

The VizieR database of astronomical catalogues

F. Ochsenbein, P. Bauer, and J. Marcout

CDS, Observatoire Astronomique, UMR 7550, 11 rue de l'Université, F-67000 Strasbourg, France

Received December 10, 1999; accepted January 17, 2000

Abstract. VizieR is a database grouping in an homogeneous way thousands of astronomical catalogues gathered for decades by the Centre de Données de Strasbourg (CDS) and participating institutes. The history and current status of this large collection is briefly presented, and the way these catalogues are being standardized to fit in the VizieR system is described. The architecture of the database is then presented, with emphasis on the management of links and of accesses to very large catalogues. Several query interfaces are currently available, making use of the *ASU* protocol, for browsing purposes or for use by other data processing systems such as visualisation tools.

Key words: astronomical data bases: miscellaneous — catalogs

1. Introduction

The Centre de Données astronomiques de Strasbourg (CDS) has a very long experience in acquiring, cross-identifying, and distributing astronomical data (Genova et al. 2000): a collaboration for the exchange of what was called *machine-readable* astronomical data started with the *NASA-GSFC* and the *Astronomisches Rechen-Institut* around 1970. This collaboration has been maintained over this 30 year period, and collaborations with other institutes for similar exchanges have been developed. The volume of data shared of course increased, at a rate which has been exploding in the recent years.

Compared to the late 60's, where the bulk of the *machine-readable* data consisted in a set of the basic catalogues carefully keypunched, the situation has changed drastically, now that every instrument or detector is generating megabytes or gigabytes of daily output. These huge data sets are hopefully not stored in data centers, but

Send offprint requests to: F. Ochsenbein,
e-mail: francois@astro.u-strasbg.fr

are processed in the observing center where the expertise exists to generate the best high-quality archives and catalogues in a form usable by astronomers who are not familiar with the instrument. The Data Centers' role is essentially to collect such "final" catalogues, or more generally high-quality data, i.e. data which either were published in the refereed scientific literature, or at least a paper describing these data and their context was accepted for publication in a refereed scientific journal.

Making an efficient usage of the data distributed by the data centers — for instance for the analysis of the statistical properties of some interesting population of stars — often requires to *combine* data coming from several data sets; this operation is far from simple, and this is why the first creation of CDS was SIMBAD, a data-base resulting from the cross-identification of the major catalogues, later expanded to thousands of catalogues and to published literature (see Wenger et al. 2000).

The VizieR system results from a different approach: the astronomical catalogues are kept in their original form, but homogeneous *descriptions* of all these data sets are provided in order to maximize their usability. In other words, VizieR relies on an homogenization of the *catalogue descriptions* — what is also called *metadata*, or data describing other data — to transform the set of *machine-readable* astronomical catalogues into a set of *machine-understandable* data. VizieR actually consists in an interface able to query this set of *machine-understandable* astronomical catalogues.

2. Astronomical catalogues

Jaschek (1989) defined a catalogue as a *long* list of *ordered data* of a specific kind, collected for a *particular purpose*. What a *long* list means has evolved dramatically in the last decade: the new way of processing data actually resulted in a tremendous increase in both the number and the volume of the astronomical catalogues. To illustrate the evolution in the domain of catalogued *surveys*, one can remember that the largest

Table 1. Evolution of the annual number of papers, and the percentage of papers with associated electronic data, for some of the main astronomical magazines

Journal	1994		1995		1996		1997		1998		1994–1998	
	Papers	%El.	Papers	%El.	Papers	%El.	Papers	%El.	Papers	%El.	Papers	%El.
A&A	1300	1.3	1223	2.9	1394	5.3	1525	6.6	1569	4.1	5711	4.8
A&AS	236	42.8	269	42.0	438	28.4	298	49.0	159	43.3	1164	38.9
ApJ(L)	2064	0.3	2121	0.4	2166	1.1	2255	0.8	2235	0.6	10841	0.6
ApJS	255	12.9	138	25.4	116	22.4	115	16.5	102	11.6	726	17.2
PASP	158	7.0	161	4.3	153	2.0	159	2.5	181	2.8	812	3.7
AJ	425	10.4	504	14.1	477	9.2	460	8.3	501	9.4	2367	10.3
MNRAS	656	1.4	752	2.7	775	0.8	833	1.6	980	1.0	3996	1.5

Table 2. Summary of the evolution of accessible digital catalogues in the last five years (number of catalogues and sizes in Mbytes). The last column gives the number of catalogues with a standardized description (see Sect. 4)

