

An online theoretical virtual observatory for hydrodynamical, cosmological simulations[☆]

Ragagnin, A.^{b,a,c,1}, Dolag, K.^{c,d}, Biffi, V.^{e,f}, Cadolle Bel, M.^{b,g}, Hammer, N. J.^a, Krukau, A.^{a,b}, Petkova, M.^{c,b}, and Steinborn, D.^a

^a*Leibniz-Rechenzentrum (LRZ), Boltzmannstrasse 1, D-85748 Garching, Germany*

^b*Excellence Cluster Universe, Boltzmannstrasse 2r, D-85748 Garching, Germany*

^c*Universitäts-Sternwarte, Fakultät für Physik, Ludwig-Maximilians Universität München, Scheinerstrasse 1, D-81679 München, Germany*

^d*Max-Planck-Institut für Astrophysik, Karl-Schwarzschild Strasse 1, D-85748 Garching bei München, Germany*

^e*Astronomy Unit, Department of Physics, University of Trieste, via Tiepolo 11, I-34131 Trieste, Italy*

^f*INAF, Osservatorio Astronomico di Trieste, via Tiepolo 11, I-34131 Trieste, Italy*

^g*Max Planck Computing and Data Facility (MPCDF), Gießenbachstrasse 2, D-85748 Garching, Germany*

Abstract

This article describes a virtual observatory hosting a web portal for accessing and sharing the output of large, cosmological, hydro-dynamical simulations with a broad scientific community. It also allows users to receive related scientific data products by directly processing the raw simulation data on a remote computing cluster.

The virtual observatory is a multi-layer structure: a web portal, a job control layer, a computing cluster and a HPC storage system. The outer layer enables users to choose an object from the simulations. Objects can be selected by visually inspecting 2D maps of the simulation data, by performing highly compounded and elaborated queries or graphically from plotting arbitrary combinations of properties. The user can apply several services to a chosen object. These services allow users to run analysis tools on the raw simulation data. The job control layer is responsible for handling and performing the analysis jobs, which are executed on a computing cluster. The inner most layer is formed by a HPC storage system which host the large, raw simulation data.

The virtual observatory provides the following services for the users: (I) CLUSTERINSPECT visualizes properties of member galaxies of a selected galaxy cluster; (II) SIMCUT returns the raw data of a sub-volume around a selected object from a simulation, containing all the original, hydro-dynamical quantities; (III) SMAC creates idealised 2D maps of various, physical quantities and observable of a selected object; (IV) PHOX generates virtual X-ray observations with specifications of various current and upcoming instruments.

Keywords: cosmology, galaxy clusters, online tools, web application, cloud computing

1. Introduction

Entering the so-called era of “precision cosmology” it becomes more and more clear that a theoretical counterpart in the form of very complex, hydrodynamical cosmological simulations is needed to interpret data from upcoming astronomical surveys and current instruments like PLANCK, South Pole Telescope (SPT), PanStars, Dark Energy Survey (DES), Euclid, LOFAR, eROSITA and many more. Such simulations follow the growth of galaxies and their associated components (like stellar population and central black hole) with their interplay with the large scale environment they are embedded in. Upcoming surveys will map large volumes of the Universe as well as record the birth of the first structures, especially galaxies and even progenitors of massive galaxy clusters at high redshift. In fact, their large potential of determining the nature of

dark matter and dark energy comes from being able to map the content and geometry of the Universe over most time in cosmic history. For theoretical models this means that simulations have to cover comparable large volumes, especially to host the rarest, most massive galaxy clusters expected to be the lighthouses of structure formation detectable at high redshifts. While the Universe makes its transition from dark matter dominated to dark energy dominated (i.e. accelerated expansion), the objects which form within it make their transition from young, dynamically active and star formation-driven systems to more relaxed and equilibrated systems observed at low redshifts. Those simulations study the internal evolution of clusters of galaxies with respect to the evolution of the cosmological background. They will be essential to interpret the outstanding discoveries expected from upcoming surveys.

However, running, storing and analysing such simulations is a challenging task, both from a technical as well as from a collaborative point of view. Recent generations

[☆]access via <https://c2papcosmosim.srv.lrz.de>.

¹ragagnin@lrz.de

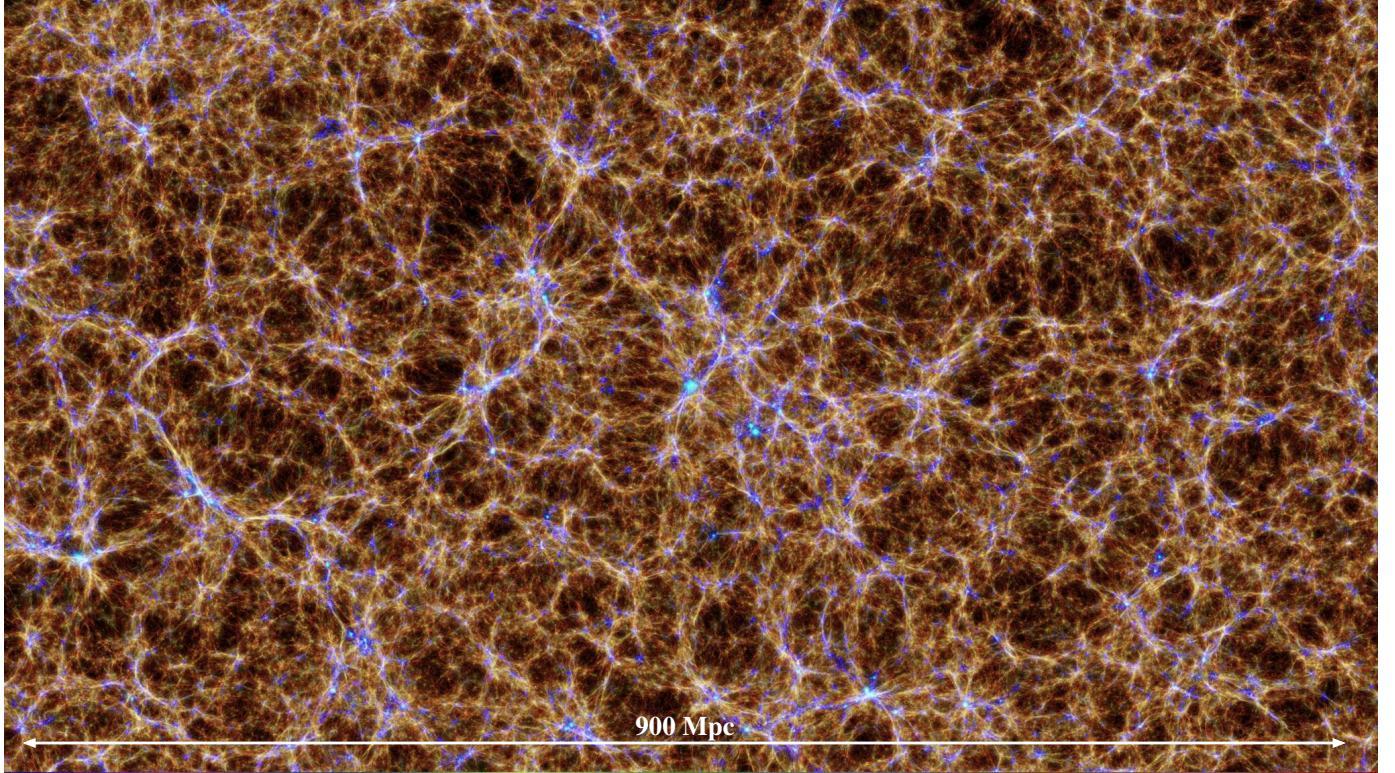


Figure 1: A visualisation of a cosmological large scale structure of the *Box2b/hr* simulation from the *Magneticum* project. This map shows diffuse baryons at $z = 0.2$, colour coded according to their temperature. The visualisation is centred on the most massive galaxy cluster in this simulation.

of HPC facilities provided within initiatives like GAUSS² or PRACE³ belong to the first generation of supercomputers which can perform cosmological, hydrodynamical simulations covering both the required large volume and high resolution requirements. Here, the largest simulation performed so far belongs to the *Magneticum* project⁴ and follows 2×4536^3 resolution elements over the whole, cosmological evolution of the universe (Bocquet et al., 2016). Such simulations model many more physical processes (e.g. star formation, cooling, winds, etc.) than the typical dark matter only counterparts used currently in computational cosmology. These simulations provide a larger set of complex data and can reach several hundreds of terabytes of raw data. Such simulations are performed within large collaborative efforts and results have to be shared with a broader scientific community. A guarantee for a deep scientific impact means that such data are made easily accessible and processable within the individual collaborating groups. It implies that data is stored on the HPC facilities for long periods of time, with the possibility to post-process the full data. In addition, it is important to make such data available to a large astrophysical community and allow the scientists to perform analysis tools via standard interfaces.

In this respect, efforts have been done in the recent years in order to share data sets of various kinds with the community. For instance, the Millennium Simulation Data Archive⁵ (Lemson and Virgo Consortium, 2006) is a pioneering work in this field. With the Millennium Simulation Data Archive, the user is able to compose SQL queries over substructure and merger-tree data in order to extract haloes and galaxies from the Millennium Simulation.

Users can also download the raw data files. The Cosmosim.org project⁶ allows users to compose additional queries over the list of particles and various post processed quantities (e.g. grid cells of density field). The Illustris Galaxies Observatory⁷ provides an application programming interface (API) where users can filter galaxies and download particle data from the Illustris simulations. The Australian Theoretical Virtual Observatory⁸ (Bernyk et al., 2016) is an online virtual laboratory where users can compose queries and run services on a selected objects (as producing mock observations or light cones).

Section 2 describes data of cosmological simulations and section 3 describes the currently available infrastructure. In section 4 we describe how users can interact

²<https://gauss-allianz.de/>

³<https://www.prace-ri.eu/>

⁴<http://www.magneticum.org>

⁵<http://wwwmpa.mpa-garching.mpg.de/Millennium/>

⁶<https://www.cosmosim.org/>

⁷http://www.illustris-project.org/galaxy_obs/

⁸<https://tao.asvo.org.au/tao/about/>

with the web interface and thereby compose science-driven queries to select objects. Section 5 describes the services currently implemented in the system.

2. The Simulations

In this section we explain in detail the simulation set from which a subset has currently made available in the virtual observatory.

2.1. The Magneticum Project

The Magneticum simulations⁹ (see Biffi et al. (2013); Saro et al. (2014); Hirschmann et al. (2014); Steinborn et al. (2015); Dolag et al. (2015b,a); Teklu et al. (2015); Steinborn et al. (2016); Bocquet et al. (2016); Remus et al. (2016)) follow the evolution of up to 2×10^{11} particles in a series of cosmological boxes ranging in size from $(50\text{Mpc})^3$ to $(4\text{Gpc})^3$. A visualisation of the second largest cosmological simulation can be seen in Figure 1. These simulations were used to interpret Sunyaev-Zeldovich (SZ) data from PLANCK (Planck Collaboration et al., 2013) and SPT (McDonald et al., 2014) as well as to predict cluster properties in X-rays for future missions such as Athena or Astro-H (Biffi et al., 2013). The first mock observations for the eROSITA cluster working group and the Athena+white book were also produced based on these simulations. Other scientific goals that were achieved with these simulations included studying the properties of the intra cluster medium (ICM) in galaxy clusters (Dolag et al., 2015b) as well as predicting the multi wavelength properties of the Active Galactic Nuclei (AGN) (Hirschmann et al., 2014; Steinborn et al., 2015). The large dynamical range probed by the combination of resolution levels and cosmological volumes also allowed us to calibrate the cosmological mass function based on hydro-dynamical simulations to a level required by future cosmological probes (Bocquet et al., 2016). The detailed treatment of all relevant physical processes allowed us to investigate dynamical properties of galaxies based on morphological classification (Teklu et al., 2015; Remus et al., 2016) for the first time. A small subset of these simulations also follows the evolution of magnetic fields.

