The RHESSI Experimental Data Center

Pascal Saint-Hilaire[1,2], Christoph von Praun[4], Etzard Stolte[3], Gustavo Alonso[3], Arnold O. Benz[1] and Thomas Gross[4]

[1] *Institute of Astronomy, ETH Zürich, CH-8092 Zurich, Switzerland* (shilaire@astro.phys.ethz.ch)

[2] *Paul Scherrer Institute, CH-5232 Villigen, Switzerland*

[3] *Institute of Information Systems, ETH Zürich, CH-8092 Zurich, Switzerland*

[4] *Laboratory for Software Technology, ETH Zürich, CH-8092 Zurich, Switzerland*

Abstract.

The RHESSI mission breaks now ground for X-ray astronomy in how the data is transmitted to the ground: The raw data does not consist of images nor count rates in certain energy channels but consists for each photon of the time, energy, and the collimator. Therefore, data analysis is also different from past missions: the analysis demands significant hardware resources, takes time, and requires software installations and expertise. The HESSI Experimental Data Center (HEDC) at ETH Zürich aims to facilitate the use of RHESSI data and has explored new ways to speed up browsing and selecting events. HEDC provides pre-processed data for on-line use and allows basic data processing remotely through the Internet. The functionality and contents of HEDC are described here, as well as first experiences by users. HEDC can be accessed at http://www.hedc.ethz.ch.

CAdditional graphical material and color pictures are available on the CD-ROM accompanying this volume.

Keywords:

1. Introduction

The Reuven Ramaty High Energy Solar Spectroscopic Imager (RHESSI) images spectroscopically the Sun in the 3-17000 keV range (soft and hard x-rays, as well as gamma rays), with unprecedented spatial, temporal and spectral resolutions. RHESSI's main objective is to deepen our understanding of energy release and particle acceleration taking place in solar flares, but it also provides spectroscopic information on cosmic gamma-ray bursts, and may even image such extra-solar objects as the Crab nebula (Lin et al., 2002).

RHESSI's nine rotating sub-collimators (RSCs), each fitted at both ends with vertical grids, modulate the incoming X-ray light on the detectors at the end of each RSC. As each RSC's grid pair has different pitches, a set of nine Fourier-like components are obtained, and are used to reconstruct the original image (Hurford et al., 2002). As each incoming photon is tagged both in time, energy and RSC; the photons can be sorted to process photons in a specified energy band for image reconstruction, to construct lightcurves for a certain energy range, or to derive a spectrum.

Image reconstruction can take from a couple of minutes (for a basic 'back projection') to a few hours or more (for a tedious 'pixon' reconstruction) on a state-of-the-art workstation. Image reconstruction must be done by each researcher using significant hardware and software resources (including the use of commercial software, e.g., IDL – the Interactive Data Language). Each flare warrants several images, at different times, energy intervals, accumulation time intervals, image sizes and resolution, etc... If we consider the importance of spectra, lightcurves and other ancillary data, we realize that large amount of computing time is needed to create derived data just to allow an observer to sift through the set of observed flares and to look for scientifically relevant datasets. These preparatory, but necessary, computational activities partially overlap among each observer and project. This observation provided the starting point for the RHESSI Experimental Data Center (HEDC): its purpose is to automatically generate an exhaustive amount of 'quicklook' data products and assemble them in an on-line data warehouse that will allow fast browsing and other services.

Section 2 is a description of the HEDC system. Section 3 explains how to use HEDC. Section 4 summarizes first experiences by users and concludes this paper. Appendix A describes the contents of the HEDC Extended Catalog: events and data products. Appendix B is a list of all user-relevant query attributes on HEDC.

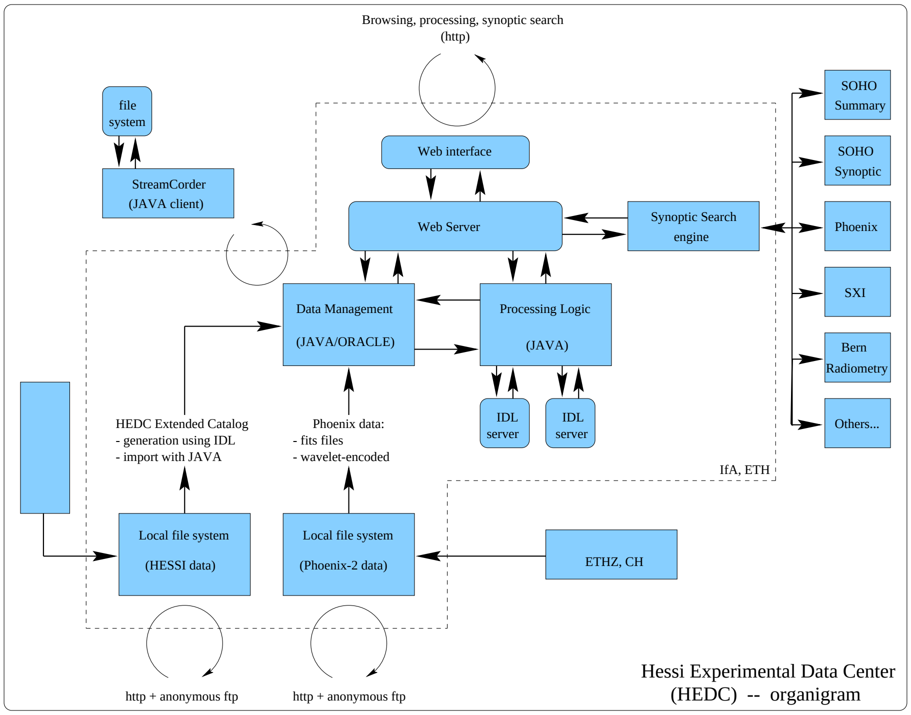

Figure 1. The HESSI Experimental Data Center at ETH Zürich

2. Description of HEDC

HEDC is a joint project of several groups at the Swiss Federal Institute of Technology's (ETH) in Zürich: the Institute of Astronomy in collaboration with the Institute of Information Systems and the Laboratory for Software Technology in the Department of Computer Science. The main services provided by HEDC are:

– on-line database for events and data products,

– on-line RHESSI data processing,

– on-line RHESSI level-0 data repository that contains the raw data,

– other on-line services such as a Synoptic Search engine that quickly retrieves other solar data of relevance.