Category	June 1994		June 1995		June 1996		Oct. 1997		Oct. 1998		Oct. 1999		Std
	N	Mb	N	Mb	N	Mb	N	Mb	N	Mb	N	Mb	
I Astrometric	151	1258	158	1292	167	1460	199	2502	207	2777	210	2798	113
II Photometric	144	307	152	320	153	305	187	467	194	525	198	563	110
III Spectroscopic	119	162	126	172	125	173	158	233	163	245	170	249	100
IV Cross-Identification	16	89	16	89	15	83	17	91	17	91	17	91	5
V Combined Data	63	367	63	372	65	365	76	557	84	728	86	842	53
VI Miscellaneous	43	157	49	188	50	502	70	634	71	634	73	653	46
VII Non-stellar	115	361	119	280	122	371	157	425	178	453	180	453	121
VIII Radio	24	269	28	269	29	269	39	414	46	615	53	853	51
IX High-Energy	—	—	—	—	—	—	6	77	8	79	10	200	10
J/A+A <i>A&A</i>	58	2	98	4	158	8	299	16	371	22	424	26	424
J/A+AS <i>A&A Suppl.</i>	123	12	235	24	350	33	544	55	698	73	817	83	817
J/AJ <i>Astron. J.</i>	15	1	91	6	126	10	252	21	295	25	345	31	345
J/ApJS <i>ApJ Suppl.</i>	13	1	36	4	52	7	111	14	147	16	165	18	165
J/(P)AZh <i>Russian Astron. J.</i>	—	—	—	—	—	—	18	0.2	26	0.5	49	0.9	49
J From Journals	233	17	517	40	766	60	1404	118	1771	151	2087	180	2087
<i>Grand Total</i>	908	2986	1228	3022	1492	3588	2313	5517	2739	6299	3084	6882	2692

catalogues in the beginning of this century, called the *Durchmusterungen* — the *Bonner*, *Cordoba* and *Cape* *Durchmusterungen* — provided only a position and a visual estimate of the brightness for $\sim 1.5 \cdot 10^6$ stars, and required over 50 years to be completed. Today, a catalogue gathering similar parameters — with an accuracy one order of magnitude better — is well represented by the USNO-A2.0 (Monet 1998) which contains roughly $5 \cdot 10^8$ sources, almost three orders of magnitude larger. Even larger catalogues are being built: let us quote the GSC-II (Greene et al. 1998) which should contain all optical sources brighter than 18th magnitude, which can be estimated to about $2 \cdot 10^9$ objects.

The existence of these new *mega-catalogues* — which are, in fact, rather *giga-catalogues* — does however not mean that the old catalogues can just be ignored: virtually any astronomical object can be subject to variability, maybe over periods of several centuries, and the discrepancies between old and newer results have therefore to be analyzed.

Another important source of tabular material consists in tables published in the astronomical literature. These tables are now almost always originally in digital form, and contain highly processed data which usage can be precious; access to these electronic data is also essential for maintaining the large databases like SIMBAD or NED.

The potential interest of the reusability of these tables led the Editors of the leading astronomical journals to distribute the tabular material in electronic form. The first realisations for *A&A* started in 1993 (see Ochsenbein & Lequeux 1995), and Table 1 summarizes the frequency of the availability of electronic tabular data among the publications in some of the main astronomical journals in the recent years: not surprisingly, the *Supplement Series*, which were created essentially for the presentation of the observational results, show a high rate of associated electronic data.

3. Astronomical catalogues in the data centers

3.1. Current contents

The growth of the collection of astronomical catalogues managed by data centers is illustrated by Table 2: the current set of available catalogues is now around 3 000, with an annual increase about 15%. Note that the entity designated as a “catalogue” can represent a table of about 100 entries (e.g. the list of galactic globular clusters), as well as a multimillion source catalogue (e.g. the USNO-A2.0).

In Table 2, the catalogues are grouped according to *categories* which were defined in the 70’s, when the bulk

of astronomical studies were dealing with the properties of stars in the optical wavelength domain. Rather than defining regularly a new classification scheme following the evolution of the discipline, it was decided, in agreement with the other data centers, to assign designations to electronic tables according to the published paper, and to reserve the assignment in the “traditional” categories to somewhat important catalogues or compilations. Simultaneously, it was decided to assign *keywords* to each catalogue, in order to allow easy retrieval of catalogues with similar contents and purposes.

Note that, if most of the catalogues contain data related to the observation of astronomical sources, other types of data are also available, generally grouped in the “Miscellaneous” (VI) category: catalogues of atomic data like wavelength tables or results of the *Opacity Project*, tabulated results of stellar evolution models, ephemeris elements, etc.

3.2. Usage of astronomical catalogues

Table 3. Yearly traffic on the CDS catalogue ftp server (*internal and mirror traffic excluded*)

Period	Files	Gbytes	Nodes
1993 Jan.–1993 Dec.	6106	1.5	458
1994 Jan.–1994 Dec.	23696	6.1	1599
1995 Jan.–1995 Dec.	57314	11.4	4022
1996 Jan.–1996 Dec.	71300	19.8	4953
1996 Oct.–1997 Sep.	143000	43.5	6279
1997 Oct.–1998 Sep.	308840	74.5	9780
1998 Oct.–1999 Sep.	538407	77.1	10146

One of the main goals of the CDS is to promote the usage of the reliable astronomical catalogues to the astronomical community. The “Catalogue Service” has been one of the major CDS services since the beginning of the CDS activity, and used to distribute catalogues on magnetic tapes and floppies; the service has been implemented on the network as a FTP server in March 1992, generating immediately a large increase in the number of distributed files. The FTP activity is still increasing at a high rate, as can be inferred from Table 3: the current traffic is equivalent to a copy of the whole collection every month.