The web portal currently hosts a subset of the full Magneticum simulation set. It hosts 28 outputs of a medium size simulation *Box2/hr*, which utilise 0.8×10^{10} particles, covering a volume of $(500\text{Mpc})^3$ as well as 11 outputs of a larger size simulation; *Box2b/hr*, which utilise 5×10^{10} particles, covering a volume of $(900\text{Mpc})^3$. For each cluster contained in the simulated volumes, the web portal stores a set of pre-computed quantities. The set of pre-computed quantities is chosen to let users select objects in categories (e.g. fossils or compact objects) that are widely studied.

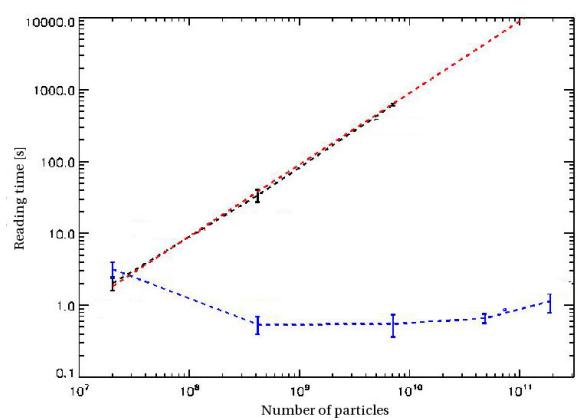


Figure 2: Wall clock time spent reading all data for the most massive galaxy cluster from a snapshot as function of the total number of particles for increasing simulation size. In black line there is the brute force approach by reading all data while the blue line is the timing of the improved algorithm. The improved algorithm does a spacial selection of the snapshot by use of key-index files. Those files allow to readout only the relevant part of the snapshot files. The percentage at the individual data points indicate the fraction of particles to be read compared to the overall particle number.

2.2. The Metadata

We use SUBFIND (Springel et al., 2001; Dolag et al., 2009) to define properties of haloes and their sub-haloes. SUBFIND identifies substructures as locally overdense, gravitationally bound groups of particles. SUBFIND starts with a halo list identified through the Friends-of-Friends algorithm. For each halo and for each of its particle the local density is estimated using adaptive kernel estimation with a prescribed number of smoothing neighbours. Starting from isolated density peaks, additional particles are added in sequence of decreasing density. When a sequence of particles contains a saddle point in the global density field that connects two disjoint over-dense regions, the smaller structure is treated as a substructure candidate, followed by merging the two regions. All substructure candidates are subjected to an iterative unbinding procedure with a tree-based calculation of the potential. These structures can then be associated with galaxy clusters, galaxy groups and galaxies and their integrated properties (like gas mass, stellar mass or star-formation rate) can be calculated.

2.3. Raw Simulation Data Layout

For the *Magneticum* project, we developed a special structured output for very large scale, n-body, cosmological, hydrodynamical simulations. It employs a special space filling curve to produce auxiliary data which allows a fast and direct access to spacial regions within the output of smoothed particle hydrodynamics (SPH) simulation. This indexing scheme was extended to work with the multi-file and multi-component output of such SPH simulations. To optimize the access to individual data, some files contain data regarding the content and length of individual data blocks.

⁹<http://www.magneticum.org>

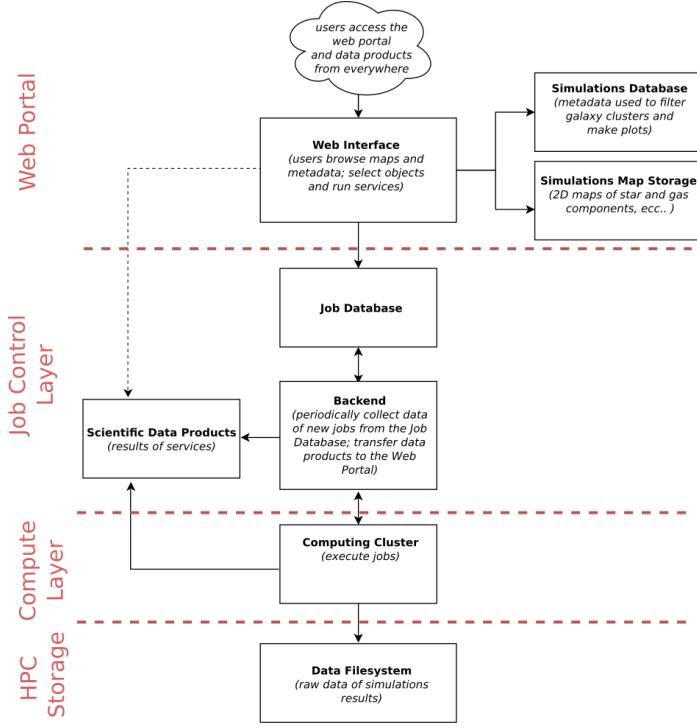


Figure 3: Schema summarising how the processes are distributed to different parts of the existing infrastructure, and how the data flows are within the fully operational prototype of the web portal.

The creation, reading and comparing of the index list add only a negligible overhead. Figure 2 shows that the reading of all particles within the virial radius of the most massive galaxy cluster in any of the simulations takes significantly less than 1 second. This is even the case for the Magneticum *Box0/mr* simulation. This simulation utilizes almost 2×10^{11} particles and the information have to be filtered out of the individual snapshots, which for one time instance occupy 20TB on disk.

In Appendix A we explain in more detail how to build and read those auxiliary data files in order to add third-party simulations or additional services to the virtual observatory.

3. Structure of the Virtual Observatory

Figure 3 illustrates our multi-layer structure (the different layers are separated by a dashed red line). Between those layers, data and processes flow over the virtual observatory, the database, the job control layer, the computing cluster (where the analysis tools are actually executed), and the storage system (where the raw simulation data are stored). The need for a separation between the web interface and the backend, arises from both the necessity of users to run personalized jobs on raw data, managed by a job scheduler of the computing cluster and the protection of the data from unauthorized access.

In Appendix B we show how to prepare data of a simulation in order to add it to the virtual observatory.

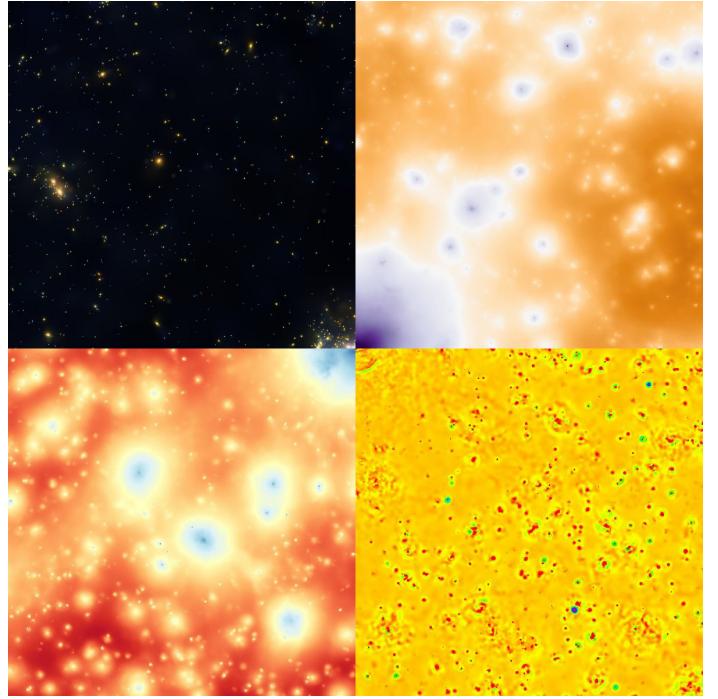


Figure 4: The four quadrants, centred on a merging cluster, report: the stellar component (*Stars*, upper left), the ICM pressure (*ComptonY*, upper right), the ICM X-ray emission (*ICM*, lower left) and the pressure fluctuations (*Shocks*, lower right).

3.1. Overview of the multi-layer architecture

The user selects a service from the web interface and this information is written into a job database¹⁰ (which in our case is implemented as a separate independent instance). The backend is triggered by the job database and will configure and submit jobs to a computing cluster which will execute them. Once a job is added, a trigger in the database will make the backend send the job to the computing cluster. Finally, the backend delivers the according data products to the user via a download link which is valid for a limited time. The computing cluster must have access to the HPC file system where the simulation data are stored, however, it does not need to store the data locally.

Almost all parts are based on common packages and available libraries except the core of the backend, which is a customized component tailored for the data flows and job requests to the specific needs of the scientific software performing the data analysis.

3.2. The Web Portal

The outer layer of the diagram in Figure 3 is the web portal. The web portal consists of two main parts, the web interface¹¹ and the simulations database. While this

¹⁰In our implementation we used PostgreSQL 9.4.6 for all the different databases

¹¹Our web server is built using the Python micro framework FLASK 0.11

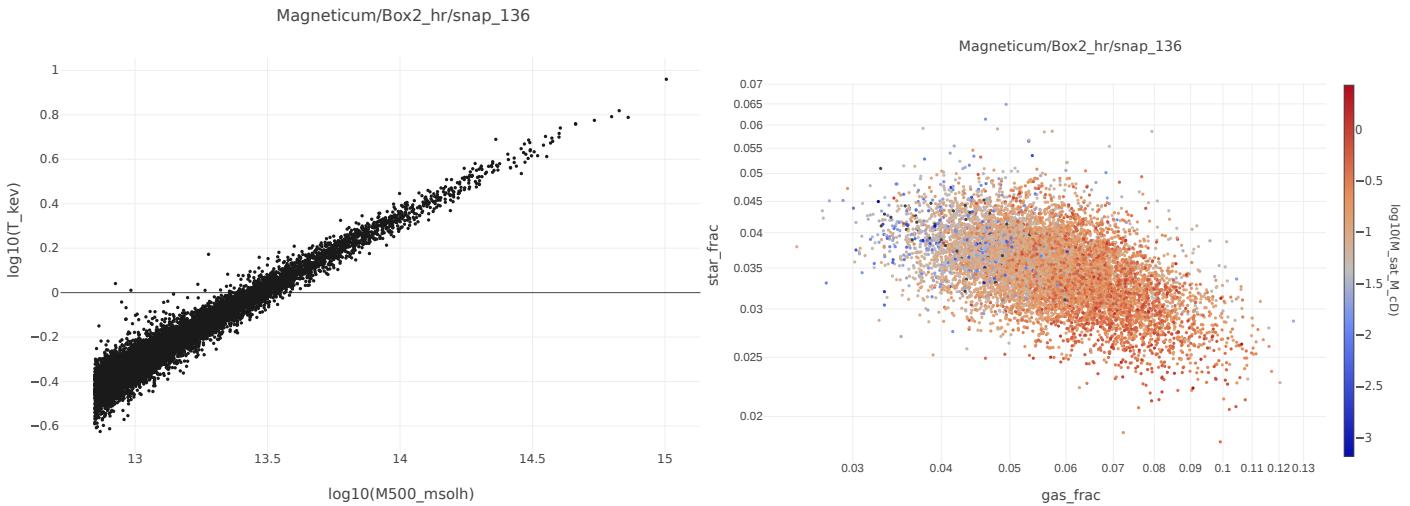


Figure 5: Mass-temperature correlation from the metadata of the *Magneticum/Box2/hr* simulation at $z = 0.14$ (left panel), where outliers can be identified and then be found by selecting their properties in the *Restrict* window. The right shows the anti correlation of stellar and gas mass fraction, coloured by the stellar mass ratio of satellite galaxies to central galaxy, which often is used as indicator for the dynamical state.

database hosts the metadata of all the simulations, the web interface allows users to select objects in various ways through a graphical interface. This interface makes possible to visualise 2D maps¹² and metadata queries. Users can navigate through different maps, as shown in figure 4 by scrolling and zooming. They can perform complex queries on simulation metadata based on object definitions (in our case based on SUBFIND) which are stored in a database. The CLUSTERFIND tool allows users to obtain all properties of clusters fulfilling the restriction criteria of the query mask. It is then possible to download these results as CSV-tables or to interactively visualize them. The user can thereby make scatter plots, 1D and 2D histograms on arbitrary expressions of the properties within the table. To account for the possibility of selecting a particularly interesting cluster, the user can either select an object from the table or directly from the scatter plots by clicking on the according data point. An example of a scatter plot between the mass and temperature is given in figure 5 where the user for example could select the cluster that clearly lies outside the mass-temperature relation. Once selected, the user can also switch to a service (if available for the current snapshot) and keep the current object as the selected one.