An 'event' (sometimes more specifically referred to as an 'HEDC event') can be a solar flare, a gamma-ray burst, or an electron precipitation. The term 'other event' is reserved for future extensions that are not yet determined. An event consists of a list of attributes (such as start and end times, total counts, peak count rates in certain energy bands, etc) and a list of associated 'data products'. Data product is the generic term that refers to all derived data: images, lightcurves, spectra, spectrograms, etc... A data product consists of a list of attributes (accumulation times, energy bands, etc...) and a picture (PNG format). The list of attributes for both events and data products are tables for database querying.

The 'E' in the HEDC acronym stands for 'Experimental'. HEDC introduces several innovations but also serves as a platform for deepening our understanding of the scientific data warehouses. HEDC is not simply another astronomical data repository or warehouse, but also a database of scientifically useful derived data. Furthermore, every derived data item (events and data products) is accessible on-line, e.g., through the use of any Web browser. And on-line data processing (on the center's servers) by users supplements the available data products. Users can add their own data products to the database, and the derived data on events of special interest thus increases in a self-organizing way. User participation should increase the scientific return of HEDC (and ultimately RHESSI).

Figure 1 gives an overview of the different parts of HEDC. Its principal components are:

- a file system for all level-0 RHESSI and Phoenix-2 (Messmer et al., 1999) raw data files
- the HEDC Extended Catalog generation (using IDL) subunit,
- the Data Management (DM) subunit,
- the Processing Logic (PL) subunit and the IDL servers,
- the Synoptic Search engine,
- a Web server providing the main user interface,
- the StreamCorder, a future alternative to the Web (browser) interface that provides more flexibility.

2.1. Architecture

HEDC has been implemented as a 3-tier architecture (Figure 2), where the intermediate application logic layer relays requests for data and/or processing by clients to an Oracle 8.1.7 RDBMS (Relational Database Management System) and a number of analysis servers at the resource

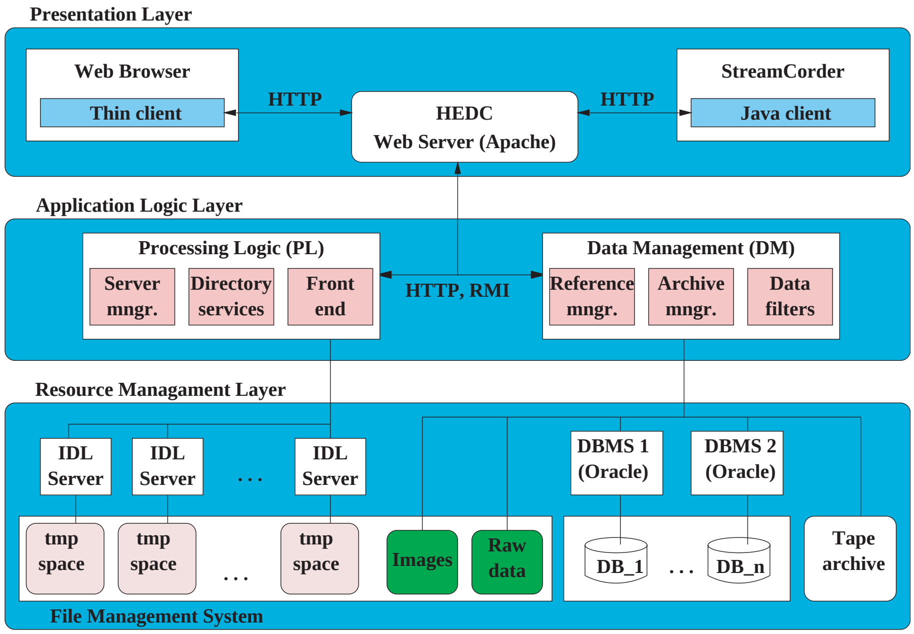

Figure 2. Architecture of HEDC

management layer. The application logic layer consists of two components. The Data Management (DM) component takes care of all data storage issues. The Processing Logic (PL) component acts as an intermediary between the DBMS and the IDL servers. Both are implemented in Java and run as stand-alone programs and servlets. HEDC can be accessed through either a Web-based client using a conventional browser[1] or a Java-based client, the StreamCorder[2]. With any of the two interfaces, users can query and download raw data, view data products and perform new processing steps. The Java-based client offers some functionality not available in the Web interface, e.g. tools for data visualization and system administration.

HEDC currently runs on a SUN Enterprise Server with 2 GBytes RAM, dual 450 MHz processors, a SUN A1000 RAID and 800 GBytes of additional hard-disk drives (storage space will grow as the data collected grows – an additional 2TB RAID was recently added). Critical data, such as the database redo logs and configuration information, is stored on RAID with tape backup. Secondarily generated data is stored on no-backup RAIDs. Raw data files are stored on hard-disks with no back-up. The IDL processing servers execute on the SUN server and on Linux PCs.

[1] On-line at http://www.hedc.ethz.ch/
[2] Available at http://www.hedc.ethz.ch/release/

2.2. On-line RHESSI data repository

RHESSI raw data (the level-0 data) are stored in the form of FITS files, from which all RHESSI derived data are produced, using the RHESSI data analysis software[3] (Schwarz et al., 2002). The raw data are mirrored daily from the Space Science Laboratory in Berkeley (CA, USA) in their entirety. All RHESSI data files are publicly available, via anonymous FTP[4] or by HTTP access through the home page. In keeping with RHESSI's open data policy, the data are immediately available to the public. The raw data come in at a rate of about 1.8 GB per day, taking about three hours to download. The data can be available as quickly as the day following the observation, although a two to three day lag is presently the norm.

2.3. HEDC Extended Catalog generation

The level-0 data files usually come with 'quicklook' derived data already appended to it (included are images and a flare list, amongst other items – see the online software documentation for further details). HEDC adds a large amount of derived data and provides an 'extended' catalog (it's called "extended" because it incorporates those quicklook data from the 'standard' catalog).

After retrieving the raw data, the raw data are first scanned for events of interest: this is presently done mostly via the flare list that is appended to the level-0 data files. Parameters of interest for this event (total counts, count rates, etc...) are then determined. They will later serve as attributes for database queries.

