PNI-HDRI PaNdata workshop

February 27

2028

Summary of the 1st joint PNI-HDRI and PaNdata ODI workshop

The 1st joint PNI-HDRI and PaNdata ODI workshop took place at DESY in Hamburg on February 27-28th 2012. The workshop was visited by 52 participants from 14 different facilities or institutions from all over Europe.

This work was partially supported by the Helmholtz Society (HGF) and the European commission under grant agreement RI-283556.

Contents

Summary of the joint PNI-HDRI and PaNdata ODI workshop 2

1. Data policies 2

Embargo period 3

Metadata 3

Retention period 3

Operational costs 3

Costs to archive data 3

2. Hardware developments 4
3. Standard data format 4
4. Analysis Frameworks and GPU acceleration 5
5. PaNdata ODI 5

Talks and Slides 7

Summary of the joint PNI-HDRI and PaNdata ODI workshop

The 1st joint PNI-HDRI and PaNdata ODI workshop took place at DESY in Hamburg on February 27-28th 2012. The workshop was visited by 52 participants from 14 different facilities or institutions from all over Europe. The presentations, list of participants and further information can be found under https://indico.desy.de/conferenceDisplay.py?confId=5517.

PNI-HDRI (http://www.pni-hdri.de) is a high data rate project of the German Photon, Neutron and Ion-Facilities of the Helmholtz-Society.

PaNdata ODI (http://www.pandata.eu) is the Open Data Infrastructure project of the PaNdata collaboration, which includes essentially all European Photon and Neutron research infrastructures.

Both projects have closely co-operated on specific topics related to scientific data management, in particular on the development and implementation of a standard data format (NeXus/HDF5) as well as data policies. However, since there were more themes of common interest, ranging from GPU-accelerated analysis tools to high speed detector development, the workshop had been organized as a joint workshop to intensify and broaden the co-operation.

To separate more technical topics from purely organizational issues, the workshop was divided into three blocks

- reports from the PNI-HDRI work packages
- technical talks and discussions on
    - data format standards and data policies
    - detector development
    - analysis frameworks
    - GPU acceleration of tomographic image reconstruction
    - data and software catalogs
- reports from the PaNdata ODI work packages and the Project Management Board

1. Data policies

After the PNI-HDRI work packages report, Rainer Gehrke (DESY) opened the discussion on the data policy as proposed by PaNdata. The lively discussion was on some points rather controversial. PaNdata foresees an embargo period of 3 years during which the group of scientists involved in the experiment retains exclusive access to the data. The embargo period can be prolonged upon request for example in case of complex and lengthy experiments. Data can however be released any time earlier, for example immediately after publication of the results based on these data, as required by some funding agencies.

Embargo period

Some users as well as facility managements consider the embargo period of 3 years as too short. Prolongation requests should be rare exceptions, and would otherwise create a significant administrative overhead. Due to the increasing complexity of experiments it might become unrealistic to complete analysis within the proposed time frame. From the user side concerns were raised on two additional points. First of all, users fear to lose control over their data. As was pointed out, the research as well as the research infrastructure is fully funded by the public, and it would hence just be consistent to make data available to the public in return. There are however not many policies issued by funding agencies or publishers promoting the release of primary data to the public, with the exception of the ERC guidelines, which clearly state that primary data are deposited to the relevant databases as soon as possible, preferably immediately after publication and in any case not later than 6 months after the date of publication, where primary data should include "unprocessed data" as well.

Metadata

The other concern from user side was the additional effort required to annotate data properly. To obtain persistent IDs for scientific data, certain quality requirements have to be met, which might indeed become an effort. On the other hand, it's a basic best practices prerequisite for publications, and the ubiquitous materials and methods section in a publication might actually be replaced by the references to properly annotated data, which could help to diminish or at least significantly reduce the effort.

Retention period

Some facilities, in particular Neutron RIs with rather modest data rates, consider archiving raw data for an unlimited time. For very large data volumes this might not be feasible or desirable. A minimum retention period of 10 years, in line with requirements from DFG, HGF and other funding agencies, was considered a reasonable figure.

Operational costs

Archiving data of external users at one of the facilities will of course create costs, which are usually not covered at all by any funding schemes. This concerns in particular development, deployment and maintenance of the data infrastructure ranging from large tape robots to network bandwidth. However, most facilities have to provide such an infrastructure for in-house users anyway. Extending the infrastructure to cover external users will certainly increase the overall costs, but considerably less than an exclusive infrastructure. However, support by funding agencies would be highly beneficial to promote such data infrastructures.