Currently, four services are available: CLUSTERINSPECT, SIMCUT, SMAC and PHOX, which are described in more detail in section 5.

The web interface will check the job database for the status of a submitted job and return results to the user.

Once the job is finished, the backend provides the web interface with a web link, where the user can download the results. The user is able to check the status of all jobs and to download results via the web interface.

Metadata from galaxies and galaxy clusters (identified with SUBFIND) are stored in a database on the web portal. Maps can be additionally overlaid. The virtual observatory is able to host and handle large simulations, with order of 10^5 objects, each containing a number of galaxy members. Note that a large galaxy cluster in the simulation can have thousands of individual galaxies as associated data. This results in a huge number of rows in the database. Therefore, to achieve high performance, we used the so called Database Indexing. Database Indexing is a solution provided by databases for speeding up queries on read-only tables. In details, Database Indexing creates a tree structure over the field of a table so that for every indexed field there is a look-up table where data are ordered according to this field. This speeds up every query where data is selected on a specific range of values (e.g. range of mass or other global quantity of galaxies and galaxy clusters). We also use a small storage system directly on the web portal in order to store *zoomify* versions of the various and large 2D maps of the simulations. The end user can therefore browse quickly through these vast simulation maps and visually seek for objects of interest.

3.3. Job Control Layer

The backend is based on Python and is activated when a new job is sent to the database or when the job scheduling system start/finish a job. New jobs are set up by preparing an unique job directory and template scripts and parameter files for the different services are used. The job is submitted to the queue of the computing cluster and executed on the cluster. Once the job is completed, the backend will collect the results and communicate a link for downloading the data to the web interface via the shared job database. Here the job database is the conjunction ring between the web portal and the backend, as can be

¹²Maps are rendered using the OPENLAYERS 3 JavaScript library

seen in figure 3. As mentioned before, for security reasons, the web interface can not directly send jobs to the compute cluster or access the raw simulation data.

3.4. Compute Layer and HPC storage

Large, hydrodynamical cosmological simulations nowadays are performed within large, collaborative efforts. While typically a large number of scientists contributes to the development of the underlying code and its auxiliary functionality needed to perform specific simulations, the actual execution of individual simulations is typically performed by a small number of directly involved scientists. It is practically impossible to grant all scientists of the collaboration (or beyond) direct access to the simulation data on the HPC facility. Therefore we follow a different concept. While the storage file system is typically assigned to the individual scientists within their HPC projects, they give reading permission of the raw simulation data to a single user of the super computing center. This can be done, as in our case, even on a much smaller, dedicated computing cluster where HPC storage system is available on the compute nodes. The execution of such jobs will be typically done via a batch system. An additional, independent, local storage keeps the data products. This local storage is represented by the box *Scientific Data Products* in figure 3. The data products of the services will be made available for the end user on the this local storage, and can be shared with the scientific community without restrictions. Within this concept, neither the web interface nor the actual user ever has (or needs) any direct access to the raw simulation data.

3.5. Implementation

Currently, the full outcome of the two Magneticum simulations is stored on a Big Data storage attached to a HPC system (in our case, it is a project space at the Leibniz Rechenzentrum). Analysis jobs run on a separate, much smaller scale computing system. Specifically, we employ to this purpose the computing cluster at the Computational Center for Particle and Astrophysics (C²PAP)¹³.

The web interface is running on virtual machines hosted at LRZ¹⁴. The database of metadata properties (cluster and galaxy properties) runs on the same virtual machines of the web interface. Users log in and register by using their e-mail addresses. User registration, encryption and reset of passwords are handled by the Flask-Login python module¹⁵.

Users must register and be approved by administrators of the virtual observatory to access its data. Users register to the web portal using their email address as username inside the web interface. The administrator can then grant *roles* to the individual users. Although at the moment

¹³<http://www.universe-cluster.de/c2pap/>

¹⁴<https://www.lrz.de/services/serverbetrieb/>

¹⁵<https://flask-login.readthedocs.io/en/latest/>

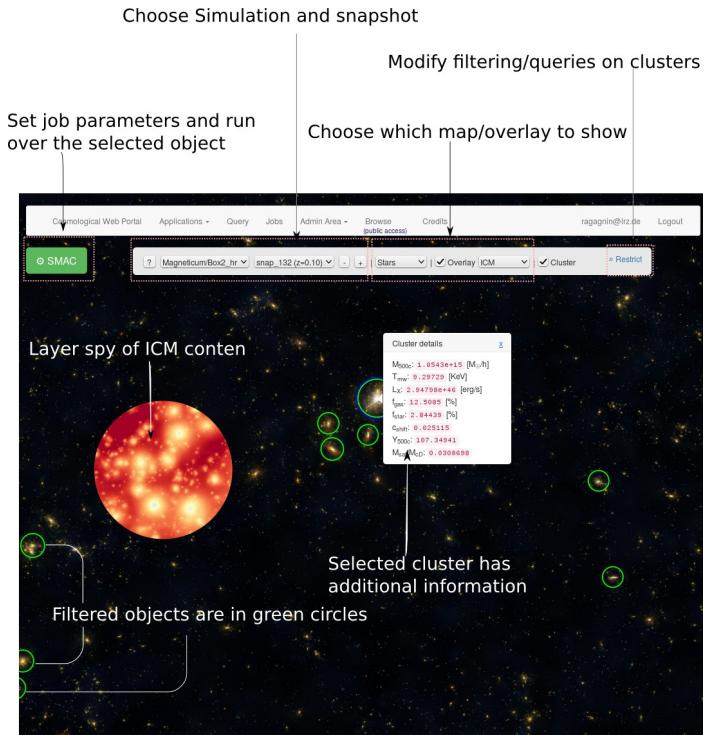


Figure 6: Overview of a given snapshot of a simulation. The layer-spy can be activated to switch from the stellar to the ICM view within a small region around the cursor.

all services are opened for all registered users, the implementation allows the administrator to make services only accessible for users with special roles. Moreover, there is the need to share data products with a wider community, where members do not necessarily have access to the system. When a job is created through the web interface by a registered user, a link containing a unique identifier of the job is delivered to the user. As mentioned above, the link can be shared, and its access does not require registration.

To allow a first exploration of the simulations data, there is the possibility to visually explore the simulations without being registered. However, the access to the meta data is limited and all services are disabled in the public browsing mode.

Jobs that are submitted through the web interface are all sent to the computing cluster.

Data products are stored on the additional FTP server which runs on the C²PAP computing cluster. Data products are guaranteed to be stored for 30 days. The data portal and its infra structure is will be available for a minimum of five years.

4. Exploring Simulations

Here we describe in detail how the web portal¹⁶ allows users to explore the simulations and how objects of interest

¹⁶ The web portal is a dynamic website and the interface is built using ANGULARJS v1.4.5, JQUERY v1.11.3, JQUERY UI v1.11.3 and BOOTSTRAP v3.3.4

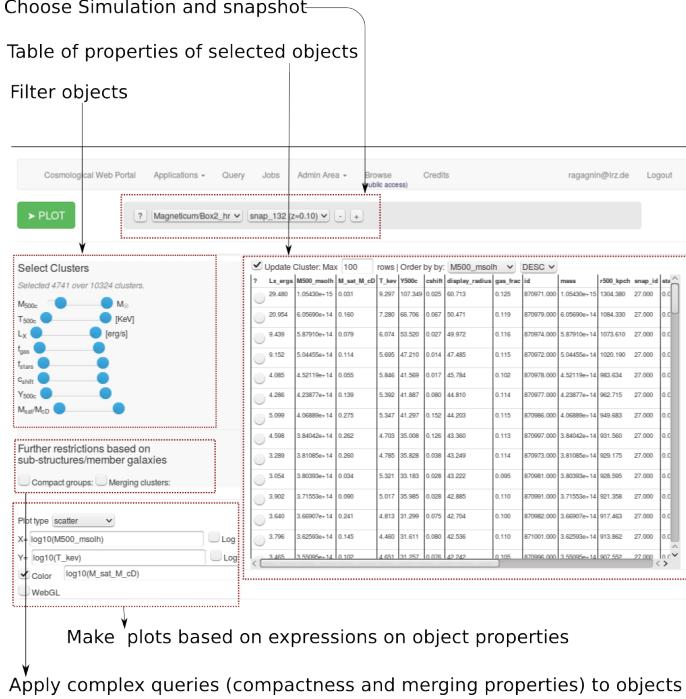


Figure 7: A screenshot of the service CLUSTERFIND. Here there is the selection panel for galaxy clusters and groups, where the objects can be filtered according to their properties.

can be interactively found by performing complex queries using *Restrict* or plot the metadata quantities using *ClusterFind*.

4.1. Selecting Objects by browsing maps

The web interface allows users to explore cosmological structures within the simulations by panning and zooming through into high resolution (256 mega pixel) images. Once a simulation is chosen, the output at various points of the cosmological history can be selected. Depending on the underlying simulation, typically between 10 and 40 different output times can be chosen. Generally, four different maps can be selected as the prime visualization, as shown already in figure 4. The diffuse baryonic medium is either visualized colour coded by it's X-ray emission (*ICM*) or by its pressure (*Compton Y*) using SMAC (Dolag et al., 2005). The stellar component is visualized by the density of stars and colour coded by the average age of the stellar population using SPOTCH (Dolag et al., 2008). Additionally, we computed a filtered visualization of the ICM pressure (it Shocks), where shocks and turbulence are visible. In figure 4 each of the segments show one of the visuals mentioned before and demonstrate the different appearance of the massive galaxy cluster in the centre. The arc-let (depicted by the green curved line over the yellow background) appearing near the centre in the *Shocks* visual (lower right segment) resembles a shock wave which spans more than 1Mpc in size and indicates a merging cluster. All visualizations are based on 16k×16k pixel sized images which can be explored using the zoomify technique. In addition, the

layer-spy can be activated to switch to a different visualization view within a small region around the cursor, (see figure 6). This immediately gives a visual impression of the dynamics and composition of the diffuse, gaseous and the stellar component within our universe. It also allows users to instantaneously see various, physical features of the simulation. For example, galaxies in less dense environment appear more often in a bluish colour, indicating a young stellar component while galaxies in more dense environment often appear in yellow and red colours, indicating an older stellar component. This reflects that the underlying simulations correctly reproduce the so called morphological density relation of galaxies, which is one of the most prominent, observed imprints of the large scale structure of galaxy formation. Additionally, the position of galaxy clusters and groups can be overlaid as green circles and an information panel on the cluster properties is visible as soon as a galaxy cluster is selected.

