, & INFN, Firenze,*  
*Institut für Theoretische Physik, Eberhard Karls Universität Tübingen,*  
E-mail: [maria.vittoria.garzelli@cern.ch](mailto:maria.vittoria.garzelli@cern.ch)

**Piergiulio Lenzi**

*Dipartimento di Fisica e Astronomia, Università degli Studi di Firenze, & INFN, Firenze,*  
E-mail: [lenzip@fi.infn.it](mailto:lenzip@fi.infn.it)

**Laura Redapi**

*GARR Consortium,*  
*Dipartimento di Fisica e Astronomia, Università degli Studi di Firenze,*  
E-mail: [laura.redapi@gmail.com](mailto:laura.redapi@gmail.com)

The accumulation of a large amount of new experimental data at an impressive rate at present and future collider experiments has led to important questions concerning data storage and organization, their public access and usability, as well as their efficient usage in order to discriminate between different theories. For the last forty years, the HEPData database has been the reference database for the worldwide community of elementary particle physicists, from DIS to fixed-target and collider experts. Using as a basis a dump of HEPData,<sup>†</sup> we discuss possible paths to enhance the capabilities of databases for High Energy Physics. Our starting point is the reorganization of the data in a different scheme, which allows for the application of OLAP techniques to automatically extract information at a multidimensional level, answering to complex queries. The feedback of the DIS community is important for understanding specific needs, aiming at a more effective storage, extraction and presentation of the data and information of their interest.

*XXVII International Workshop on Deep-Inelastic Scattering and Related Subjects - DIS2019*

*8-12 April, 2019*

*Torino, Italy*

---

\*Speaker.

<sup>†</sup>the HEPdata dump on which this work is based was kindly provided us by K. Ellis, F. Krauss and G. Watt (IPPP, Durham), with whom further collaboration is foreseen for the future.

## 1. Introduction

Physics experiments are accumulating data at an impressive rate, and questions about the most efficient ways of storing, accessing and preserving these data in view of their future re-usage and of the possibility of performing consistent comparisons of the results of different experiments, are open. In this respect, the Group of scientists involved in the discussion of Data Preservation in High Energy Physics (DPHEP), considering the flow of information in a typical experimental analysis at an high-energy accelerator, has suggested a classification of data into four different levels, characterized by different contents [1]. In particular, going from larger data sizes to lower ones, they distinguish raw data collected during the measurement and reconstruction software (level 4), analysis-level data (level 3), i.e. the data and software sufficient to perform a complete re-analysis, simplified-event-level data (level 2), i.e. sets of vectors describing the particles of each selected event in a simplified format, and published data (level 1), i.e. the data reported in scientific works.

In this work, we focus on level-1 data, which are characterized by the fact of being 1) multisource, corresponding to the outcome of many different experiments, 2) multidimensional and 3) heterogeneous. This means that they are characterized by different dimensions, whose number vary according to the observable/experiment considered and the way of reporting the information, and they are organized in different kinds of data structures. These features represent in general a challenge for the organization of all these data in a common database and their retrieval, also considering that information of the same quality might have been reported by different experiments in different ways (for example, the information concerning correlated uncertainties, crucial for QCD analyses interesting for the DIS community, can be stored in covariant matrices or in other formats) and that information reported in some analyses can be absent in other ones.

The most complete open-access collection of level-1 data on high-energy scattering reactions at experiments conducted from the early days of particle physics to present, is the HEPData database, which can be regarded as the reference database for the worldwide community of elementary particle physicists at accelerators. This database, mostly developed at the Institute for Particle Physics Phenomenology (IPPP) of the University of Durham and, more recently, also in collaboration with the Open Access group at CERN, has a long history, from the sixties to the present. A legacy version [2] was organized according to the hierarchical Berkeley Database Management System (DBMS), with information stored in tree-like structures, with papers (including experimental data) as record units. The user could access the stored information through Fortran routines. This version was superseded by a more modern one, called HepData [3], developed within the CEDAR project [4], with the information reorganized according to the relational DBMS MySQL, in a complex structure of many tables with different kinds of relations (one-to-many, many-to-many). At that time, the HepML data format [5], based on XML, was created to enter data, as well as a web interface in Java for allowing user queries [6]. The parallel development of further tools, like e.g. Rivet and JetWeb in the CEDAR context [7], was driven by the ultimate objective of automatizing the comparison of the collected data with theory predictions as much as possible. Finally, the most recent HEPData version [8] has been developed in the last few years. Differently from previous versions, hosted on Durham servers, the last version [9] is hosted on the CERN OpenStack infrastructure. It makes use of the PostgreSQL DBMS and of the Invenio3 open-source framework for large-scale digital repositories. An interface has been created to allow the