Costs to archive data

To keep data on the long term, more in the sense of bit-stream preservation, creates some cost on media and devices. However, the costs were quoted to be in the order of 100€ per Terabyte for the total lifetime of the data. Some facilities quote even lower figures. The costs are considerably lower than the typical user based solution (portable media) and the infrastructure much more reliable. So even if facilities would consider charging research

groups for the archival of data, the costs would still be lower than for any unmanaged solution.

2. Hardware developments

Matthias Balzer (KIT/IPE) presented current dedicated hardware developments at the PNI facilities involved. The talk was focusing on the KIT Camera EVA, which is based on a FPGA Evaluation board and an exchangeable sensor carrier, which enables a wide field of applications. Do to the high speed interfaces, it is capable do deliver images rapidly to an analysis cluster, for example supporting real time reconstruction of tomographic samples.

The second part of the presentation dealt xTCA, which becomes the new generation platform for PNI applications. xTCA provides a reliable, well established and scalable platform with interfaces for infiniband and camera link. Several developments at KIT and DESY are currently exploiting the potential of xTCA in conjunction with new generations of FPGA and ARM (e.g. Xilinx). Such developments can have a high impact on future generations of detectors and data acquisition systems, a topic of general relevance and more intense co-operations among the facilities involved in PNI-HDRI and/or PaNdata ODI.

3. Standard data format

The first larger block of presentations circled around NeXus, the elected standard data format. Eugen Wintersberger (DESY) and Halil Pasic (KIT) provided an overview of the recent developments to consolidate and extend NeXus.

E.Wintersberger presented a complete rewrite of NeXus' C++ API. The API provides an elegant interface to the full set of NeXus classes and applications, and in addition extends the original NeXus API for example on several useful datatypes, adds a stdc++ compliant UTF-8 string handling, and enables a greatly facilitated integration of data sources like e.g. any tango device. The source is freely available under http://www.pni-hdri.de.

H.Pasic proposed extensions to NeXus, in particular to provide more flexibility on file handling. An important aspect is the splitting and repacking of very larger data volumes. NeXus has only rather limited capabilities in this direction. Recording the data layout or structure as meta data within each chunk of data forming a complete dataset enables a very flexible handling of large data volumes. This can be further extended by introducing relationship base classes, which eases integration of multiple channels into the data acquisition chain.

Rainer Gehrke (DESY) outlined a generic NeXus application definition for scattering experiments. He emphasized the need to rapidly converge towards a final definition. It is planned to proceed with the definition and in particular the implementation of the NeXus format at selected beamlines at Petra-III starting with small angle scattering. The session on standardization of data formats and NeXus implementation was concluded by two talks on enhancement and practical usage of NeXus.
Tobias Richter (DLS) summarized the use of NeXus at Diamond, the current status of the implementation and NeXus capable software available.

Jens-Uwe Hoffmann (HZB) presented the developments for the HZB-facilities to easily customize NeXus layouts and to aggregate data streams from different sources in a single NeXus-file.

NeXus has been generally accepted as the common standard and is being widely adopted at the PaNdata facilities. The different developments appear to converge nicely and implementations are progressing well.

4. Analysis Frameworks and GPU acceleration

The afternoon session was devoted to several aspects related to data analysis. The presentations by Gunthard Benecke (DESY), Andy Goetz (ESRF) and Mark Basham (DLS) dealt with the development of data analysis framework. The developments by ESRF and DLS named Data Analysis Workbench (DAWB) and Scientific Data Analysis (SDA) resp., will actually converge into a single, joint development named DAWN. For more information see http://www.dawnsci.org/.

DPDAK is a collaborative effort of HASYLAB and MPI Golm to enable high-throughput data evaluation especially for micro-beam SAXS/WAXS beamlines. Although focusing on Small Angle X-ray scattering, the concept is sufficiently modular and generic to be expanded to other experimental techniques. Naturally, DAWN and DPDAK developments can mutually exploit synergies. Exchange of the concepts might permit to specify APIs in a homogenous fashion, which could facilitate integration of the complementary approaches.