4.2. Composing Queries

To select special galaxy clusters and groups, the *Restrict* window can be used to perform complex queries among the metadata of clusters and groups, as shown in Figure 7. Every time a value is modified by the sliders, a new database query is performed and the selected objects are shown as green circles in the web interface.

4.2.1. Restrict sample by Value

The upper row of sliders allows users to choose minimum and maximum value for various global quantities, like mass (M_{500c}), temperature (T), bolometric X-ray luminosity (L_x), gas and stellar fractions (f_{gas} and f_{stars} respectively). The results also can be displayed or downloaded as ASCII tables. This allows the user to perform simple analysis, for example plotting scaling relations like the well known mass-temperature relation (shown in in the left panel of figure 5). For example, a prominent outlier in the figure can then be selected in the web interface by restricting $M_{500c} > 14.1$ and $T > 3.05$. The resulting cluster will then be the only marked cluster. By a closer inspection, it can be recognized as a major merger system, already visible in both visualization, as double emission peak in the diffuse medium as well as second, very large infalling group of galaxies in the stellar component.

4.2.2. Select by Dynamical State

There are two classical measures (Rasia et al., 2013; Jones et al., 2003; Cui et al., 2011) of the dynamical state of galaxy clusters and groups. One is the ratio of the total stellar mass in satellite galaxies with respect to the mass in the central galaxy. The other is the so called centre shift, which is a measure of the displacement of the “centre of emission” of the diffuse baryons compared to the position of the potential minimum and is typically measured in units of the size (e.g. R_{500c}) of the system. These two measurements can be used to select relaxed and

unrelaxed systems. For example, the outlier in the above example shows a very large centre shift of $c_{shift} = 0.13$ and a large stellar satellite fraction of $M_{sat}/M_{cD} = 0.7$, clearly classifying it as major merger. Usage of the metadata tables allows users also to verify other correlations, like the anti correlation between stellar mass and gas mass fraction as shown in the right panel of figure 5. In that case we used one of the merger indicators to colour code the objects across the anti-correlation.

4.2.3. Finding different Classes

As seen in the examples before, the different filters allow users to select objects with different values or with different dynamical states. However, such filters can also be used to select different object classes.

Fossil groups for example have typically a very large, dominant central galaxy and only a very small amount of satellite galaxies. Such objects can be selected via the M_{sat}/M_{cD} parameter, which is the mass ratio between the sum of all satellite galaxies and the cD galaxy. In this case, choosing a small value in will select fossil groups.

Compact groups can be selected by marking the according check-box in the *Restrict* window. If done, an additional query form appears. This query is more complex than the ones before because it involves constraints on galaxy members of each cluster. In this case, all member galaxies of a group or cluster within the chosen radius (100 kpc in the example shown in figure 7) are first ranked by mass. Then only objects where the Nth galaxy (the 4th in the example in figure 7) has a stellar mass of at least 10% of the central galaxy (selecting $\log_{10}(M_{cD}/M_{Nth}) < 1$ in the example in figure 7) are selected. Without any further restrictions there are 142 compact group candidates out of 7428 objects fulfilling this criteria at $z = 0.67$ within our example of *Magneticum/Box2/hr*.

Merging clusters can be also selected by marking the corresponding check-box in the *Restrict* window. This query allows to select clusters or groups where at least one sub-structure fulfils the given criteria. The user can select the range of stellar and gas mass content, the relative, tangential or radial velocity and a distance of the sub structure to the centre. To find bullet cluster like systems, one would select a large stellar and gas mass, a large, outgoing velocity (v_r positive !) and a recent distance from the centre.

5. The Services

The results of the different services typically come with different number of files, depending on the configuration of the workflow specified by the user. Therefore, the results are made available in form of a tar ball, which the user can download. Additionally, the web interface provides small iPython examples for every available result. They can be used to have a quick look at the obtained results.

Block	Type	Size	Blocks
POS	FLOATN	3	0,1,4,5
VEL	FLOATN	3	0,1,4,5
ID	LLONG	1	0,1,4,5
MASS	FLOAT	1	0,1,4,5
U	FLOAT	1	0
RHO	FLOAT	1	0
HSML	FLOAT	1	0
SFR	FLOAT	1	0
AGE	FLOAT	1	4,5
BHMA	FLOAT	1	5
BHMD	FLOAT	1	5
BHPC	LONG	1	5
iM	FLOAT	1	4
Zs	FLOATN	11	0,4
CLDX	FLOAT	1	0
TEMP	FLOAT	1	0

Table 1: Data block names, type and sizes currently produced by SIMCUT.

5.1. ClusterInspect

CLUSTERINSPECT allows users to browse the member galaxies of the selected cluster. As in CLUSTERFIND, it is then possible to make plots on the properties of the member galaxies¹⁷.

5.2. SimCut

The SIMCUT service allows users to create artificial GADGET snapshot files. They are produced following the same format of P-GADGET2.¹⁸ The file produced by the SIMCUT service contains a list of particles within a region centred on the selected galaxy cluster or the group. For every quantity that is stored (e.g. position, velocity, mass), these files have a so called “block”. Every block has a code name (respectively, POS, VEL) composed of a string of maximum of 4 characters and an array of the size of the number of particles. In addition, the snapshot contains an extension to the P-GADGET2 output files the so called *INFO* block. This block contains information on data types and dimensions of the different values stored in the snapshot, as shown in table 1.

5.3. SMAC

The SMAC service allows users to construct maps from the simulations using the map making program SMAC (Dolag et al., 2005), which allows to integrate various physical and observational quantities throughout the simulation. Once a galaxy cluster or group is chosen, the service allows the user to select various different map making options:

¹⁷Within the CLUSTERFIND and CLUSTERINSPECT service, interactive plots are made using the JavaScript version of the library PLOTLY.JS v1.9.0

¹⁸For more details about the Gadget formats and parameter file configurations, see: <https://wwwmpa.mpa-garching.mpg.de/gadget/users-guide.pdf>.

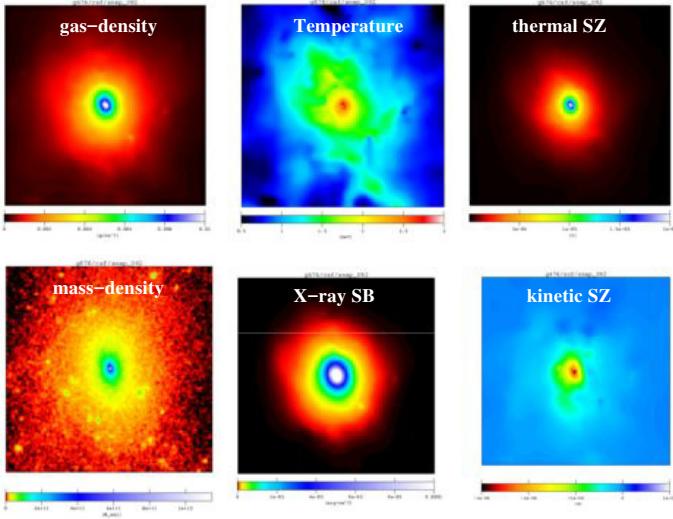


Figure 8: A sample of maps currently available within the SMAC service for a given selected cluster.

- *Baryonic density map [g/cm²]*
- *Total mass density map [M_⊙/cm²]*
- *Mass-weighted Temperature [keV]*
- *Bolometric, X-ray surface brightness [erg/cm²]*
- *thermal SZ effect [Compton Y-parameter]*
- *kinetic SZ effect [Compton w-parameter]*

The size of the image as well as the integration depth along the z-axis can be chosen. The data are returned as standard FITS¹⁹ files containing an image. Figure 8 shows some example maps obtained by this service.

5.4. PHOX

PHOX is a virtual X-ray telescope devoted to obtaining X-ray synthetic observations from hydrodynamical, numerical simulation outputs (Biffi et al., 2012, 2013). As a first step, the simulation box is converted into a virtual box of photon packages, generated by the sampling of the X-ray emission spectra calculated for each gas element as well as for each AGN within the simulation.

5.4.1. The ICM emission

In order to compute the X-ray emission spectrum of each gas element, the APEC model (Smith et al., 2001) for the emission of a collisionally-ionized, chemically enriched plasma implemented within the external, publicly available package XSPEC²⁰ Arnaud (1996) is utilized. The properties (namely density, temperature and metallicity,

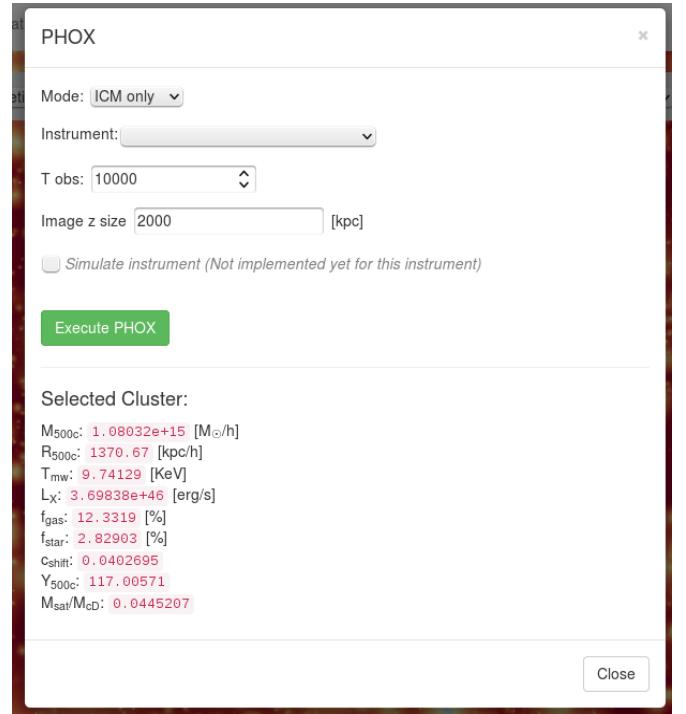


Figure 9: Selection panel for executing PHOX, where details on the virtual observation and the instrument properties can be selected

and even chemical composition) of the gas elements are directly used as input for this model. Here, no external libraries of model spectra are constructed, but the emission spectrum generated by XSPEC is processed immediately, gas element per gas element, which allows users to vary the amount of parameters involved to describe the spectral emission without enormously increasing the computational effort and memory requirements. Each spectrum is then populated with a predicted number of photons (according to fiducial, general values for collecting area and exposure time). The photons collected from all the X-ray-emitting gas elements are eventually stored in terms of photon packages, each of them being characterized by the position and velocity of the original emitting element and by the array of energies of the associated photon sample.