experimental collaborations to fill the database with their data, in files written in `yaml`, a superset of the `json` markup language. The current emphasis is on the inclusion of more data, not only limited to HEP reactions, and on the development of statistical models and of studies of Physics Beyond the Standard Model.

Our work starts from the analysis of the present status of HEPData, and shows how a re-organization of the data into a different structure, inspired by OnLine Analytical Processing (OLAP) techniques [10], allows the user to perform queries, centered on the extraction and comparison of data from different papers, more complex than those allowed by the present HEPData structure.

## 2. The MineHEP search-engine prototype


Our final goal is organizing the already available level-1 HEP data in a way to automatically extract from them as much information as possible. Using as a starting point the MySQL dump file `hepdatapublic.dmp.gz_01Apr18` of HEPData, we have built a search-engine prototype, dubbed MineHEP, capable of queries more complex than those possible nowadays through the HEPData web interface.

In relational databases, data are organized in tables structured in fields, where each field corresponds to a column and each data entry to a row. There is not a unique recipe to organize the tables and to connect them among each others, i.e. to define the structure of a database. However, the information retrieval capabilities and performances are strictly connected to the structure of the database.

The structure of HEPData, as inferred from the aforementioned dump, is shown in Fig. 1, where one can see that there are two prominent tables called *Datasets* and *Papers*, connected between each other. Each of them is then connected to many other tables, which, in turn, can be connected to further ones. Although this structure and the relations between tables reflect an intuitive way of organizing the information provided by the experimental collaborations in papers including their data, the complexity of the schema makes it difficult, if not impossible, to write and efficiently execute queries capable of extracting information from tables distant many hops from each other.

In order to allow for complex queries, we have reorganized the information in a star-schema, as shown in Fig. 2. Considering each data point as a collection of  $(\{x_i\}, y)$  values, where  $y$  is a function of the  $\{x_i\}$  coordinates, we have chosen to center the star-schema around the table of the  $y$ -values of all data points published by the experiments, to ensure the maximum granularity in information retrieval. The central table, also called *fact* table according to the OLAP terminology, is surrounded by *dimensional* tables, each of which refers to a group of attributes associated to each entry of the fact table. We found it convenient to reorganize the relevant information concerning these attributes in eight dimensional tables, called 1) *XAxesDim*, containing the information on the  $\{x_i\}$  coordinates, to which each value  $y$  stored in the fact table refers; 2) *YErrorDim*, containing the information on the uncertainties on  $y$ ; 3) *ObservableDim*, containing the information on the observable to which the  $y$  value refers; 4) *ReactionDim*, containing the information on the reaction in which the datum was measured; 5) *ReactionStateDim*, containing the information on each single particle in the initial and final state of this reaction; 6) *PropertiesDim*, containing the information on the properties of the reaction, from the center-of-mass energy to the analysis cuts; 7) *PaperDim*,





**Figure 2:** Star-schema structure at the basis of the MineHEP search engine.

containing various information (arXiv number, doi, collaboration, etc.) on the paper to which the data point belongs; 8) *KeywordDim*, containing the keywords associated to each data point.

We have started testing the reliability and the robustness of the implementation, by developing a set of queries of increasing complexity, which can be interesting for typical HEP users. An example of a query for selecting data on inclusive production of  $B^+$  mesons in  $pp$  collisions, interesting for DIS experts aiming at performing PDF fits, is reported in Fig. 3. The complexity of this query can be further increased by requiring that the experimental points are obtained under additional analysis selection cuts.