Then, a set of data products is automatically generated for each event: spectra at different times, images at different times using different energy bands, lightcurves, etc. Whenever possible, additional data relevant to the event are added, e.g., Phoenix-2 radio spectrograms, or quicklook images or spectra.

Appendix A gives the current list of the data products computed automatically for each event.

All new events and data products are inserted in the HEDC database, and may be recovered by anyone, via the Web interface or the StreamCorder.

The whole process of creating the extended catalog and inserting its elements into the database is fully automated. This setup provides a standardized set of data products for each event, but the setup is currently error-prone. E.g., when the roll solution is not good, or when

[3] http://www.RHESSI.ethz.ch/software/
[4] ftp://hercules.ethz.ch/pub/hessi/data/

the brightest pixel in a full-Sun map corresponds to the spin axis instead of the actual flare (this problem occurs about 40% of the time), images of the wrong region are made. Future software improvements are expected to handle these cases.

2.4. The data management (DM) component

The Data Management (DM) component handles data requests by external clients and the processing logic (PL). It offers HTTP and RMI (Remote Method Invocation) interfaces that hide the complexity of the lowest tier and provide an abstraction from the database schema. The DM uses JDBC to communicate with the DBMSs. Therefore most of the database related code is vendor independent, and the current Oracle 8i server could be replaced by any other database system.

Web clients access the DM through the HEDC Web server. Requests are first analyzed to determine the sub-systems needed to create the response. Then the relevant data and data references are gathered from the database and data repositories, and a dynamic HTML response page is generated. StreamCorder clients access the DM directly through RMI, so that no HTML pages need to be created.

Most existing data repositories in astrophysics have a very simple system architecture. Often they are no more than file systems accessed through FTP servers, tied together by some home-grown perl scripts. During the development of HEDC, on the other hand, the higher development costs, which incur if a commercial database management systems is used, have been easily offset by the many services provided. Most important, a DBMS provides fast and reliable storage and access from little to very large amounts of data. Furthermore, consistency guarantees prevent conflicting data access by concurrent clients. Tools for system maintenance offer storage allocation, backup and recovery even in the event of serious hard- or software failure. Last but not least, any part of the system can be replicated and extended, to match the resources required as the number of HEDC users grows.

2.5. The processing logic (PL) component

The Processing Logic (PL) allows each user to compute data products beyond those that are part of the pre-rendered extended catalog provided by HEDC. Currently, three types of processing are supported, corresponding to the three main objects of the RHESSI software: imaging, lightcurves and spectroscopy; each processing activity can be configured with a set of basic parameters.

Users access the PL through a Web-based interface, where each user interacts with the system in a personalized session. Processing steps are specified as tasks that are handled by the system in accordance with the availability of computing resources. Processing is done in the background so that users can submit several tasks at once. Each task is an individual batch operation; hence the system does not allow incremental refinement of analyses. Unlike in an interactive setting, previously computed results are not reused. This design choice has been taken to avoid user-specific resource reservations on the server. The absence of heavyset state on the server leaves the scheduling of individual tasks unconstrained, leading to improved resource utilization.

Online processing is integrated with the DM in the following two ways: (1) At the user interface, the PL can be easily accessed while browsing so that the standard attributes of a data product from the extended catalog are copied into the task submission form of the PL. (2) Processing results (PNG and FITS format) can be submitted to the DM and stored permanently in the database.

The implementation of the PL is based on an object-oriented software framework. The system has been designed to easily accommodate changes and additions in the supported types of processing activities; the task control and scheduling system is strictly separated from application specific issues, such that more than 95% of the code are independent of the 3 currently implemented processing activities types. The structure of the system is based on service modules that execute independently and allow distributed task execution in a network of workstations.

Our first experience shows that the duration of individual processing tasks can vary significantly. Hence we limited the total CPU time per task to 20 minutes and constrained the set of admissible input parameters. This policy cannot entirely rule out processings that exceed the CPU limit or specify mismatched parameter combinations. However we observed during a workshop with 30 concurrent users that only around 10% of the submitted processings failed due to one of the aforementioned reasons. The session and task management overhead is marginal compared to the cost of data transfers and computation.

2.6. Other services

2.6.1. *Synoptic Search*

The synoptic search subsystem serves to locate ancillary data related to a particular event in remote astro-archives. The query mechanism resembles a Web-crawler: First, online requests are issued to several remote archives in parallel. Then the results are collected, grouped and displayed to the user. Currently, the only search criterion is the observation time.

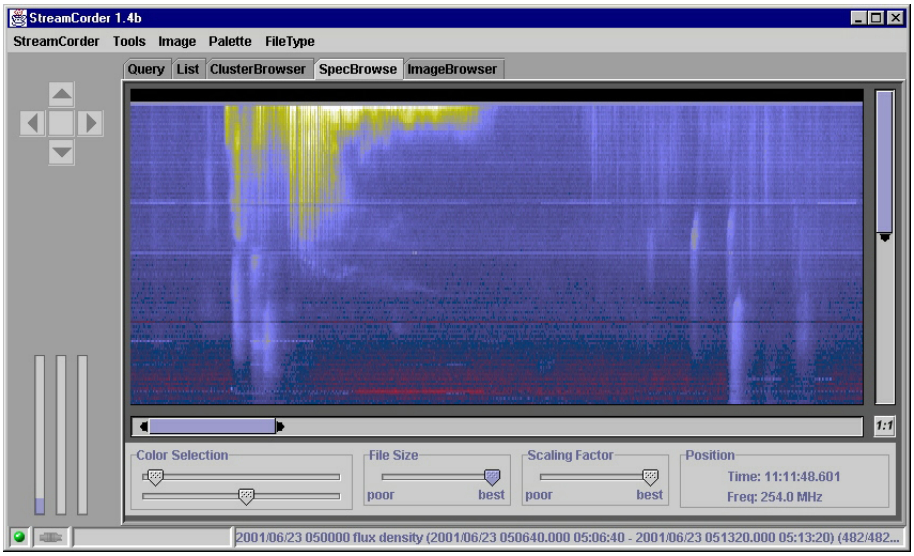

Figure 3. The StreamCorder is an alternative to the Web interface. It can search and display all HEDC data products. Here, it is looking at a wavelet-encoded radio spectrogram from Phoenix-2.