5.4.2. The AGN emission

For the AGN component (see Biffi et al. 2016, in prep.), the procedure followed to convert the simulation box into a box of ideal photons emitted by all the AGN sources is similar to the one used for the gas, except for the spectral model utilized. Namely, we model the AGN emission with an intrinsically absorbed power law, constrained as follows. We convert the bolometric luminosities L_{bol} of the AGN into rest-frame SXR and HXR luminosities assuming the bolometric corrections proposed by (Marconi et al., 2004), which can be approximated by the following third-degree polynomial fits

$$\log(L_{hxr}/L_{bol}) = -1.54 - 0.24\mathcal{L} - 0.012\mathcal{L}^2 + 0.0015\mathcal{L}^3$$

$$\log(L_{sxr}/L_{bol}) = -1.65 - 0.22\mathcal{L} - 0.012\mathcal{L}^2 + 0.0015\mathcal{L}^3$$

¹⁹<http://fits.gsfc.nasa.gov/>

²⁰<https://heasarc.gsfc.nasa.gov/xanadu/xspec/>.

with $\mathcal{L} = \log(L_{bol}/L_\odot) - 12$, derived for the range $8.5 < \log(L_{bol}/L_\odot) < 13$. (see fig. 3(b) in Marconi et al. (2004)). Here, we mimic the observed scatter in these relations by adding a random scatter of 0.1 to the SXR and HXR luminosities, in logarithmic scale. Then we construct an intrinsic redshifted power law spectrum

$$A(E) = [K(1+z)][E(1+z)]^{-\alpha} \left(\frac{1}{1+z} \right), \quad (1)$$

where K is the normalization constant, α is the photon index and z is the redshift of the source. For every AGN in the simulation, the two parameters K and α can be constrained by integrating the spectrum from the expected values of L_{SXR} and L_{HXR} . In this approach, the obtained distribution of photon indexes α reasonably reflects the observed Gaussian distribution in the range 1.2–2.8, which peaks around ~ 2 (e.g. Zdziarski et al., 2000)

Many observational works suggest that AGN sources also show evidences for the presence of obscuring material (i.e. the torus) in the vicinity of the central BH, which leads to the partial absorption of the emitted radiation. In order to account for this phenomenon, we assign to each AGN in the simulation a value of the obscurer column-density (N_h) by assuming the estimated column-density distribution presented in the study by Buchner et al. (2014) (see top-left panel of fig. 10, in their paper) and based on a sample of 350 X-ray AGN in the 4Ms Chandra Deep Field South. Within the PHOX code, we include this in the construction procedure of the X-ray emission model from AGN sources. The resulting absorbed SXR and HXR luminosity functions are found to be in overall good agreement with those observed (see Biffi et al. 2016, in prep).

5.4.3. Performing an X-ray Observation

With this service, the field of the selected photon packages (either only “ICM”/“AGN” or both can be selected) is further processed by taking into account the geometry of the mock study and the *idealized* instrument characteristics. In particular, given the ideal cube of virtual photon packages associated to the simulation output, it is here possible to select the size of the subregion of interest, centred on the selected galaxy cluster or group, and to choose an instrument, which defines a field of view and a nominal effective area of the X-ray telescope, as shown in Figure 9. This also defines the line-of-sight (l.o.s.) direction and accounts for the according Doppler shift in the photon energies due to the peculiar motion of the original emitting particles. The virtual observation then returns the photon list expected for the idealized instrument chosen (e.g. assuming the nominal effective area over all energies, no beam smearing of the position and keeping the exact energy) and returning the data in the form of a FITS file in SIMPUT²¹ format, which allows users to directly utilise this

file for more sophisticated instrument simulations. The time of the virtual observation is restricted, and a limit of $A_{instrument} \times T_{obs} < 10^9$ is enforced. This process typically takes only few seconds to be performed.

5.4.4. Including the mock X-ray observation

We also installed some publicly available instrument simulators, which can be added by clicking on the *Simulate Instrument* check-box, for some of the chosen instruments.

For *XMM*(EPIC), *eROSITA* and *Athena*(XIFU/WFI) we are using SIXTE²², whereas for *SUZAKU*(XIS) we are using XISSIM, which is part of the HEADAS package (Ishisaki et al., 2007). For *XraySurveyor*(HDXI/XCAL) and *Hitomi* (SXS/SXI/HXI) we are using SIMX²³ while for *Chandra* (ACIS-S/I) instruments we are using MARX (Davis et al., 2012). Depending on the instrument and time chosen for the observation, this process can take up to several minutes. It returns an event file which can be then analyzed with standard X-ray analysis tools.

Figure 10 shows the result for observations centred on the same massive galaxy cluster at $z = 0.3$ while choosing different instruments. The exposure time is set to 100ks (reduced to 50ks for *Athena*) and a slab of 100Mpc along the line of sight is used. Both, the ICM as well as the AGN component are taken into account in this example. For the *eROSITA* case, a simulation of all 7 CCD chips is preformed, delivering 7 event files, which are additionally combined to one single event file, which leads to the appearance of edges from the rotated frames in the combined image as seen in the map. For the *Athena*(WFI) simulation, the 4 individual chips are simulated, as still visible by the gaps in the maps. However, as for real observations, the telescope is made to dither during the integration time to smears out the gaps between the chips, which is controlled by a realistic attitude file which defines the pointing of the instrument.

In Appendix C we show configuration files for SIXTE, XISSIM, SIMX and MARX used for each instrument we included.

6. Conclusions

In this work we present an online theoretical virtual observatory based on a multi-layer infrastructure for large cosmological hydro-dynamical simulations. It will give a wide scientific community the possibility to perform analysis tools on data from several large simulations.

The increasing amount of upcoming astronomical surveys makes it necessary to increase resolution and volume of cosmological simulations to interpret the results from such observational campaigns. These simulations may include different matter contents and object types (gas, dark matter, stars, black holes) and various physical processes

²¹<http://hea-www.harvard.edu/heasarc/formats/simput-1.1.0.pdf>.

²²<http://www.sternwarte.uni-erlangen.de/research/sixte>.

²³<http://constellation.gsfc.nasa.gov/science/simulatorInstall.html>

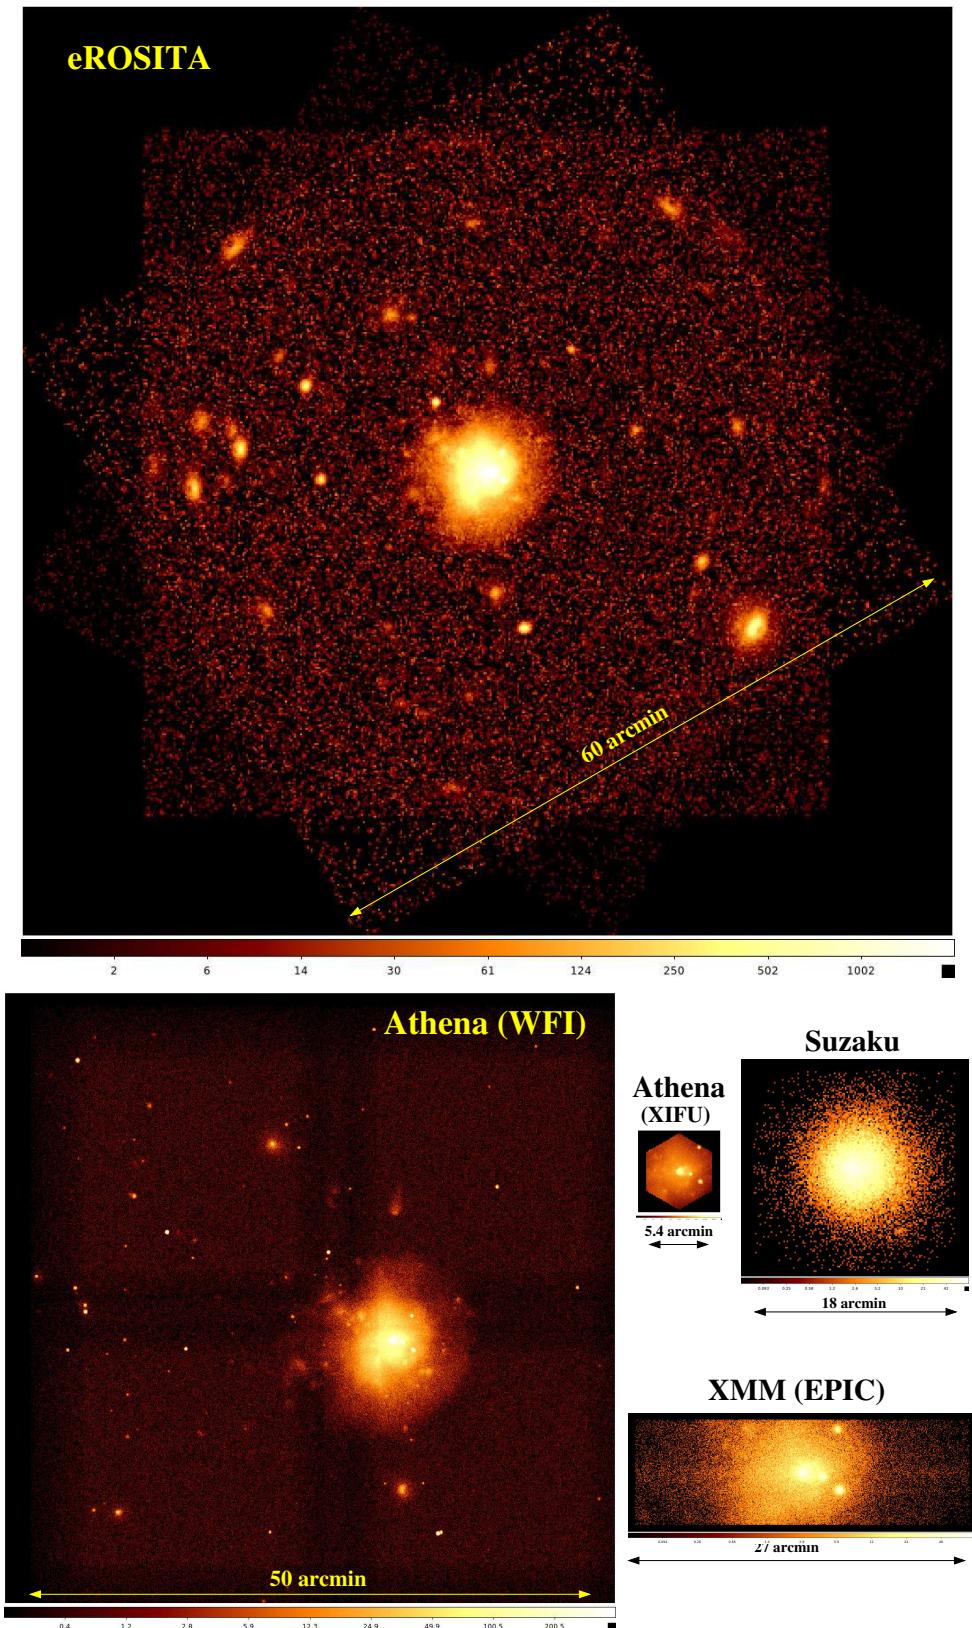


Figure 10: Observation of the X-ray emission obtained from the combined ICM and AGN contribution, centred on the most massive cluster at $z = 0.3$, performed with different, current and future X-ray instruments, including the actual instrument simulations.

that shapes the formation of objects and structures in our universe. All those factors dramatically increase the output size of current simulations.

As a result, the storage and sharing of the largest simulations that are currently available represents a real challenge.

For example, the size of an individual snapshot of the *Box0/mr* simulation of the *Magneticum* project (which follows almost 2×10^{11} particles), is more than 20TB.

Although we currently only store results from the *Magneticum* project, the virtual observatory infrastructure is capable of including other simulations as well.