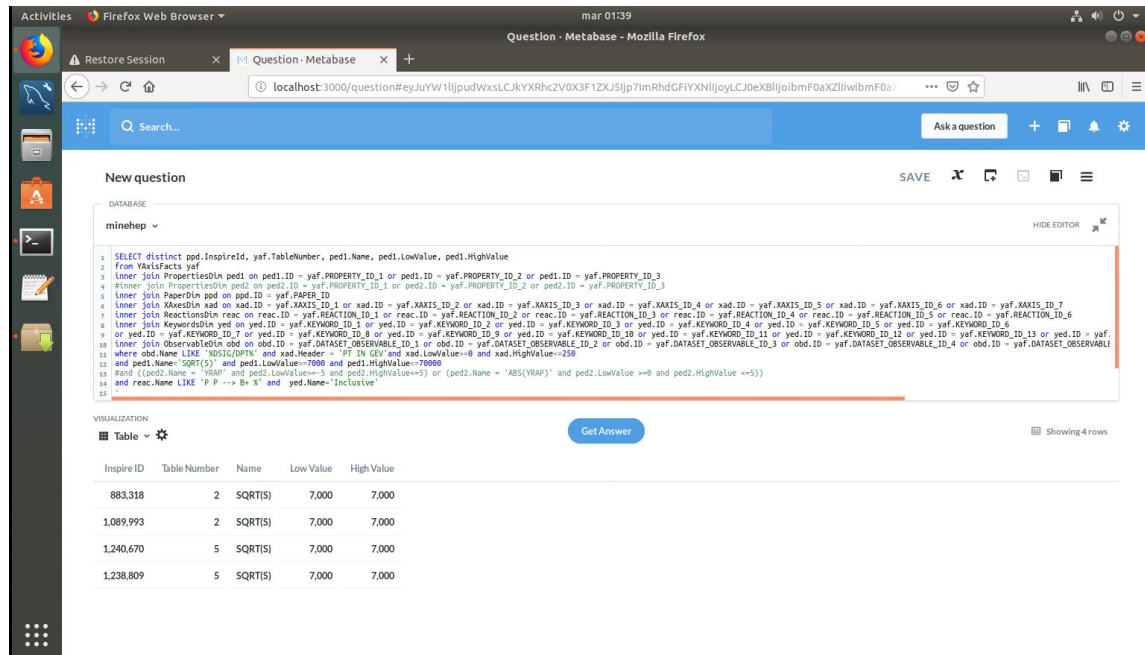
In case of queries involving multiple entries of a same dimensional table (e.g. multiple keywords, or multiple properties), it is worth exploring the performance of variants of the star schema, making use of intermediate bridge tables. We are currently working on this topic.

In the future we plan to host the MineHEP search engine prototype on a public server, granting access to users, who will be able to profit of its query capabilities and will provide very useful feedback for its further development.

**Acknowledgements** We are grateful to the members of IPPP, Durham, for useful discussions at different stages of this project.

## References

- [1] Z. Akopov *et al.* [DPHEP Study Group], *Status Report of the DPHEP Study Group: Towards a Global Effort for Sustainable Data Preservation in High Energy Physics*, arXiv:1205.4667 [hep-ex].