It is also interesting to quote those catalogues which are the most frequently copied from the CDS archives, summarized in Table 4 for the last two years: not surprisingly, *surveys*, and what Jaschek (1989), in his Sect. 5.2, designates as *General Compilation Catalogues*, are among the most popular catalogues. It is also interesting to note the large number of copies of the GSC catalogue (about 300 Mbytes): it was copied by over 500 nodes in the last 12 months, which is 4 times more than in the previous

year; this could indicate that catalogues of this size can be quite easily managed on small computers nowadays.

Table 4. Catalogues which have been the most frequently copied

Number of Nodes 1999	(1998)	Catalogue designation and short title
879	(750)	(I/239) Hipparcos & Tycho Catalogues
502	(123)	(I/220) The HST Guide Star Catalog, V1.1 (Lasker+ 1992)
293	(165)	(VI/87) Planetary Ephemerides (Chapront+ 1996)
284	(241)	(I/131A) SAO Star Catalog J2000 (SAO Staff 1966; USNO, ADC 1990)
248	(60)	(I/197) Tycho Input Catalogue, Revised version (Egret+ 1992)
203	(221)	(VII/118) NGC 2000.0
195	(162)	(V/50) Bright Star Catalogue, 5th Revised Ed. (Hoffleit+, 1991)
173	(145)	(VI/80) Opacities from the Opacity Project (Seaton+, 1995)
169	(134)	(I/246) The ACT Reference Catalog (Urban+ 1997)
126	(142)	(V/70A) Nearby Stars, Preliminary 3rd Version (Gliese+ 1991)
124	(120)	(VI/81) Planetary Solutions VSOP87 (Bretagnon+, 1988)
112	(73)	(VII/207) Quasars and Active Galactic Nuclei (8th Ed.) (Veron+ 1998)
102	(134)	(II/214A) Combined General Catalogue of Variable Stars (Kholopov+ 1998)
101	(76)	(VII/155) Third Reference Cat. of Bright Galaxies (RC3) (de Vaucouleurs+ 1991)
100	(75)	(VI/79) Lunar Solution ELP 2000-82B (Chapront-Touze+, 1988)
99	(153)	(VI/69) Atomic Spectral Line List (Hirata+ 1995)
97	(149)	(V/95) SKY2000 - Master Star Catalog (Myers+ 1997)
90	(118)	(I/196) Hipparcos Input Catalogue, Version 2 (Turon+ 1993)

4. Standardized description of astronomical catalogues

Making use of the data contained in a set of rapidly evolving catalogues, as illustrated by Table 2, raises the problem of accessing and *understanding* accurately the parameters contained in catalogues which are constantly improved. Typical questions to be addressed are: does the catalogue contain colours; if yes what is their reliability; are they expressed in a well-known standard system; are they taken from other publications or catalogues; how can the associated data file be processed? All these details which describe the data — the *metadata* — are traditionally presented in the introduction of the printed catalogue, or detailed in one or several published papers presenting and/or analyzing the catalogued data.

Metadata play therefore a fundamental role: first the scientists have to get information about the *environment* of the data in order to make their judgement about the suitability of the data for their project, such as: date and/or method of acquisition, related publications, estimation of the internal and external errors, purpose of the data collection, etc.; but also a minimal knowledge of the

```

I/221          The Magellanic Catalogue of Stars - MACS      (Tucholke+ 1996)
=====
The Magellanic Catalogue of Stars - MACS
  Tucholke H.-J., de Boer K.S., Seitter W.C.
  <Astron. Astrophys. Suppl. Ser., 119, 91-98 (1996)>
  <The Messenger 81, 20 (1995)>
  =1996A&AS..119...91T =1995Msngr..81...20D
=====
ADC_Keywords: Magellanic Clouds ; Positional data

Description:
  The Magellanic Catalogue of Stars (MACS) is based on scans of ESO
  Schmidt plates and contains about 244,000 stars covering large areas
  around the LMC and the SMC. The limiting magnitude is B<16.5m and the
  positional accuracy is better than 0.5" for 99% of the stars. The
  stars of this catalogue were screened interactively to ascertain that
  they are undisturbed by close neighbours.

File Summary:
=====
  FileName      Lrecl      Records      Explanations
-----
ReadMe          80          .           This file
lmc             52         175779      The Large Magellanic Cloud
smc             52         67782      The Small Magellanic Cloud
=====

Byte-by-byte Description of file: lmc smc
=====
  Bytes Format  Units  Label  Explanations
-----
  1- 12  A12    ---   MACS   Designation
  14- 15  I2     h      RAh    Right Ascension J2000 , Epoch 1989.0 (hours)
  17- 18  I2     min    RAm    Right Ascension J2000 (minutes)
  20- 25  F6.3   s      RAs    Right Ascension J2000 (seconds)
  27      A1    ---   DE-    Declination J2000 (sign)
  28- 29  I2     deg    DEd    Declination J2000 , Epoch 1989.0 (degrees)
  31- 32  I2     arcmin DEm    Declination J2000 (minutes)
  34- 38  F5.2  arcsec DEs    Declination J2000 (seconds)
  40      I1    ---   Npos   Number of positions used
  42- 46  F5.2  mag    Mag    []?=99.00 Instrumental Magnitude
          (to be used only in a relative sense)
  48      I1    ---   PosFlag [0/1] Position Flag (0: ok,
          1: internal error larger than 0.5")
  50      I1    ---   MagFlag [0/1] Magnitude Flag (0: ok,
          1: bad photometry or possible variable)
  52      I1    ---   BochumFlag *[0] Bochum Flag
=====

Note on BochumFlag: 1 if in Bochum catalog of astrophysical information
on bright LMC stars (yet empty)
=====

Author's address:
  Hans-Joachim Tucholke <tucholke@astro.uni-bonn.de>
=====
(End)          Hans-Joachim Tucholke [Univ. Bonn]          20-Nov-1995

```

Fig. 1. Example of a documentation `ReadMe` file

metadata is required by the data processing system in order to merge or compare data from different origins — for instance, the comparison of data expressed in different units requires a unit-to-unit conversion which can be performed automatically only if the units are specified unambiguously.