This virtual observatory allows users to configure workflows to run individual analysis tools for specific objects of a given snapshot on the raw simulation data. A web interface helps the user to build compounded queries on metadata of the different simulations, which in our case are galaxies and galaxy clusters as obtained by SUBFIND. This allows users to select a target object in the different simulations with the desired properties. Those queries can be restrictions to global properties like mass or other observables, as well as complex queries on various internal aspects. This gives the possibility to select different general classes of objects (like merging clusters, compact groups or fossil groups). Our initial version of the virtual observatory provides so far four services: CLUSTERINSPECT plots and shows data of member galaxies of a given galaxy cluster, SIMCUT gives the possibility of downloading the raw simulation data of all particles belonging to the selected object, SMAC creates 2D maps of different physical or observational quantities, PHOX generates virtual X-ray observations.

A number of considerations about accessibility, security and performance led us to a multi-layer infrastructure: the web portal, the job control layer, a computing cluster and the HPC storage system. The outer layer lets people perform highly compounded and elaborated queries on SUBFIND results; the computing cluster runs the chosen job, reading the raw simulation data directly via access to the HPC storage system and sends the results back to the web portal.

Concerning the cluster selection, there is the possibility of interlocking a number of different subqueries. Some very basic queries are filtering by mass, temperature, bolometric X-ray luminosity, fraction of stars and gas masses. The user can also select an object by browsing 2D maps of the simulation, or by clicking objects in scatter plots of SUBFIND data.

It is possible to run queries for the selection of clusters in different dynamical states. For instance depending on their satellite fraction and displacements between baryon and potential centre. Other subqueries allow for the selection of clusters with different degree of compactness by choosing the ratio between masses of the central galaxy and the n-th satellite galaxy within a given distance.

Services themselves come with different parameters that the user is free to choose. SMAC can project onto three dif-

ferent axes and map different matter properties (e.g. density, temperature, X-ray temperature); PHOX can simulate a number of different instruments (e.g. XMM, Suzaku, Chandra, eROSITA, Athena, XraySurveyor, Hitomi) and compute X-ray emission due to AGN or ICM, separately or together.

Services will be extended in the future to incorporate additional virtual observatories. Possibility to browse observationally motivated realizations of light-cones is foreseen. Finally, the virtual observatory will permanently grow in size and amount of simulation data which will be made available for the general scientific community. Although some modifications on the different components would be needed to adapt the concept to other infrastructures and simulations, we are happy to provide the underlying source code on request and to give advice for adapting the package to other institutions. The current infrastructure for the virtual observatory is secured for five years.

7. Credits

The results obtained via the virtual observatory are free to use. Please acknowledge the *Magneticum Project*²⁴ (Dolag et al. 2016, in prep) and the *Cosmological Web Portal* (this work, Ragagnin et al. 2016). Please also give the references for the individual simulations used (Box2/hr: Hirschmann et al. (2014), Box0/mr: Bocquet et al. (2016)). If results produced with PHOX are used, please cite Biffi et al. (2012), and in case of SMAC please cite Dolag et al. (2005). In the case that results of the X-ray instrument simulators are used, please refer to SIXTE²⁵, XISSIM (Ishisaki et al., 2007), SIMX²⁶ and MARX (Davis et al., 2012), depending on the simulator used for the simulation.

Acknowledgements

We thank G. Lemson for various discussions and his contribution to various VIO proto-types, which largely inspired this work. We thank T. Dauser, G. Lamer, N. Clerc, A. Rau and E. Pointecouteau for helping with the X-ray data and especially with the instrument simulators, which can be additionally involved within the “PHOX” service. AR thanks Nathan James Turner for carefully checking the english grammar of this manuscript. We also thank Ben Hoyle for carefully reading the manuscript. This work has been supported by the DFG Cluster of Excellence “Origin and Structure of the Universe” (www.universe-cluster.de) and the SFB-Tansregio TR33 The Dark Universe. We are especially grateful for the support by the Computational Center for Particle and Astrophysics (C²PAP). Computations have been performed at the Leibniz Supercomputing Centre (LRZ) with CPU time assigned to the Project

²⁴<http://www.magneticum.org>

²⁵<http://www.sternwarte.uni-erlangen.de/research/sixte>

²⁶<http://constellation.gsfc.nasa.gov/science/simulatorInstall.html>

“pr83li” and “pr86re”. The project received technical support from LRZ as well. Information on the *Magneticum Pathfinder* project is available at <http://www.magneticum.org>.

References

- Arnaud, K.A., 1996. XSPEC: The First Ten Years, in: Jacoby, G.H., Barnes, J. (Eds.), Astronomical Data Analysis Software and Systems V, p. 17.
- Bernyk, M., Croton, D.J., Tonini, C., Hodkinson, L., Hassan, A.H., Garel, T., Duffy, A.R., Mutch, S.J., Poole, G.B., Hegarty, S., 2016. The Theoretical Astrophysical Observatory: Cloud-based Mock Galaxy Catalogs. *ApJS* 223, 9. doi:10.3847/0067-0049/223/1/9, [arXiv:1403.5270](https://arxiv.org/abs/1403.5270).
- Biffi, V., Dolag, K., Böhringer, H., 2013. Investigating the velocity structure and X-ray observable properties of simulated galaxy clusters with PHOX. *MNRAS* 428, 1395–1409. doi:10.1093/mnras/sts120, [arXiv:1210.4158](https://arxiv.org/abs/1210.4158).
- Biffi, V., Dolag, K., Böhringer, H., Lemson, G., 2012. Observing simulated galaxy clusters with PHOX: a novel X-ray photon simulator. *MNRAS* 420, 3545–3556. doi:10.1111/j.1365-2966.2011.20278.x, [arXiv:1112.0314](https://arxiv.org/abs/1112.0314).
- Bocquet, S., Saro, A., Dolag, K., Mohr, J.J., 2016. Halo mass function: baryon impact, fitting formulae, and implications for cluster cosmology. *MNRAS* 456, 2361–2373. doi:10.1093/mnras/stv2657, [arXiv:1502.07357](https://arxiv.org/abs/1502.07357).
- Buchner, J., Georgakakis, A., Nandra, K., Hsu, L., Rangel, C., Brightman, M., Merloni, A., Salvato, M., Donley, J., Kocevski, D., 2014. X-ray spectral modelling of the AGN obscuring region in the CDFS: Bayesian model selection and catalogue. *A&A* 564, A125. doi:10.1051/0004-6361/201322971, [arXiv:1402.0004](https://arxiv.org/abs/1402.0004).
- Cui, W., Springel, V., Yang, X., De Lucia, G., Borgani, S., 2011. Properties of fossil groups in cosmological simulations and galaxy formation models. *MNRAS* 416, 2997–3008. doi:10.1111/j.1365-2966.2011.19248.x, [arXiv:1102.4269](https://arxiv.org/abs/1102.4269).
- Davis, J.E., Bautz, M.W., Dewey, D., Heilmann, R.K., Houck, J.C., Huenemoerder, D.P., Marshall, H.L., Nowak, M.A., Schattenburg, M.L., Schulz, N.S., Smith, R.K., 2012. Raytracing with MARX: x-ray observatory design, calibration, and support, in: Space Telescopes and Instrumentation 2012: Ultraviolet to Gamma Ray, p. 84431A. doi:10.1117/12.926937.
- Dolag, K., Borgani, S., Murante, G., Springel, V., 2009. Substructures in hydrodynamical cluster simulations. *MNRAS* 399, 497–514. doi:10.1111/j.1365-2966.2009.15034.x, [arXiv:0808.3401](https://arxiv.org/abs/0808.3401).
- Dolag, K., Gaensler, B.M., Beck, A.M., Beck, M.C., 2015a. Constraints on the distribution and energetics of fast radio bursts using cosmological hydrodynamic simulations. *MNRAS* 451, 4277–4289. doi:10.1093/mnras/stv1190, [arXiv:1412.4829](https://arxiv.org/abs/1412.4829).
- Dolag, K., Hansen, F.K., Roncarelli, M., Moscardini, L., 2005. The imprints of local superclusters on the Sunyaev-Zel'dovich signals and their detectability with Planck. *MNRAS* 363, 29–39.
- Dolag, K., Komatsu, E., Sunyaev, R., 2015b. SZ effects in the Magneticum Pathfinder Simulation: Comparison with the Planck, SPT, and ACT results. ArXiv e-prints [arXiv:1509.05134](https://arxiv.org/abs/1509.05134).
- Dolag, K., Reinecke, M., Gheller, C., Imboden, S., 2008. Splotch: Visualizing Cosmological Simulations. *NJoP*, submitted [arXiv:arXiv:astro-ph/xxxxx](https://arxiv.org/abs/arXiv:astro-ph/xxxxx).
- Hirschmann, M., Dolag, K., Saro, A., Bachmann, L., Borgani, S., Burkert, A., 2014. Cosmological simulations of black hole growth: AGN luminosities and downsizing. *MNRAS* 442, 2304–2324. doi:10.1093/mnras/stu1023, [arXiv:1308.0333](https://arxiv.org/abs/1308.0333).
- Ishisaki, Y., Maeda, Y., Fujimoto, R., Ozaki, M., Ebisawa, K., Takahashi, T., Ueda, Y., Ogasaka, Y., Ptak, A., Mukai, K., Hamaguchi, K., Hirayama, M., Kotani, T., Kubo, H., Shibata, R., Ebara, M., Furuzawa, A., Iizuka, R., Inoue, H., Mori, H., Okada, S., Yokoyama, Y., Matsumoto, H., Nakajima, H., Yamaguchi, H., Anabuki, N., Tawa, N., Nagai, M., Katsuda, S., Hayashida, K., Bamba, A., Miller, E.D., Sato, K., Yamasaki, N.Y., 2007. Monte Carlo Simulator and Ancillary Response Generator of Suzaku XRT/XIS System for Spatially Extended Source Analysis. *PASJ* 59, 113–132. doi:10.1093/pasj/59.sp1.S113, [arXiv:astro-ph/0610118](https://arxiv.org/abs/astro-ph/0610118).
- Jones, L.R., Ponman, T.J., Horton, A., Babul, A., Ebeling, H., Burke, D.J., 2003. The nature and space density of fossil groups of galaxies. *MNRAS* 343, 627–638. doi:10.1046/j.1365-8711.2003.06702.x, [arXiv:astro-ph/0304257](https://arxiv.org/abs/arXiv:astro-ph/0304257).
- Lemson, G., Virgo Consortium, t., 2006. Halo and Galaxy Formation Histories from the Millennium Simulation: Public release of a VO-oriented and SQL-queryable database for studying the evolution of galaxies in the LambdaCDM cosmogony. ArXiv Astrophysics e-prints [arXiv:astro-ph/0608019](https://arxiv.org/abs/arXiv:astro-ph/0608019).
- Marconi, A., Risaliti, G., Gilli, R., Hunt, L.K., Maiolino, R., Salvati, M., 2004. Local supermassive black holes, relics of active galactic nuclei and the X-ray background. *MNRAS* 351, 169–185. doi:10.1111/j.1365-2966.2004.07765.x, [arXiv:astro-ph/0311619](https://arxiv.org/abs/arXiv:astro-ph/0311619).
- McDonald, M., Benson, B.A., Vikhlinin, A., Aird, K.A., Allen, S.W., Bautz, M., Bayliss, M., Bleem, L.E., Bocquet, S., Brodwin, M., Carlstrom, J.E., Chang, C.L., Cho, H.M., Clocchiatti, A., Crawford, T.M., Crites, A.T., de Haan, T., Dobbs, M.A., Foley, R.J., Forman, W.R., George, E.M., Gladders, M.D., Gonzalez, A.H., Halverson, N.W., Hlavacek-Larrondo, J., Holder, G.P., Holzapfel, W.L., Hrubes, J.D., Jones, C., Keisler, R., Knox, L., Lee, A.T., Leitch, E.M., Liu, J., Lueker, M., Luong-Van, D., Mantz, A., Marrone, D.P., McMahon, J.J., Meyer, S.S., Miller, E.D., Mocanu, L., Mohr, J.J., Murray, S.S., Padin, S., Pryke, C., Reichardt, C.L., Rest, A., Ruhl, J.E., Saliwanchik, B.R., Saro, A., Sayre, J.T., Schaffer, K.K., Shirokoff, E., Spieler, H.G., Stalder, B., Stanford, S.A., Staniszewski, Z., Stark, A.A., Story, K.T., Stubbs, C.W., Vanderlinde, K., Vieira, J.D., Williamson, R., Zahn, O., Zenteno, A., 2014. The Redshift Evolution of the Mean Temperature, Pressure, and Entropy Profiles in 80 SPT-Selected Galaxy Clusters. *ApJ* 794, 67. doi:10.1088/0004-637X/794/1/67, [arXiv:1404.6250](https://arxiv.org/abs/1404.6250).
- Planck Collaboration, Ade, P.A.R., Aghanim, N., Arnaud, M., Ashdown, M., Atrio-Barandela, F., Aumont, J., Baccigalupi, C., Balbi, A., Banday, A.J., et al., 2013. Planck intermediate results. V. Pressure profiles of galaxy clusters from the Sunyaev-Zeldovich effect. *A&A* 550, A131. doi:10.1051/0004-6361/201220040, [arXiv:1207.4061](https://arxiv.org/abs/1207.4061).
- Rasia, E., Meneghetti, M., Ettori, S., 2013. X-ray Morphological Estimators for Galaxy Clusters. *The Astronomical Review* 8, 40–70. doi:10.1080/21672857.2013.11519713, [arXiv:1211.7040](https://arxiv.org/abs/1211.7040).
- Remus, R.S., Dolag, K., Naab, T., Burkert, A., Hirschmann, M., Hoffmann, T.L., Johansson, P.H., 2016. The Co-Evolution of Total Density Profiles and Central Dark Matter Fractions in Simulated Early-Type Galaxies. ArXiv e-prints [arXiv:1603.01619](https://arxiv.org/abs/arXiv:1603.01619).
- Saro, A., Liu, J., Mohr, J.J., Aird, K.A., Ashby, M.L.N., Bayliss, M., Benson, B.A., Bleem, L.E., Bocquet, S., Brodwin, M., Carlstrom, J.E., Chang, C.L., Chiu, I., Cho, H.M., Clocchiatti, A., Crawford, T.M., Crites, A.T., de Haan, T., Desai, S., Dietrich, J.P., Dobbs, M.A., Dolag, K., Dudley, J.P., Foley, R.J., Gangkofner, D., George, E.M., Gladders, M.D., Gonzalez, A.H., Halverson, N.W., Hennig, C., Hlavacek-Larrondo, J., Holzapfel, W.L., Hrubes, J.D., Jones, C., Keisler, R., Lee, A.T., Leitch, E.M., Lueker, M., Luong-Van, D., Mantz, A., Marrone, D.P., McDonald, M., McMahon, J.J., Mehl, J., Meyer, S.S., Mocanu, L., Montroy, T.E., Murray, S.S., Nurgaliev, D., Padin, S., Patej, A., Pryke, C., Reichardt, C.L., Rest, A., Ruel, J., Ruhl, J.E., Saliwanchik, B.R., Sayre, J.T., Schaffer, K.K., Shirokoff, E., Spieler, H.G., Stalder, B., Staniszewski, Z., Stark, A.A., Story, K., van Engelen, A., Vanderlinde, K., Vieira, J.D., Vikhlinin, A., Williamson, R., Zahn, O., Zenteno, A., 2014. Constraints on the CMB temperature evolution using multiband measurements of the Sunyaev-Zel'dovich effect with the South Pole Telescope. *MNRAS* 440, 2610–2615. doi:10.1093/mnras/stu575, [arXiv:1312.2462](https://arxiv.org/abs/1312.2462).
- Smith, R.K., Brickhouse, N.S., Liedahl, D.A., Raymond, J.C., 2001. Collisional Plasma Models with APEC/APED: Emission-Line Diagnoses of Hydrogen-like and Helium-like Ions. *ApJ* 556, L91–L95. doi:10.1086/322992, [arXiv:astro-ph/0106478](https://arxiv.org/abs/arXiv:astro-ph/0106478).