The talk on "Practical Experiences with GPUs for High-Throughput Computing" by Suren Chilingaryan (KIT) exhibited another approach. The UFO framework is tailored to schedule processes on GPUs and CPUs to optimize performance and utilization of multi-core resources, but also provides a customizable framework for tomography, which can easily be implemented as a web service, which might give the other developments a twist toward high performance and real time analysis.

The session was concluded with two talks on HDF5 performance tuning and its application to high throughput experiments. Ulrik Pedersen (DLS) presented the work at Diamond light source on Parallel HDF5 writing and future developments. One aspect is the measurement of performance under realistic conditions, which would be greatly facilitated by a benchmarking library to standardize performance tests for a wide spectrum of different experimental techniques and beamlines. Finally, Mark Basham (DLS) showed example implementations for HDF5, in particular the optimal layout (slicing) of 3D-volumetric data to achieve maximal performance.

5. PaNdata ODI

The second day of the joint HDRI and PaNdata workshop was devoted to the PaNdata Open Data Infrastructure project. As an open session the first section concentrated on topics of more general interest.

After an introduction by the PaNdata ODI project leader, Juan Bicarregui (STFC), Brian Matthews (STFC) presented the current status of iCAT/topCAT and its plans for future developments. iCAT is currently used at a small number of Neutron and Photon sources, and

is the prime candidate for a PaNdata implementation of a data catalogue. The possibility to federate iCAT instances across facilities, one of the core requirements, has been demonstrated.

Jean-Francois Perrin (ILL) continued with the status of the PaNdata software catalogue. Questions were raised to extend the software catalogue for community channels, in particular in view of rapid developments in various fields like FPGA programming.

The so called Umbrella, a shibboleth based authentication prototype system currently tested by user communities, is progressing well, as Stephan Egli (PSI) reported. The "friendly user phase", which aims to collect feedback and experiences with the system from the user communities is on-ongoing. So far, comments have been very positive. The identity management is further being discussed in the EIROforum workshops on federated identity systems for scientific communities.

The work package on Virtual Labs aims to integrate the different modules like authentication, software and data catalogues or data analysis tools like frameworks or workflows. Thorsten Kracht (DESY) summarized the current requirements from the different facilities and work package. The requirements catalogue was essentially finalized, and will be documented later on in the WP5.1 deliverable.

Talks and Slides

HDRI WP1 Status Report	Rainer Gehrke (DESY)	Slides
HDRI WP2 Status Report	Andreas Kopmann (KIT)	Slides
HDRI WP3 Status Report	Joachim Wuttke (JCNS)	Slides
Discussion Data Policy	Rainer Gehrke (DESY)	Slides
Dedicated Hardware in HDRI WP2	Matthias Balzer (KIT)	Slides
NeXus API	Eugen Wintersberger (DESY)	Slides
Extending NeXus for HDRI	Halil Pasic (KIT)	Slides
NeXus Application Definitions	Rainer Gehrke (DESY)	Slides
NeXus at Diamond	Tobias Richter (DLS)	Slides
Collection Software as Link Between NeXus and the Instrument	Jens-Uwe Hoffmann (HZB)	Slides
Directly Programmable Data Analysis Kit (DPDAK)	Gunthard Benecke (DESY)	Slides
Data Analysis Workbench (DAWB)	Andy Goetz (ESRF)	Slides
SDA and DAWN - Scientific Data Analysis	Mark Basham (DLS)	Slides
Practical Experiences with GPUs for High-Throughput Computing	Suren Chilingaryan (KIT)	Slides
DLS Parallel HDF5 File Writing	Ulrik Pedersen (DLS)	Slides
HDF5 Parallel Reading and Tomography Processing at DLS	Mark Basham (DLS)	Slides
PaNdata ODI	Juan Bicarregui (STFC)	Slides
ICAT	Brian Matthews (STFC)	Slides
PaNdata Software catalogue	Jean-Francois Perrin (ILL)	Slides
Umbrella and PaNdata ODI Status report WP3	Stephan Egli (PSI)	Slides

PaNdata ODI Status report WP4	Milan Prica (Elettra)	Slides
PaNdata ODI Status report WP6	Brian Matthews (STFC) , Erica Yang (STFC)	Slides
PaNdata ODI Status report WP7	Jean-Francois Perrin (ILL)	Slides
Virtual Labs requirements and WP5	Thorsten Kracht (DESY)	Slides