This service operates largely independent from other subsystems of HEDC. The service is best effort (if a query to a remote archive times out, no results are available). Query results are not cached, and there is no data synchronization between HEDC and the remote archives. This light-weight approach of rendering synoptic data accessible through HEDC has proved to be practical and robust. In its current configuration, six popular remote archive sites are searched, including the SOHO synoptic data archive and the RAPP archive at ETH Zurich. Due to its flexible software architecture, additional Internet-archives can be easily integrated.

2.7. The StreamCorder

HEDC is not only accessible through a Web browser, but also through the StreamCorder (see Figure 3). Except for some performance sensitive native routines, it has been completely implemented in Java. The architecture is extensible and modules are loaded according to the current data context. Modules may access core services such as stream management, request queues and local analysis programs. Currently available modules support browsing and download of all datatypes stored in HEDC, allow local and remote processing and offer administrative tools. During RHESSI data processing, the StreamCorder coordinates the asynchronous download, caching, decoding and processing of the data. A local database transparently caches query results and manages downloaded files. The local DBMS schema and the structure of the local file-system archive are identical to the ones on the server. Thus, in combination with semantic caching, global tuple identifiers and synchronization, offline work is possible.

2.8. Phoenix-2 archive

HEDC also holds the Phoenix-2 radio spectrometer archive, both in FITS format and in a wavelet-encoded format (the latter for speedy spectrogram viewing with the StreamCorder).

3. Using HEDC

The term 'browsing' is used in this paper to refer to one of these activities: making database queries (either by event or by data products), exploring the result set by going through the links, making another query (perhaps a finer one, or an entirely new one), etc...

The standard workflow model for users on HEDC is to browse back and forth for events and/or data products, eventually to make new data products online, and to add them to the HEDC database. Once a user has zeroed in on a dataset of interest, he or she will have to make a thorough scientific analysis on the user's workstation, perhaps downloading some of the images previously made on HEDC. Of course, a user may decide to use only parts of HEDC for browsing, processing or synoptic searching.

3.1. Browsing with the Web interface

The HEDC Web interface currently offers two types of querying: by event or by data products (Figure 4). Each of those two possibilities can be done using either a standard Web form, or a more advanced one.

The 'standard' Web form is intended for use by casual users. It is simple and intuitive: there should be no need to consult the on-line documentation. However, only a handful of query fields are available. The 'advanced' form should be used by people more familiar with the system, and offers the full range of user-accessible database queries.

Once the query has been submitted, a 'list page' appears. There is a current limit of 100 entries on this list. Each entry is a different event, respectively data product, with a few attributes (time intervals, energy bands, imaging algorithm, etc...) to guide the next choice of the user: each entry is actually a link to an event (or data product).

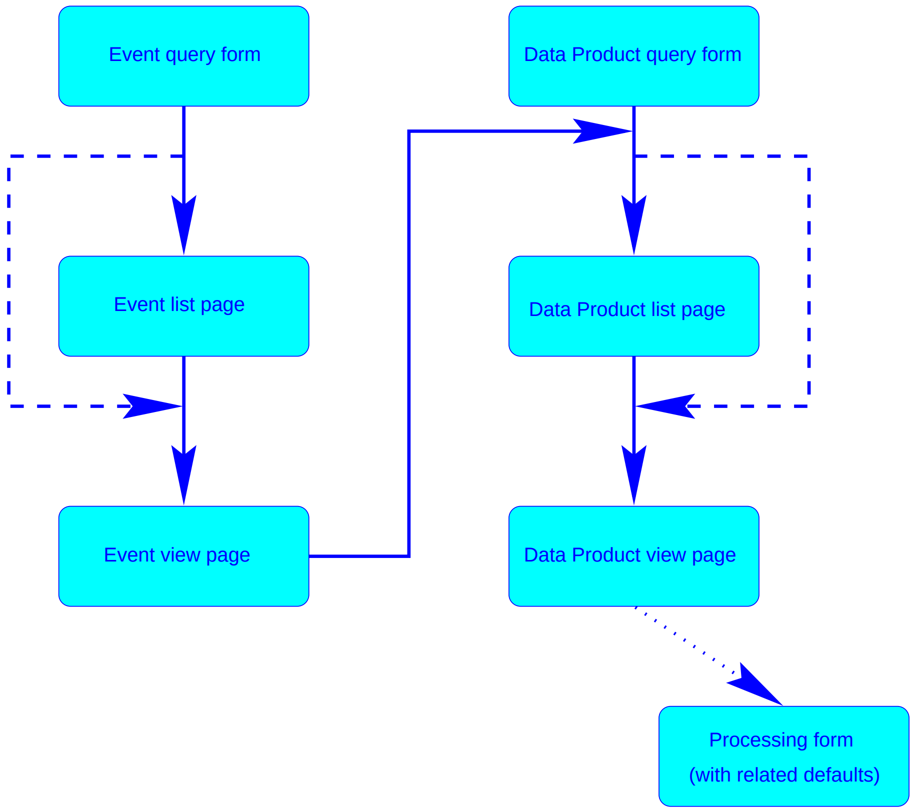

Figure 4. Browsing the HEDC using the Web interface. There are two possible entry points: either querying for events, or for data products. After submitting the appropriate 'query form', a 'list page' containing the result set appears. Clicking on one of its elements leads to the 'view page', where the full set of database attributes can be examined, as well as a picture in the case of data products. If a query results in a single set, then the 'list page' is bypassed. To the 'event view page' is appended a 'data product list page', containing the data products associated with that event. The 'data product view page' also has icons to access the processing form: using those instead of directly going to the processing form via the main link on the left of the Web page has the advantage that most of the processing form's attributes are already defaulted to those of the data product just examined.

Clicking on an event will lead the user to an 'event view page', an HTML page that displays all the event's database attributes (count rates, ...), as well as a list of all its associated data products (and only those). A data product view form also lists that data product's attributes (imaging algorithm, ...), but also a picture.

A comprehensive set of examples are available in the online documentation.

All RHESSI images on HEDC may be downloaded in FITS format (as produced by the RHESSI software's 'fitswrite' method), by clicking on the appropriate icon in the 'data product view page'.