This need for a description which is readable both by a computer and by a scientist led to a standardized way of documenting astronomical catalogues and tables, promoted by CDS from 1993 in the form of a dedicated `ReadMe` file associated to each catalogue (Ochsenbein 1994). An example of such a file is presented in Fig. 1: it is a plain ascii file, quite easy to interpret for a scientist, and at the same time structured enough to be interpreted by a

dedicated software. The `ReadMe` description file starts with a *header* specifying the basic references — title, authors, references — and contains a few key sections introduced by standard titles like `Description:` or `Byte-by-byte Description of file:`. Such a file is relatively easy to produce by someone who knows the catalogue contents. The example of Fig. 1 represents the documentation of a very simple catalogue, made of just two data tables, each with a small set of parameters. The output catalogue of the *Hipparcos* mission¹ is an example of a much more complex catalogue: it is composed of two fundamental large tables (HIP with 10^5 stars and TYC with 10^6 stars) and

¹ <http://vizier.u-strasbg.fr/cgi-bin/Cat?I/239>

includes a dozen of annex tables, but can still be described by the the same kind of simple standardized documentation.

The most important part of the *ReadMe* file is the *Byte-by-byte Description* which details the table structures in terms of *formats*, *units*, column naming or *labels*, *existence of data* (possibility of unspecified or *null* values), and brief explanations. Among the conventions, some fundamental parameters are assigned fixed labels like sky coordinates (components of right ascension RA... and declination DE... in Fig. 1); a *prefix* convention, detailed in Table 5, is also used to specify obvious relations between a value, its mean error, its origin, etc.

Table 5. Conventions used for label *prefixes*

Symbol	Explanation
<i>a_label</i>	<i>aperture</i> used for parameter <i>label</i>
<i>E_label</i>	mean error (upper limit) on parameter <i>label</i>
<i>e_label</i>	mean error (σ) on parameter <i>label</i>
<i>f_label</i>	<i>flag</i> on parameter <i>label</i>
<i>l_label</i>	<i>limit flag</i> on parameter <i>label</i>
<i>m_label</i>	<i>multiplicity index</i> on parameter <i>label</i> to resolve ambiguities
<i>n_label</i>	<i>note</i> (remark) on parameter <i>label</i>
<i>o_label</i>	number of <i>observations</i> on parameter <i>label</i>
<i>q_label</i>	<i>quality</i> on parameter <i>label</i>
<i>r_label</i>	reference (source) for parameter <i>label</i>
<i>u_label</i>	<i>uncertainty flag</i> on parameter <i>label</i>
<i>w_label</i>	<i>weight</i> of parameter <i>label</i>
<i>x_label</i>	unit in which parameter <i>label</i> is expressed

This standardized way of presenting the metadata proved to be extremely useful, especially for *data checking* and *format conversion*: many errors were detected in old catalogues simply because a general checking mechanism became available. Tools have been developed for generating a Fortran source code which loads the data into memory, or for converting the data into the FITS format which is presently the most “universal” data format understood by data processing systems in astronomy — but unfortunately a data format which is not convenient outside this context (see e.g. Grøsbøl et al. 1988).

During the six years since this standardized way of describing astronomical catalogues has been defined, over 2 600 astronomical catalogues have been described by means of this *ReadMe* file, and the same conventions have been adopted by the other astronomical data centers and journals for the electronic publication of tables. The present (October 1999) figures of the amount of standardized catalogues are summarized in the rightmost column of Table 2; previous figures were presented in an earlier paper (Ochsenbein 1997).

It is expected, in the future, that the authors will supply the documentation of their data in this simple form; it is already the case for a very significant fraction of the tables mailed to the CDS, and in order to help the authors, template files as well as a few tips on how to create

the *ReadMe* file are accessible on the Web². The *ReadMe* files and the data files are then checked by a specialist, who contacts the authors if errors are detected or when changes are necessary to increase the clarity or homogeneity of the description.

5. VizieR organisation

VizieR³ is a natural extension of the usage of the metadata stored in the *ReadMe* files, as an implementation of these metadata in terms of tables managed by a relational database management system (RDBMS).

The first prototype of VizieR was the result of a fruitful collaboration between *ESIS* (European Space Information System, a project managed by ESRIN, a department of the European Space Agency) and the CDS; VizieR has been under full responsibility of CDS since January 1996. It was presented at the 1996 AAS meeting (Ochsenbein et al. 1996), and became fully operational in February 1996. This prototype has been significantly upgraded in May 1997, just in time for the implementation of the final catalogues of the Hipparcos mission. The number of catalogues accessible within the VizieR system has grown since that time to 2 374 catalogues (Table 6).