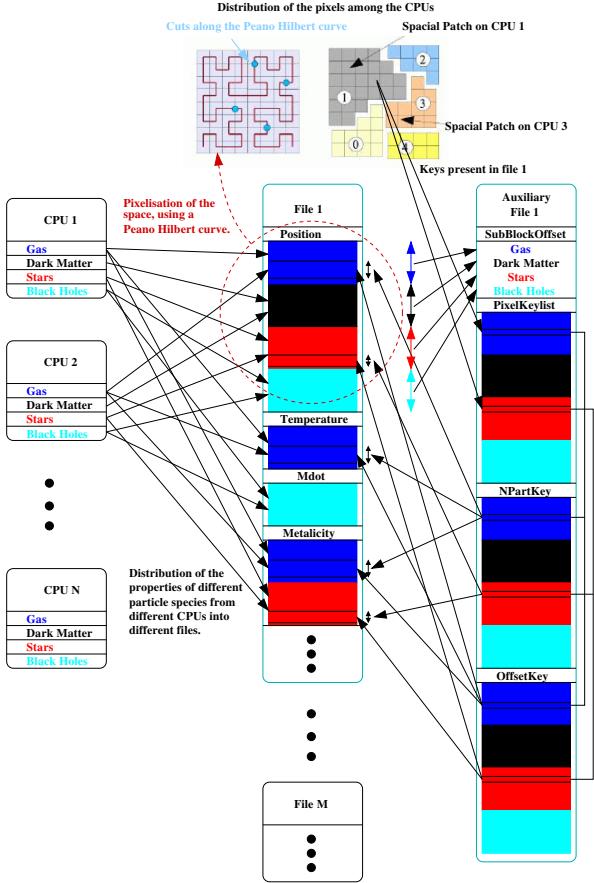


Figure A.11: This diagram illustrate how data of the different properties, attached to the different particle types are collected from several CPUs and stored into different files (left part). Once the position of the particles are associated (and sorted) according to their associated elements along a space filling curve, the index of such *pixels*, together with the number of particles associated with each *pixel* and the according offset in the files can be stored as auxiliary data.

- Springel, V., White, S.D.M., Tormen, G., Kauffmann, G., 2001. Populating a cluster of galaxies - I. Results at [formmu2]z=0. MNRAS 328, 726–750. doi:10.1046/j.1365-8711.2001.04912.x, arXiv:arXiv:astro-ph/0012055.
- Steinborn, L.K., Dolag, K., Comerford, J.M., Hirschmann, M., Remus, R.S., Teklu, A.F., 2016. Origin and properties of dual and offset active galactic nuclei in a cosmological simulation at z=2. MNRAS 458, 1013–1028. doi:10.1093/mnras/stw316, arXiv:1510.08465.
- Steinborn, L.K., Dolag, K., Hirschmann, M., Prieto, M.A., Remus, R.S., 2015. A refined sub-grid model for black hole accretion and AGN feedback in large cosmological simulations. MNRAS 448, 1504–1525. doi:10.1093/mnras/stv072, arXiv:1409.3221.
- Teklu, A.F., Remus, R.S., Dolag, K., Beck, A.M., Burkert, A., Schmidt, A.S., Schulze, F., Steinborn, L.K., 2015. Connecting Angular Momentum and Galactic Dynamics: The Complex Interplay between Spin, Mass, and Morphology. ApJ 812, 29. doi:10.1088/0004-637X/812/1/29, arXiv:1503.03501.
- Zdziarski, A.A., Poutanen, J., Johnson, W.N., 2000. Observations of Seyfert Galaxies by OSSE and Parameters of Their X-Ray/Gamma-Ray Sources. ApJ 542, 703–709. doi:10.1086/317046, arXiv:astro-ph/0006151.

Appendix A. Indexing files

Below we explain how to build those auxiliary data files in order to add third-party simulations to the virtual observatory. Figure A.11 shows a sketch of how the indexing scheme works for the multi-file and multi-component output of our SPH simulations. In such cosmological SPH simulations, the data to be stored is quite complex. The particle data are distributed among many thousands of CPUs. Each CPU holds various particles types, representing different components of the simulated system (e.g. *dark matter*, *gas*, *stars* and *black hole* particles). For these particles, different properties have to be stored. Some (like *temperature*) are special to individual particle species (like *gas* particles in this case). Other properties (like *position*) are necessary for all components. To optimize the I/O and to avoid bottlenecks, only a subset of CPUs write in parallel. Data from the individual CPUs are stored in individual sub-files, as illustrated in the left part of figure A.11. To optimize the access to individual data, files are structured with additional data regarding the content and length of individual data blocks (e.g. Gadget files).

We implemented an algorithm which sorts the particles of all CPUs among a coarse-grained space filling curve before writing them. In addition it produces an auxiliary file which allows identification of the sub-data volume elements of any stored property among all particle species associated with each element of the space filling curve. This allows the user to effectively collect all data associated with a given volume in space introducing minimal overhead. In detail, the process works as follows:

- The particles will be ordered along a space filing curve (e.g. a Peano Hilbert curve) with a defined graininess (e.g. *pixel size*)²⁷ among all CPUs. This order will be used to decompose the spacial region of the simulation among the different CPUs so that all particles within an individual *pixel* are associated to a CPU. During this process, the amount of particles of different species falling within each *pixel* is stored as well.
- The particle data of the various CPUs are written into the sub-files in the same order as above. Therefore particles belonging to the same *pixel* are written in a consecutive order into the files.
- An additional auxiliary file stores the information needed to re-construct the position of each sub-data-block for particles located in the a specific *pixel*. It contains the offsets between different particles species, the list of *pixels* present in the file, the number of particles in the related *pixel* and the offset of each *pixel* within each particle species.

²⁷ Note that in contrast to the standard domain decomposition in Gadget which is using a 64 bit long key, we are using a shorter, 32 bit long key which strongly reduces the later reading overhead.

- Finally a super index file is created, which contains the compressed information of *pixel* indexes and the files where they are contained. This super index is later used to recover which sub-files have to be accessed for the reading process.

These additional auxiliary files are produced for each output of the simulations and allows the user to investigate very efficiently individual objects within such cosmological simulation through post-processing tasks. The post-processing software has been adapted to this new output and for reading sub-volumes of the simulation. The following strategy is applied for a given sub-volume of the simulation which should be analysed (e.g. given by a galaxy cluster position and its size):

- First, a list of *pixels* (e.g. the elements of the space filling curve which falls within the region of the space of interest) is prepared.
- Then the super index file is read and a list of the output files which hold the individual subsets of the *pixel* index list is produced.
- For each sub-file in that list, first the *pixel* list is read from the auxiliary file and then it is compared with the *pixels* needed for the current task. For this subset of pixel, the additional information is read from the auxiliary file.
- Now the post-processing can reconstruct the position of the sub-blocks which have to be read from the individual blocks within the data file. This sub-blocks can then be read directly without loading the full data block. As further optimization, consecutive sub-blocks within individual files, are joined to larger sub-blocks to prevent unnecessary fragmentation of the reading process.