3.2. Processing with the Web interface

Whereas browsing is open to the public, an account is needed to perform processing steps on HEDC (account requests are made online).

Once a properly filled processing form is submitted, a 'job list' appears. It is a listing of all job requests that were sent with their states (pending, running, finished, failed). If a job successfully ended, an icon that displays the resulting picture appears in the job list. Clicking on this icon allows the user to view the full picture, as well as to obtain relevant attributes pertaining to it. One attribute of interest is the IDL output (particularly useful for failed jobs). The current setup allows a maximum of 10 jobs at any time per user.

Jobs stay on the job list until the user officially logs out, up to a maximum of one week. Pushing the "update" button will update the job list to its latest status.

3.2.1. *User events: folders for users' data products*

Once a new data product has been generated and is being viewed by a user, it is possible to store it permanently in the HEDC database. Each user-made data product must be saved in a 'user event', which is just another event on HEDC, and serves as folder (or directory) for users' data products. In this manner, individual users can create several different folders, one for each of their projects, and put in them whatever data products they process on HEDC. User events do not have any attributes, except for a 'code' starting with the user's username. This means that a query for events using time intervals do not reveal user-made events, even if the user-made data products stored inside are within that time interval.

If one is interested in all data products ever generated concerning a certain time interval, one should browse using the data product query form: both user-made and HEDC-made data products are shown (HEDC-made data products have a code similar to their parent event, always starting with 'HX').

3.3. Synoptic engine

Using the HEDC's Synoptic Engine is straightforward: a user enters an approximative date and time of interest and submits the request. A list of available pictures appears.

3.4. The StreamCorder

The StreamCorder gives most of the previously described functionalities, but is much more flexible manner than the browser-based interface. As it is a client application, it runs on a user's local workstation (and therefore uses local resources for processing, in contrast to to the browser-based interface that employs the HEDC server(s)). Additionally, a high level of extensibility exists. E.g., a 'MovieCordlet' extension allows one to rapidly view a sequence of images made on HEDC, the 'Spectrum Browser' enables one to look at wavelet-encoded Phoenix-2 radio spectrograms, IDL sessions can be run locally or remotely, and a 'Cluster Browser' allows one to look at the density of tuple population of HEDC in phase space (i.e. database attribute space). The StreamCorder is fully operational, although slight improvements are still being applied to increase ease-of-use.

4. First experiences and conclusions

At the time of submission of this article, HEDC contains more than 10000 data products of over 400 flares.

The automatic generation of the extended catalog is still complicated by the constantly changing IDL software environment and the lack of reliable flare positions. It will be constantly improved and completed over the next months, and derived data generated for the extended catalog will be reprocessed regularly until a final satisfactory stage is reached.

Having a database instead of a file system to classify all pertinent RHESSI data products has already been recognized as a big advantage: a single event can warrant so many data products that users get lost trying to sort them again if they rely on a standard file system. HEDC not only sorts them but makes them available for easy searches.

For the astrophysics researcher, the quick determination of flare position has been found to be most useful. Many small flares not in the catalog are located by users and stored. The demand for this function will greatly increase once the aspect solution will be more reliable. Also in high demand are lightcurves of all RHESSI data and the visualization of the data in the observing summaries of satellite orbits.

Last, but not least, HEDC is appreciated by those working at home who do not have available the necessary software or sufficient transmission bandwidth to download raw data files.

Acknowledgements

We thank the RHESSI software team for continuous encouragements and support, in particular Jim McTiernan, Kim Tolbert and André Csillaghy for their help, explanations and goodwill.

The RHESSI work at ETH Zurich is supported by the Swiss National Science Foundation (grant nr. 20-67995.02) and ETH (grant TH-W1/99-2).

References

Lin R.P. et al., 2002, Solar Phys. this volume

Hurford, G. et al., 2002, Solar Phys. this volume

McTiernan J.M. et al., 2002, Solar Phys. this volume

Schwarz R.A. et al., 2002, Solar Phys. this volume

Messmer P., Benz A.O., Monstein C., 1999, Solar Phys. 187 (2), 335

Stolte E., Alonso G., "Approximated Trial and Error Analysis in Scientific Databases", Journal of Information Systems, 2002, Vol. 27, No. 8

Stolte E., Alonso G., "StreamCorder: Fast Trial-and-Error Analysis in Scientific Databases", ICDE International Conference on Data Engineering (ICDE), 2002, San Jose, USA

Appendix

A. HEDC Extended Catalog contents

This appendix describes the current state of what is being generated and stored on HEDC. It is liable to change. Consult the on-line documentation for the latest updates. Currently, the generation of the HEDC extended catalog is done about a week after observation by RHESSI. Later reprocessings will occur periodically and incrementally, following improvements or additions to the catalog generator, or major modifications to the raw data or the flare list. The newest, reprocessed versions of HEDC events and their associated data products will replace previous versions. Of course, user-made events and data products will never be reprocessed.

A.1. Detection of events

Currently, only solar flares and some 'other' flares (i.e. with parameters still undefined) are being looked for and generated. Later, this might be extended to gamma-ray bursts and electron events.

Solar flares are given by the flare list attached to the level-0 data. Basically an increase in photon count rates in the 12-25 keV energy band is looked for. The signal must also be strongly modulated in RHESSI's two coarsest detectors (number 8 and 9). See the RHESSI Data Analysis Software pages[5] for more details on this.

'Other flares' are those enhancements in the count rates as seen by HEDC or the flare list, and for which no other classification was (yet) found.

After an event is detected, and its type determined, a set of attributes are determined which characterize the event for later database queries.

A.2. Determination of event attributes

Attributes for each events (such as start & end times, total counts, peak count rates in certain energy bands, SAA and eclipse flags, etc...) are determined as each event is generated. Those attributes can be used as search fields during database queries.

Appendix B gives a full listing of HEDC event attributes.

A.3. Data Products automatically generated with each event

For all events:

- lightcurves of the whole event, in different energy bands.
- three spectra in the 3-2500 keV range, with one minute accumulation time. One done at peak time, one midway between start time and peak time, and one midway between peak and end time.
- images made from Observing Summary (McTiernan et al., 2002) data: count rates in different energy bands; RHESSI trajectory on a mercatorian projection of Earth; Modulation variance lightcurves; flags; geomagnetic latitude (Figure 5 and 6).
- RHESSI spectrograms are generated. If possible, they are also superposed with radio spectrograms from Phoenix-2 (Figure 7).
- background-subtracted time series of the 25-50 keV over 12-25 keV counts ratio for the whole event.