The core of VizieR consists in the organisation of the *meta dictionary*, i.e. the set of metadata extracted from the standardized *ReadMe* descriptions discussed in Sect. 4. There are however two main problems which had to be solved: the access to very large catalogues (larger than a few million rows) for which RDBMS proved to be inefficient, requiring therefore dedicated search methods, and the generation of *links* allowing to connect two related pieces of information, like other tables in the same catalogue, or spectra, images from remote services, etc.

Table 6. Summary of the VizieR contents (November 1999)

VizieR contents in terms of:	All Catalogues	Dealing with objects having positions
<i>Catalogues:</i>	2374	1247
<i>Tables:</i>	6071	1929
<i>Columns:</i>	77260	30261
<i>Rows:</i>	1.17 10 ⁹	1.16 10 ⁹
(<i>without megacatalogs</i>)	40.3 10 ⁶	31.6 10 ⁶

5.1. META dictionary

The meta-dictionary consists in 3 main tables detailed below, and about 20 annex tables, all stored in a relational database:

² <http://vizier.u-strasbg.fr/doc/submit.htm>

³ <http://vizier.u-strasbg.fr/>

1. *METAcatalog* describes the *catalogues*, a *catalogue* being defined as a set of related tables published together: typically a catalogue gathers a table of observations, a table of mean values, a table of references, a list of related images, etc.; *METAcatalog* details the authors, reference, title, explanations of each stored catalogues. This table contains currently 2 374 rows (Table 6);
2. *METAtab* describes each *data table* stored in VizieR: table caption, number of rows, how to access the actual data, the equinox and epoch of the coordinates, etc. This table contains currently 6 071 rows (Table 6) — i.e. the average catalogue is made of 2.6 tables;
3. *METAcolum* details each of the 77 260 Cols. (Table 6) currently stored in VizieR: column name or *label*, the textual *explanation* of the column contents, *datatypes* (numeric or character) and storage mode within the database (integer or floating-point, maximal length of strings, etc.), *units* in which the data are stored in the data-base and *units* in which the data are presented to the user, edition formats, and a few flags used for searches (e.g. column used as primary key) or data presentation (e.g. column to be displayed in the default presentation of the result). The average table is therefore made of ~ 12.7 columns — in fact ~ 11.7 because each table contains an *identification* column in addition to the original set of columns.

Note that, since the set of *META* tables is itself described in VizieR, the meta-dictionary can be viewed and queried like any of the catalogues stored in VizieR — allowing to locate easily e.g. tables with a large number of rows, or catalogues having the words *mass loss* in the description of one of their columns, etc.

The annex tables of the meta-dictionary contain some definitions, like the list of known data-types (*METAtypes*) and keywords (*METAkwd*); or other details like the acronyms used to designate well-known catalogues like *HIP*, *GSC...* (*METAacro*), the keywords associated to each catalogue (*METAkwd*), detailed notes and remarks (*METAnot*), or the list of those objects which are individually quoted in the *ReadMe* files (*METAobj*). A special indexing scheme (*METAcell*), explained briefly in Sect. 5.5, was built to locate the existing objects in all catalogues in a single run. Details on how to generate links are stored in the *METAamor* table.

5.2. Links in VizieR

The interest of having a *link*, or an *anchor* in HTML terms, becomes obvious when a table contains a column representing a reference to an original paper, as for example in Véron and Véron's compilation of quasars⁴: once the

⁴ <http://vizier.u-strasbg.fr/cgi-bin/VizieR?-source=7207/table1>

rules to transform the contents of this column into an actual link to e.g. the ADS bibliographic service⁵ is set up, details about the authors and references, or even the full article, can then be displayed on the screen by a simple mouse click. Another frequent example is the possible expansion of some footnote symbol into the lengthy note detailed in some other table.

The links existing in VizieR may be classified in the following categories:

1. *hard-wired links* which are part of the standard description presented in Sect. 4, like the existence of notes (stored in the *METAnot* table), or the *r_* prefix (Table 5) which indicates a reference which may be detailed in a table of references;
2. *internal links* which connect tables of the same catalogue: such links may be expressed in terms of *keys* in the RDBMS terminology (definitions of columns as primary and/or foreign keys), by the existence of *note flags*, or by more complex relations stored in the *METAamor* table. Another type of *internal link* allows one to retrieve the spectra or images which are part of the catalogue, but which are stored as separate files;
3. *VizieR links* which refer to another catalogue within the VizieR system;
4. *external links* which refer to any other service, like bibliographic services, external databases or archives, image servers, etc.

While links of the first 3 categories can easily be maintained, the maintenance of the *external links* depends on modifications which are completely outside VizieR's control. These external links are maintained by the *GLU* system (Fernique et al. 1998), a system which (i) allows one to use *symbolic names* instead of hard-coded URLs, and (ii) translates these symbolic names with the help of a *distributed* dictionary in which the service providers keep up the descriptions of their own services only in terms of URL addresses and actual presentation of the query parameters.