The screenshot shows the Metabase interface with a MySQL query and its results. The query is as follows:

```

1 SELECT distinct ppp.InspireID, yaf.TableNumber, ped1.Name, ped1.LowValue, ped1.HighValue
2 from VxAxisFacts yaf
3 inner join PropertiesDtn ped1 on ped1.ID = yaf.PROPERTY_ID_1 or ped1.ID = yaf.PROPERTY_ID_2 or ped1.ID = yaf.PROPERTY_ID_3
4 inner join PropertiesDtn ped2 on ped2.ID = yaf.PROPERTY_ID_1 or ped2.ID = yaf.PROPERTY_ID_2 or ped2.ID = yaf.PROPERTY_ID_3
5 inner join PaperDtn ppp on ppp.ID = yaf.PAPER_ID
6 inner join XAxisDtn xad on xad.ID = yaf.XAXIS_ID_1 or xad.ID = yaf.XAXIS_ID_2 or xad.ID = yaf.XAXIS_ID_3 or xad.ID = yaf.XAXIS_ID_4 or xad.ID = yaf.XAXIS_ID_5 or xad.ID = yaf.XAXIS_ID_6 or xad.ID = yaf.XAXIS_ID_7
7 inner join ReactionsDtn reac on reac.ID = yaf.REACTION_ID_1 or reac.ID = yaf.REACTION_ID_2 or reac.ID = yaf.REACTION_ID_3 or reac.ID = yaf.REACTION_ID_4 or reac.ID = yaf.REACTION_ID_5 or reac.ID = yaf.REACTION_ID_6
8 inner join KeywordsDtn yed on yed.ID = yaf.KEYWORD_ID_1 or yed.ID = yaf.KEYWORD_ID_2 or yed.ID = yaf.KEYWORD_ID_3 or yed.ID = yaf.KEYWORD_ID_4 or yed.ID = yaf.KEYWORD_ID_5 or yed.ID = yaf.KEYWORD_ID_6
9 or yed.ID = yaf.KEYWORD_ID_7 or yed.ID = yaf.KEYWORD_ID_8 or yed.ID = yaf.KEYWORD_ID_9 or yed.ID = yaf.KEYWORD_ID_10 or yed.ID = yaf.KEYWORD_ID_11 or yed.ID = yaf.KEYWORD_ID_12 or yed.ID = yaf.KEYWORD_ID_13 or yed.ID = yaf
10 inner join ObservablesDtn obd on obd.ID = yaf.DATASET_OBSERVABLE_ID_1 or obd.ID = yaf.DATASET_OBSERVABLE_ID_2 or obd.ID = yaf.DATASET_OBSERVABLE_ID_3 or obd.ID = yaf.DATASET_OBSERVABLE_ID_4 or obd.ID = yaf.DATASET_OBSERVABLE
11 where obd.Name LIKE 'B51G10P1%' and xad.Header = 'P' IN GeV and xad.LowValue=0 and xad.HighValue=250
12 and ped1.Name='SQRT(S)' and ped1.LowValue=7000 and ped1.HighValue=70000
13 and ((ped2.Name = 'YRAP' and ped2.LowValue=5 and ped2.HighValue=5) or (ped2.Name = 'ABS(YRAP)' and ped2.LowValue >=0 and ped2.HighValue <=5))
14 and reac.Name LIKE 'P p --> B+ B+' and yed.Name='Inclusive'
15

```

The results table is as follows:

Inspire ID	Table Number	Name	Low Value	High Value
883,318	2	SQRT(S)	7,000	7,000
1,089,993	2	SQRT(S)	7,000	7,000
1,240,670	5	SQRT(S)	7,000	7,000
1,238,809	5	SQRT(S)	7,000	7,000

**Figure 3:** Example of a MySQL query to the MineHEP search engine, executed through its preliminary interface to Metabase [11]. It selects all tables of papers reporting the  $p_{T,B^+}$  distribution in inclusive  $B^+$  meson production in  $pp$  collisions at center-of-mass energies  $\sqrt{s} = 7$  TeV, for  $0 < p_{T,B^+} < 250$  GeV. The output of the query, reporting the InspireID of the paper and the relevant Table number, is also shown.

- [2] W. J. Stirling and M. R. Whalley, *A Compilation of Drell-Yan cross-sections*, J. Phys. G **19**, D1 (1993).
- [3] A. Buckley and M. Whalley, *HepData reloaded: Reinventing the HEP data archive*, PoS ACAT **2010**, 067 (2010) [arXiv:1006.0517 [hep-ex]].
- [4] J. M. Butterworth, S. Butterworth, B. M. Waugh, W. J. Stirling and M. R. Whalley, *The CEDAR project*, hep-ph/0412139.
- [5] S. Belov, L. Dudko, D. Kekelidze and A. Sherstnev, *HepML, an XML-based format for describing simulated data in high energy physics*, Comput. Phys. Commun. **181**, 1758 (2010) [arXiv:1001.2576 [hep-ph]].
- [6] <http://hepdata.cedar.ac.uk/>
- [7] A. Buckley, W. J. Stirling, M. R. Whalley, J. M. Butterworth, J. Monk, E. Nurse and B. Waugh, *HepData and JetWeb: HEP data archiving and model validation*, hep-ph/0605048.
- [8] E. Maguire, L. Heinrich and G. Watt, *HEPData: a repository for high energy physics data*, J. Phys. Conf. Ser. **898**, no. 10, 102006 (2017) [arXiv:1704.05473 [hep-ex]].
- [9] <https://www.hepdata.net/>
- [10] R. Kimball, M. Ross, *The Data Warehouse Toolkit, 3rd Edition: The Definitive Guide to Dimensional Modeling*, Wiley, Indianapolis 2013
- [11] <https://metabase.com/>