Additionally, for 'solar flare' events only:

- full-Sun image (Figure 8)

[5] http://www.RHESSI.ethz.ch/software/

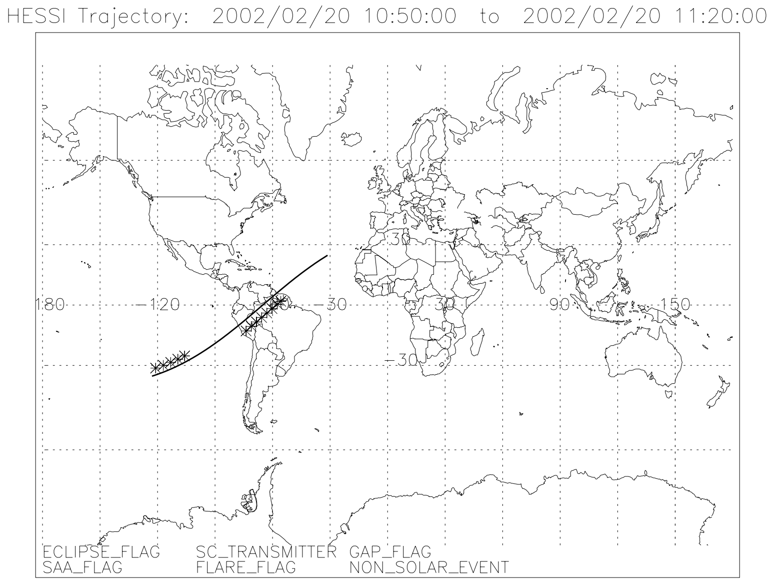

Figure 5. For each event, HEDC has a projection of RHESSI's subpoint on a mercatorian view of the Earth. A list of Observing Summary flags is shown at the bottom of the plot. Those are the flags that are being checked for and displayed with stars along the trajectory. This normally color-coded map is available with the CD-ROM material accompanying this volume.

- movies, i.e. series of images in the following energy bands: 3-12, 25-50, 50-100 and 100-300 keV

- 'quicklook' images and spectra (i.e. those that are included with the raw data) are also extracted and inserted in the database.

- A panel of up to 5x5 images (up to 5 different time intervals, in 5 different energy bands) of the region of interest (Figure 9).

- A panel of 3x3 images, one for each RHESSI sub-collimator, of the region of interest, using the back projection imaging algorithm (Figure 10).

Appendix B gives a full listing of data product database attributes.

The time it takes to generate a single 'solar flare' event and its associated data products is less than one hour for the above list of data products. More images per event will certainly be generated later on, increasing the processing time accordingly.

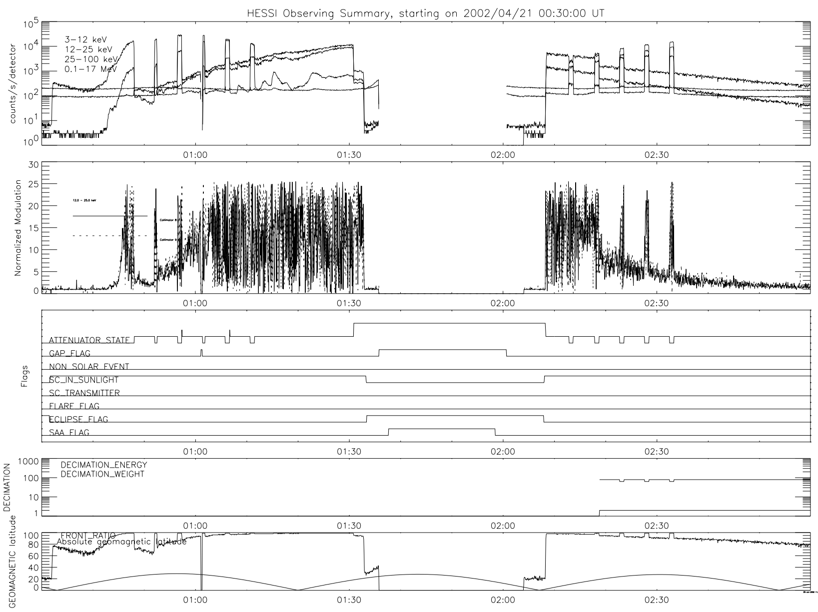

Figure 6. For each event, HEDC has an 'Observing Summary page', showing several products available in the Observing Summary, as well as RHESSI's geomagnetic latitude. The usual color-coded version is available owith the CD-ROM material accompanying this volume.

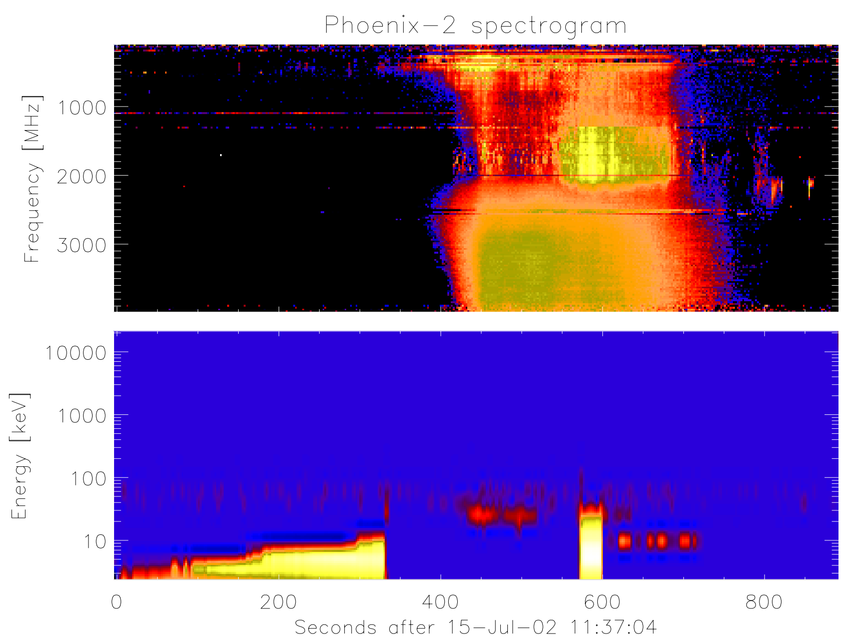

Figure 7. For each event, HEDC has a spectrogram from both RHESSI and Phoenix-2 radio data.