5.3. VizieR feeding pipeline

On the average, about one new catalog – or 2.6 tables – is added *daily* into VizieR. Such figures imposed the following constraints on the addition of new tables into VizieR:

1. no human intervention is required to populate the database (the meta dictionary and the data tables): all meta-data related to a catalogue can be found or computed on the basis of documentation and configuration files which are read by the VizieR feeding pipeline;
2. we rely as much as possible on the *standardized description* of the catalogues presented in Sect. 4: this means that the configuration file associated to each

⁵ <http://adswwww.harvard.edu/>

catalogue should be minimized, i.e. as few *ad-hoc* details as possible should be needed besides the `ReadMe` files.

The actual delay required to ingest a new catalogue into the system is currently estimated to something between a few minutes and several days for the preparation of the `ReadMe` description file, depending on the initial presentation supplied by the authors and on the catalogue complexity — the delay can be occasionally longer when problems are encountered, requiring interactions with the authors; and a few seconds up to an hour for the actual ingestion into VizieR from the standardized files.

Table 7. Large catalogues currently implemented in VizieR

Acronym	Rows ($\times 10^6$)	Catalogue designation
USNO-A1.0	488.0	The USNO-A1.0 Catalog (Monet 1997)
USNO-A2.0	526.3	The USNO-A2.0 Catalog (Monet 1998), calibrated against Tycho data
GSC1.1	25.2	HST Guide Star Catalog, 1992 version
GSC1.2	25.2	HST Guide Star Catalog, 1996 version
GSC-ACT	25.2	HST Guide Star Catalog, calibrated against Tycho data [†]
2MASS	20.2	2 μm All Sky Survey, Spring 1999 release (Skrutskie et al. 1997)
DENIS	17.5	Deep Near-IR Survey first release (Epchtein et al. 1999)

[†] calibration made by the *Pluto project* (http://www.projectpluto.com/gsc_act.htm)

5.4. Access to very large catalogues

The second challenge is to open a fast access for querying the *mega-catalogues* introduced in Sect. 2. This denomination was somewhat arbitrarily assigned to catalogues having 10^7 or more rows. Such large catalogues are essentially surveys used as *reference catalogues*, typically to find all objects detected in some region of the sky under some conditions of wavelength, time, object structure, etc. The set of such catalogues currently implemented is summarized in Table 7, but this set will grow rapidly in the near future with the continuation of the infra-red surveys, and the emergence of surveys presently in preparation (SLOAN, GSC-II, NVSS, ...).

The limit of 10^7 rows corresponds to a limit in performance and time required to ingest the tables into the relational databases; the largest table, in terms of number of rows, currently stored in VizieR is the *AC2000* catalog (Urban et al. 1997), with $4.62 \cdot 10^6$ rows.

The method used to access these very large catalogues consists in grouping the objects within carefully designed groups based essentially on the location in the sky, followed by a lossless compression obtained by replacing the actual values by offsets within the group; details about the actual results and performances are described in another

paper (Derriere & Ochsenbein 1999). Each very large catalogue has presently its own organisation which depends on its actual column contents, and therefore requires a dedicated program for accessing it. VizieR stores in its *META* dictionary (see Sect. 5.1) which program has to be called to actually access the catalogue, and the description of the columns as they are returned from the dedicated program.

5.5. Accessing all catalogues from a position in the sky

In order to allow a fast answer to the question: *find out all objects for all available catalogues around some target position*, an indexing mechanism is necessary. The total number of object positions currently stored in VizieR, excluding the *megacatalogues*, is about $32 \cdot 10^6$ (Table 6); a classical indexation, in terms of relational DMBS, shows very poor performances especially in the updating phase: the addition of a new catalogue can require up to 4.6 millions modifications or additions – which becomes dramatically slow.

The method adopted for this indexation consists first in a mapping of the celestial coordinates into a set of boxes using a hierarchical spherical-cubic projection similar to the technique used by SIMBAD (Wenger et al. 2000), but down to a level 8 which corresponds to a granularity of about $20'$, or $6 \cdot 4^8$ ($\simeq 4 \cdot 10^5$) individual boxes. The list of catalogues which exhibit sources in the region of the sky covered by the box is then stored for each of the defined boxes, allowing therefore a fast answer to the question: “what is the list of catalogues which have a fair chance of having at least one source close to a specified target?” The final step consists in looking successively into the matching catalogues.

The method offers the particularity of being *hierarchical*: 6 boxes are defined at level 0, 24 at level 1, ..., and going down one step in the hierarchy consists in dividing each box into four parts. The indexing mechanism recursively groups contiguous non-empty boxes represented by a single box at the upper level, meaning that a dense survey covering the whole sky is just represented by the 6 boxes of level 0 in this index. In practice, the 1247 catalogues with positions are summarized in this index by $3.9 \cdot 10^6$ elements (to be compared to the $31.6 \cdot 10^6$ sources in Table 6), i.e. an average of 3 000 elements per catalogue.

5.6. Current contents

The status of VizieR contents is presented in Table 6, where we distinguished those tables representing data about actual astronomical objects which can be accessed by their position in the sky. In terms of number of available records, those containing celestial positions represent over 78% even when the *megacatalogs* are omitted, even though only 32% of the tables are concerned. In other




Fig. 2. Histogram of the number of rows among VizieR tables (the darker bars correspond to tables containing celestial coordinates)

words, the average table dealing with actual astronomical objects contains around 16 000 rows — a theoretical mean, as can be seen from the histogram of the table populations in VizieR represented in Fig. 2 which shows a modal value around tables of 100 objects.

6. VizieR interfaces

Several interfaces are currently available for an access to the data stored in VizieR: directly from a Web browser, via a construction of the query using the *ASU* conventions, or the developing *XML* interfaces.

6.1. Access from a browser

From a WWW-browser, a “standard query” in VizieR consists in a few steps:

1. Locate the interesting catalogues in the VizieR Service⁶. This can be done in various ways illustrated in Fig. 3: from well-known catalogue acronyms like *HIP* or *GSC*, from a choice in the set of predefined keywords, from authors’ names, or from a self-organizing (or Kohonen) map constructed on the basis of the keywords attached to the catalogues (Poinçot et al. 1998). New possibilities for locating catalogues of interest for the user are currently under development;
2. Once a catalog table – or a small set of catalog tables — is located (for instance the *Hipparcos* Catalog⁷ resulting from the *Hipparcos* mission), *constraints* about what to search and how to present the results can be specified, as:




Fig. 3. Excerpt of the VizieR first search page

- constraints based on the celestial coordinates, i.e. location in the neighbourhood of a target specified by its actual coordinates in the sky, or by one of its name as known in SIMBAD (see Wenger et al. 2000);
 - any other constraint on any of the columns of the table(s), like a minimal flux value, or the actual existence of some parameter (non-NULL value);
 - which columns are to be displayed, and in which order the matching rows are to be presented.
- By pushing the appropriate buttons, it is for instance easy to get the list of *Hipparcos* stars closer than 5 parsecs to the Sun, ordered by their increasing distance⁸.
3. Obtaining full details about one row is achieved by a mouse click in the first column of the result: for instance, the first row of the search for nearby stars described above leads to the VizieR Detailed Page with *Hipparcos* parameters and their explanations concerning Proxima Centauri⁹.
 4. Finally, there may be correlated data, like notes or remarks, references, etc. In our example, Proxima

⁶ <http://vizier.u-strasbg.fr/cgi-bin/VizieR>

⁷ http://vizier.u-strasbg.fr/cgi-bin/VizieR?-source=I/239/hip_main

⁸ http://vizier.u-strasbg.fr/cgi-bin/VizieR?-source=I/239/hip_main&-sort=-Plx&Plx=%3e=200

⁹ http://vizier.u-strasbg.fr/cgi-bin/VizieR-5?-source=I/239/hip_main&HIP=70890

Centauri is related to the α Cen multiple star system, which components can be viewed from the link to the double and multiple stars (CCDM)¹⁰ that appears in the detailed page.

The quantitative monthly usage of VizieR is presently (October 1999) about 40 000 external requests from 2700 different nodes; mirror copies were installed recently in the US¹¹ and in Japan¹² in order to overcome the transcontinental network congestions.

6.2. The ASU protocol

The uniform access to all catalogues is based on the so-called ASU¹³ (Astronomical Standardized URL) protocol resulting from discussions between several institutes (CDS, ESO, CADC, Vilspa, OAT). The basic concept of ASU is a standardized way of specifying queries to remote catalogues in terms of HTTP requests: the target catalogue is specified by a `-source=catalog_designation` parameter, the target sky position by a `-c=name_or_position,rm=radius_in_arcmin` parameter, the output format by `-mime=type`, and general constraints on parameters by `column_name=constraint`. It should be noticed that the representation of a target by the name of an astronomical object (typically a star or galaxy name, e.g. *3C 273*) implies the usage of a *name server* converting a target name into a position in the sky, which is typically achieved by a call to SIMBAD.

6.3. The XML interface

The output of a query to VizieR as presented in Sect. 6.1 can hardly be used by an independent application for further data processing, such as the ALADIN¹⁴ visualisation tool (Bonnarel et al. 2000) which allows to superimpose the catalogued sources on top of actual image of the sky: the application requires an accurate interpretation of the catalogued output in terms of celestial positions in order to find out the exact location of each source. This means that ALADIN has to figure out not only which are the columns representing the celestial coordinates, but also accurate definitions of the system used to express the coordinates, their accuracy, etc. — in other words the *metadata* about the celestial coordinates.

XML (*eXtensible Markup Language*) is an emerging standard which allows to embed markup “*tags*” within a document; the key advantages of this language are that

the *same document* can either be parsed by simple-minded programs (XML uses hierarchical structuring), or can be displayed in the new generation of browsers (via an XSL style sheet which maps the markup “*tags*” into typographical specifications). This language presents other potential interests, especially regarding interoperability issues facilitated by the emergence of generic tools able to process XML documents.

The XML layout of astronomical tables was discussed extensively with interested collaborators, and the agreed definitions were presented at a recent ADASS meeting (Ochsenbein et al. 1999). The output of VizieR is readily available in this format¹⁵, currently used by the Aladin image applet; it is hoped that it will facilitate the usage of the astronomical data in new contexts.

6.4. Current developments

With the large set of homogenized catalogues, VizieR plays a central role in a *data-mining project* currently in development as a collaboration of ESO and CDS, in two main directions: (i) make use of the VizieR large set of described columns (over 70 000 currently) to build up new methods for locating the catalogues which are the best suited to a particular research topic; and (ii) develop automated cross-correlation tools which can take into account the largest possible set of meaningful parameters (Ortiz et al. 1999).