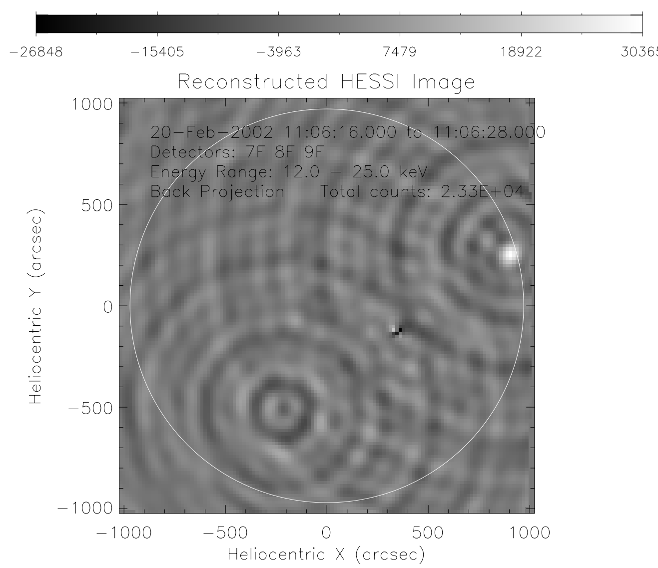

Figure 8. For each event, HEDC has a full-Sun back-projected image, at the peak of the 12-25 keV flux.

A.4. Others

- RHESSI mission-long daily lightcurves in different energy bands are available through the home page.

B. Attributes used for browsing queries

Tables I and II are lists of the attributes that may be used by users to query for data on the HEDC using the Web interface's 'expert' query form. The on-line documentation provides an up-to-date listing, as well as additional details.

Table I. List of query fields for HEDC events

code	the 'name' of an event. Typically 'HXS202261026', where HX stands for Hessi eXperimental (i.e. an event made by HEDC, 'S' for solar flare, and '202261026' for the peak time of the February 26th, 2002 10:26 flare). Another possible format is 'hadar', 'hadar001', etc.. for a USER event (= folder) made by user 'hadar'.
event ID	An internal, unique ID number for each event.
event type	'S' for solar flares, 'G' for gamma-ray bursts, 'E' for electron events, 'O' for other flares.
Flarelist	time-concurrent flare list number for a solar flare.
minimum energy	lower edge of highest energy band where flare counts were seen.
maximum energy	upper edge of highest energy band where flare counts were seen.
total counts	total counts of the flare, in the 12-25 keV energy band.
distance to sun	solar flare's offset from Sun center, in arcseconds.
x pos	solar flare's West-East offset on the Sun, in arcseconds.
y pos	solar flare's North-South offset on the Sun, in arcseconds.
creation date	creation date of the event.
start DATE+TIME	date and time of the start of the flare in the 12-25 keV band.
end DATE+TIME	date and time of the end of the flare in the 12-25 keV band.
start time-of-day	time, in seconds since midnight, of the start of the flare.
end time-of-day	time, in seconds since midnight, of the end of the flare.
duration	time between flare's start and flare's end, in seconds.
peak D+T (3-12keV)	date and time of the peak of the flare in the 3-12 keV band.
peak t-o-d (3-12keV)	time, in seconds since midnight, of the peak of the flare in the 3-12 keV band.
total counts (3-12keV)	total counts of the flare, in the 3-12 keV band.
peak rate (3-12keV)	count rate at peak time, in the 3-12 keV band.
peak D+T (12-25keV)	date and time of the peak of the flare in the 12-25 keV band.
peak t-o-d (12-25keV)	time, in seconds since midnight, of the peak of the flare in the 12-25 keV band.
total counts (12-25keV)	total counts of the flare, in the 12-25 keV band.
peak rate (12-25keV)	count rate at peak time, in the 12-25 keV band.
peak D+T (25-100keV)	date and time of the peak of the flare in the 25-100 keV band.
peak t-o-d (25-100keV)	time, in seconds since midnight, of the peak of the flare in the 25-100 keV band.
total counts (25-100keV)	total counts of the flare, in the 25-100 keV band.
peak rate (25-100keV)	count rate at peak time, in the 25-100 keV band.
ratio 25-50/12-25	ratio of counts in the 25-50 keV band over those in the 12-25 keV band at peak time.
source multiplicity	number of sources in a solar flare. Not operational yet.
active region	where the flare occurred, as given by the flare list.
is simulated data	0/1 or NO/YES flag.
S/C in SAA flag	0/1 or NO/YES flag. S/C stands for spacecraft (i.e. RHESSI).
S/C in night flag	0/1 or NO/YES flag. S/C stands for spacecraft (i.e. RHESSI).
backgound rate	background count rate. Not operational yet.
comments	Made automatically by HEDC (e.g. highest geomagnetic latitude during an event), or by a user for a user-made event.
reserves	unused yet.