7. Conclusions

VizieR is an illustration of the benefits resulting from an homogeneous documentation of the existing astronomical catalogues, facilitating the transformation of a set of heterogeneous data into a fully interactive database, furthermore able to interact with remote services. The interoperability issues between the databases, in astronomy and probably in connected disciplines, will most likely be among the key developments necessary to allow the scientists to make use of the existing high-quality data without the prerequisite of being familiar with the data.

Acknowledgements. The long-term exchanges of data have been fundamental for these developments; more specifically, we wish to thank Jaylee Mead, Nancy G. Roman, Wayne H. Warren and Gail Schneider at NASA/ADC for decades of collaborative work, and the present director Cynthia Y. Cheung; and Olga Dluzhnevskaya at INASAN, the Russian data center.

The support of INSU-CNRS and CNES is acknowledged, as well as the contribution of ESA-ESIS for the initial developments of VizieR, and more specifically Salim Ansari and Isabelle Bourekeb. The development of VizieR also resulted from fruitful discussions with Françoise Genova, Michel Crézé and Daniel Egret; the enthusiasm of James Lequeux and its implication for

¹⁰ <http://vizier.u-strasbg.fr/cgi-bin/VizieR-6?-source=1239&corr=PK=CCDM&CCDM==14396-6050>

¹¹ <http://adc.gsfc.nasa.gov/vizier/>

¹² <http://z13.mtk.nao.ac.jp/vizier/>

¹³ <http://vizier.u-strasbg.fr/doc/asu.html>

¹⁴ <http://aladin.u-strasbg.fr/>

¹⁵ <http://vizier.u-strasbg.fr/cgi-bin/asu-xml>

the emergence of electronic tables in the *A&A* publication had a large impact on the accessibility of the astronomical data.

We are also grateful to those who contributed in the more tedious, although critical, part of data standardisation: Simona Mei, Joseph Florsch and Patricio Ortiz at CDS; Gail Schneider and collaborators at NASA/ADC; Koichi Nakajima at ADAC/Japan; Veta S. Avedisova and collaborators at INASAN; and we would like to thank also the authors who participated in the elaboration of the documentation about their data, and answered patiently to all our questions.

References

- Bonnarel F., Fernique P., Bienaymé O., et al., 2000, *A&AS* 143, 33 (Aladin)
- Derriere S., Ochsenbein F., ADASS IX Proc. (in press)
- Epchtein N., Deul E., Derriere S., et al., 1999, *A&A* 349, 236
- Fernique P., Ochsenbein F., Wenger M., 1998, “CDS GLU, a tool for managing heterogeneous distributed web services”, in: Albrecht R., Hook R.N. and Bushouse H.A. (eds.), ASP Conf. Ser. 145, 466
- Genova F., Egret D., Bienaymé O., et al., 2000, *A&A* 143, 1
- Greene G., McLean B., Lasker B., 1998, “Development of the Astronomical Image Archive and Catalog Database for Production of GSC-II” in *Future Generation Computers* (in press)
- Grøsbøl P., Harten R.H., Greisen E.W., Wells D.C., 1988, *A&AS* 73, 359; see also <http://fits.gsfc.nasa.gov/>
- Jaschek C., 1989, “Data in Astronomy”. Cambridge Univ. Press
- Monet D., 1997, 1998, The USNO A1.0 and A2.0 catalogues, <http://ftp.nofs.navy.mil/projects/pmm>
- Ochsenbein F., 1994, *Bull. Inf. CDS* 44, 19; see also <http://cdsweb.u-strasbg.fr/doc/catstd.txt>
- Ochsenbein F., 1997, *Proc. International Cooperation in Dissemination of Astronomical Data*, *Baltic Astron.* 6, 221
- Ochsenbein F., 1998, “The VizieR System for Accessing Astronomical Data”, in: Albrecht R., Hook R.N. and Bushouse H.A. (eds.) ASP Conf. Ser. 145, 387
- Ochsenbein F., Albrecht M., Brighton A., et al., 1999, in *Proc. ADASS-IX* (in press); see also <http://cdsweb.u-strasbg.fr/doc/astrores.txt>
- Ochsenbein F., Genova F., Egret D., et al., 1996, *Bull. Am. Astron. Soc.* 187 #9103
- Ochsenbein F., Lequeux J., 1995, *Vistas Astron.* 39, 227
- Ortiz P.F., Ochsenbein F., Wicenc A., Albrecht M., 1999, in ASP Conf. Ser. 172, ADASS VIII, Mehringer D.M., Plante R.L., Roberts D.A. (eds.), p. 379
- Poinçot P., Lesteven S., Murtagh F., 1998, *A&AS* 130, 183-191
- Skrutskie M.F., Schneider S.E., Stiening R., et al., 1997, in *Proc. Workshop The Impact of Large Scale Near-IR Sky Surveys*
- Urban S.E., Corbin T.E., Wycoff G.L., 1997, “The AC2000: the Astrographic Catalogue on the Hipparcos System” CD-ROM, US Naval Observatory
- Wenger M., Ochsenbein F., Egret D., et al., 2000, *A&AS* 143, 9 (the SIMBAD astronomical database)