Table II. List of query fields for data products

code	the 'name' of a data product. For a data product associated with an HEDC-made event, the data product's code is usually the same as the HEDC event's (e.g. HXS202261026). For data products made by users, any combination of 12 characters are possible.
product ID	An internal, unique ID number for each data product stored on HEDC.
product type	'IM' for images, 'SP' for spectra, etc... see the online documentation for a complete listing.
imaging algr	'BACK' for back projection, etc... see the online documentation for a complete listing.
movie code	most of the RHESSI images made on HEDC are meant to be viewed in sequence, i.e. they share energy bands, imaging algorithm, etc..., and differ only by their time ranges. All those images have the same movie code.
movie frame	the order in which an image which is part of a movie appears.
creation date	creation date of the data product.
start DATE+TIME	date and time of the start of the accumulation time for the data product.
end DATE+TIME	date and time of the end of the accumulation time for the data product.
start time-of-day	time, in seconds since midnight, of the start of the accumulation time for the data product.
end time-of-day	time, in seconds since midnight, of the end of the accumulation time for the data product.
duration	accumulation time for the data product.
min energy	lower edge of the energy bands used for the data product.
max energy	upper edge of energy bands used for the data product.
time resolution	time binning for lightcurves (corresponds to LTC_TIME_RES).
front segments used?	a 0/1 or NO/YES flag.
rear segments used?	a 0/1 or NO/YES flag.
subcollimator used	ex: 101111100.
distance to sun center	angular offset (in arcseconds) of the center of an image with respect to Suncenter (image data products only).
xpos	angular x-offset from Suncenter of the center of an image (image data products only).
ypos	angular y-offset from Suncenter of the center of an image (image data products only).
xdimension	number of horizontal pixels in an image (image data products only).
ydimension	number of vertical pixels in an image (image data products only).
xpixel size	horizontal size (in arcseconds) of a pixel (image data products only).
ypixel size	vertical size (in arcseconds) of a pixel (image data products only).
data quality	unused yet.
is simulated data	a 0/1 or NO/YES flag.
is background-subtracted	a 0/1 or NO/YES flag. Not used yet.
other alg. params	text containing information on some other parameters of the data product.
comments	text added by HEDC or by users, for their own data products.
reserves	unused yet.

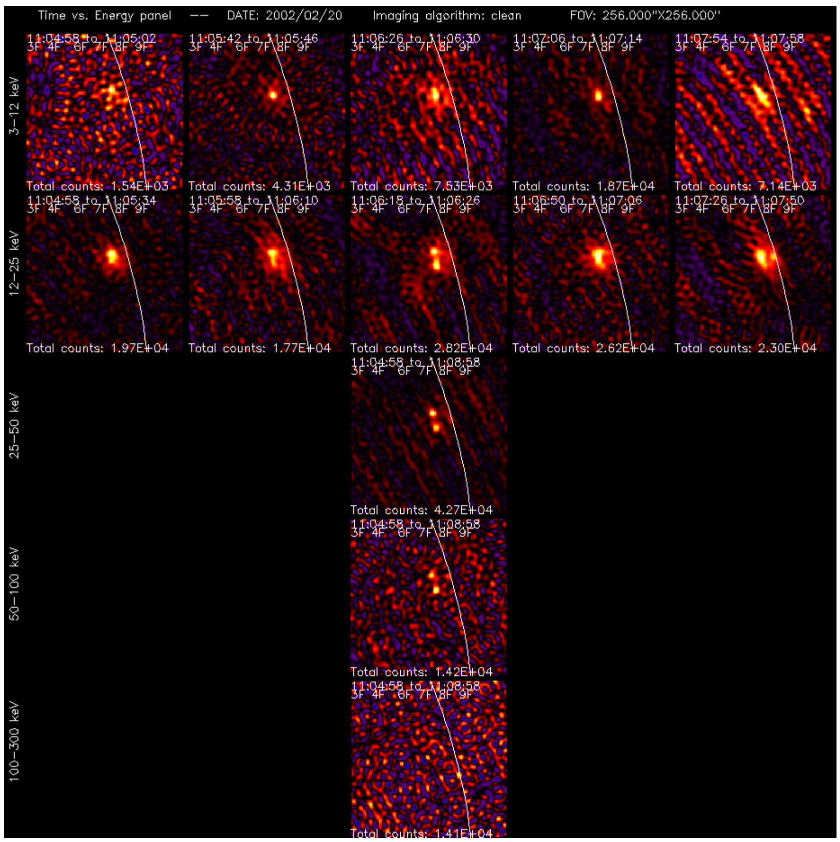

Figure 9. For each event, HEDC has a panel of images showing the evolution of the flare in time (horizontal) and energy (vertical). Only images with a minimal number of counts are made. Hence, small flares do not necessarily have five images in every energy band.

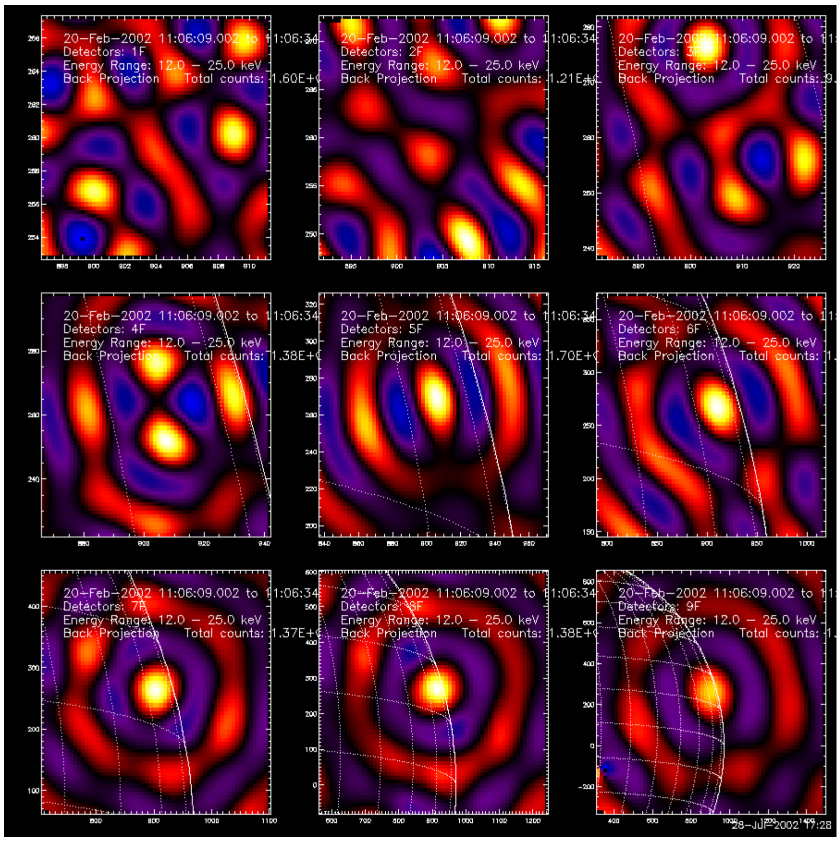

Figure 10. For each event, HEDC has a panel of back-projected images at the peak of the 12-25 keV flux, one for each sub-collimator.