Appendix B. Preparing Simulations for the virtual observatory

In this appendix we describe how to add simulations to the data centre. To this scope, several pre-computed data products, meta data and informations have to be provided, as we briefly describe in the following.

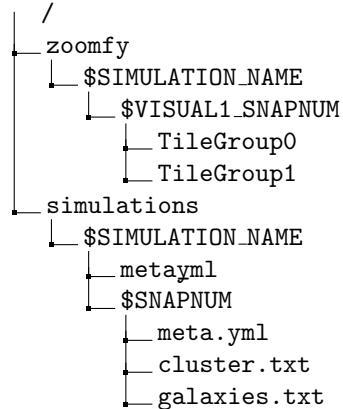
Appendix B.1. Describing the Simulations

Every simulation consists of one or more time instances (snapshots), where each should consist of a set of pre-computed 2D maps in the *zoomify*²⁸ format, a list of galaxy clusters and their galaxy members. Metadata regarding galaxy cluster properties and member properties are stored in *yaml*²⁹ files. Metadata and maps are stored in files and

Property	Description
name	Simulation name
code	Description of the Code used
size_mph	Mpc/h
n_particles	number of initial particles
nfiles	number of files the snapshot
mdm_msun	mass dark matter particles [Msun/h]
mgas_msun	mass gas particle [Msun/h]
epsilon_dm	softening of dark matter particles
epsilon_gas	softening of gas particles
epsilon_stars	softening of stars particles
f_bary	cosmological baryon fraction
omega_m	total matter content
omega_lambda	Cosmological constant
hubble	Hubble constant
sigma_8	Normalization of matter power-spectrum
n_primordial	Primordial power spectrum index
path_exec	file system path to simulation

Table B.2: Properties used to describe a simulation.

folder in a way that for every simulation there is a folder for its 2D maps (*zoomify* folder) and one for the metadata (*simulations* folder). Here follows a sample of the folder structure:



The *zoomify* files are stored on the local disk space of the web portal and contain the images which can be browsed by the web interface. Table B.2 shows a list of parameters used to describe a simulation.

Every time instance of a simulation, which is stored in the virtual observatory and made available via the web interface, contains its own sub directory which holds the prepared metadata for clusters and galaxies. They contain the information as described in the previous section. In addition, it also contains a general *yaml* metadata file with snapshot details. Table B.3 contains the parameters that typically describe a snapshot.

Such information and definitions can be extended easily to any other cosmological simulation and will allow the final user to flexibly add new simulations to the system. Also, services can be readily added or disabled for individual simulations.

²⁸<http://www.zoomify.com/free.htm>

²⁹see <http://yaml.org/>

Property	Description
name	snapshot name
mask_path	not used
width_pixel	total number of pixels
height_pixel	total number of pixels
redshift	Redshift
angular_diameter	Angular diameter distance
phox_max	Maximum GRASP (for Phox)
phox_avail	Phox service available
smac_avail	Smac service available
simcut_avail	SimCut service available
simcut_plot	SimPlot service available

Table B.3: Properties used to describe a snapshot within a simulation.

quantity	unit	name
ID	[integer]	id
x	[kpc/h]	x
y	[kpc/h]	y
z	[kpc/h]	z
M_{500c}	[M_\odot/h]	M500_msolh
R_{500c}	[kpc/h]	r500_kpch
f_{gas}	[fraction]	gas_frac
f_{stars}	[fraction]	star_frac
L_x	[10^{44}erg/s]	Lx_ergs
Y_{500c}	[$\Delta T/T$]	Y500c
M_{sat}/M_{cD}	[fraction]	M_sat_M_cD
c_{center}	[R_{500c}]	c_shift

Table B.4: Metadata for galaxy clusters and groups, from top to bottom: cluster ID, position (in comoving coordinates), mass and radius in respect to 500 times the critical density, gas and star fraction, bolometric X-ray luminosity, stellar mass fraction of satellite galaxies to central galaxy, and weighted centre shift between X-ray emission and mass distribution within R_{500c} .

quantity	unit	name
ID	[integer]	id
x	[kpc/h]	x
y	[kpc/h]	y
z	[kpc/h]	z
M_{star}	[M_\odot/h]	M_solh
M_{gas}	[M_\odot/h]	M_gas
sfr	[M_\odot/year]	sfr_msoly
host ID	[integer]	cluster_id
radial distance	[kpc/h]	dist
peculiar v_x	[km/s]	vx
peculiar v_y	[km/s]	vy
peculiar v_z	[km/s]	vz
velocity	[km/s]	dv
radial velocity	[km/s]	dr
tangential velocity	[km/s]	dt
mass ratio to cD	[km/s]	log10_mcD_m

Table B.5: Metadata for galaxies, from top to bottom: galaxy ID, position (in comoving coordinates), mass, star-formation rate, ID of cluster or group where the galaxy belongs to, distance to the center of the cluster or group it belongs to, different velocity components relative to the cluster center (as peculiar velocities) and the ratio of the stellar mass in satellite galaxies in respect to the central one.

Appendix B.2. The Metadata

Each galaxy cluster object has its metadata content, that is a list of properties associated with it (e.g. mass within a certain radius). The metadata associated with each galaxy cluster contain table which describes its galaxy members. We extract a reduced subset of halos (e.g. galaxies, groups or clusters, depending on the size and resolution of the underlying simulation) by applying a lower mass cut and providing the relevant part of the available global properties as metadata for user queries as shown in table B.4. In addition, for each halo (e.g. cluster, group or galaxy) we also store a list of all member galaxies (or satellites). For each of these galaxies/satellites we store some additional metadata as shown in table B.5, which can be used to further refine user queries.

Appendix C. Instrument Configurations

In this Appendix we report example configuration files for XISSIM, SIXTE, SIM and MARX, used for each instrument currently included within the PHOX service. Note that the exposure time will be replaced by the value chosen in the web interface.

Suzaku(XIS):

```
xissim \
  clobber=yes \
  instrume="XIS1" \
  ea1=0 ea2=0 ea3=90 \
  infile1="phlist_xissim.fits" \
  infile2=none \
  date_obs ="2009-09-01T00:00:00" \
  xis_rmf_file=suzaku/xis/cpf/ae_BI_ao4_20090901.rmf" \
  xis_contam_ifile="suzaku/xis/bcf/ae_xi1_contami_20061016.fits" \
  outfile="suzaku_xis_events.fits"
```

XMM(EPIC):

```
runsixt \
  EventList="sixtxmm_events.fits" \
  PatternList="sixtxmm_pattern.fits" \
  Mission="XMM" Instrument="EPICPN" \
  Mode="FFTHIN" \
  XMLFile="xmm/epicpn/fullframe_thinfilter.xml" \
  Simput="phlist.fits" \
  Exposure=1.0e4E \
  RA=10.0 Dec=10.0 \
  MJDREF=50814.0
epicpn_events \
  PatternList="sixtxmm_pattern.fits" \
  EPICPnEventList="xmm_epic_events.fits"
```

eROSITA:

```
erosim \
  prefix="erosita_" \
  PhotonList=events_pv.fits \
  RawData=events_allpv.fits \
  background=yes \
  XMLFile="srg/erosita_1.xml" \
  XMLFILE1="srg/erosita_1.xml" \
  XMLFILE2="srg/erosita_2.xml" \
  XMLFILE3="srg/erosita_3.xml" \
  XMLFILE4="srg/erosita_4.xml" \
  XMLFILE5="srg/erosita_5.xml" \
```

```

XMLFILE6="srg/erosita_6.xml" \
XMLFILE7="srg/erosita_7.xml" \
Simput="phlist.fits" \
Exposure=1.0e4 \
SkipInvalids=yes \
seed=-1 \
clobber=yes \
RA=10.0 Dec=10.0 \
MJDREF=50814.0

```

Athena(XIFU):

```

xifupipeline \
prefix="athena_xifu_" \
PixImpactList=impact.fits \
XMLfile=athena/1469mm_xifu/xifu_baseline.xml \
AdvXml=athena/1469mm_xifu/ \
xifu_detector_hex_baseline.xml \
Background=yes \
RA=10.0 Dec=10.0 \
Simput="phlist.fits" \
Exposure=1.0e4 \
UseRMF=T \
clobber=yes

```

Athena(WFI):

```

athenawfisim \
prefix="athena_wfi_" \
XMLFile0="athena/1469mm_wfi_w_filter/ \
depfet_b_11_ff_chip0.xml" \
XMLFile1="athena/1469mm_wfi_w_filter/ \
depfet_b_11_ff_chip1.xml" \
XMLFile2="athena/1469mm_wfi_w_filter/ \
depfet_b_11_ff_chip2.xml" \
XMLFile3="athena/1469mm_wfi_w_filter/ \
depfet_b_11_ff_chip3.xml" \
Simput="phlist.fits" \
Exposure=1.0e4 \
Background=yes \
Attitude="athena/attitude_wfi_ra10_dec10.fits" \
RA=10.00 Dec=10.00 \
chatter=0 \
MJDREF=52000.0 \
clobber=yes

```

XraySurveyor(HDXI)

```

pset simx mode=hl
pset simx Exposure=1.0e4
pset simx UseSimput=yes
pset simx MissionName=XraySurveyor
pset simx InstrumentName=HDXI
pset simx ScaleBkgnd=0.0
pset simx RandomSeed=24
pset simx SimputFile=phlist.fits
pset simx PointingRA=10.0
pset simx PointingDec=10.0
pset simx OutputFileName=surveyor_hdxi_events
simx

```

XraySurveyor(XCAL)

```

pset simx mode=hl
pset simx Exposure=1.0e4
pset simx UseSimput=yes
pset simx MissionName=XraySurveyor
pset simx InstrumentName=XCAL
pset simx ScaleBkgnd=0.0
pset simx RandomSeed=24
pset simx SimputFile=phlist.fits
pset simx PointingRA=10.0
pset simx PointingDec=10.0
pset simx OutputFileName=surveyor_xcal_events
simx

```

Hitomi(SXS)

```

pset simx mode=hl
pset simx Exposure=1.0e4
pset simx UseSimput=yes
pset simx MissionName=Hitomi
pset simx InstrumentName=SXS
pset simx ScaleBkgnd=0.0
pset simx RandomSeed=24
pset simx SimputFile=phlist.fits
pset simx PointingRA=10.0
pset simx PointingDec=10.0
pset simx OutputFileName=hitomi_sxs_events
simx

```

Hitomi(SXI)

```

pset simx mode=hl
pset simx Exposure=1.0e4
pset simx UseSimput=yes
pset simx MissionName=Hitomi
pset simx InstrumentName=SXI
pset simx ScaleBkgnd=0.0
pset simx RandomSeed=24
pset simx SimputFile=phlist.fits
pset simx PointingRA=10.0
pset simx PointingDec=10.0
pset simx OutputFileName=hitomi_sxi_events
simx

```

Hitomi(HXI)

```

pset simx mode=hl
pset simx Exposure=1.0e4
pset simx UseSimput=yes
pset simx MissionName=Hitomi
pset simx InstrumentName=HXI
pset simx ScaleBkgnd=0.0
pset simx RandomSeed=24
pset simx SimputFile=phlist.fits
pset simx PointingRA=10.0
pset simx PointingDec=10.0
pset simx OutputFileName=hitomi_hxi_events
simx

```

Chandra(ACIS-S)

```

marx S-SIMPUT-Source="phlist.fits" \
ExposureTime=1.0e4 TStart=2012.5 \
GratingType="NONE" DetectorType="ACIS-S" \
DitherModel="INTERNAL" RA_Nom=10 Dec_Nom=10 \
Roll_Nom=50 \
SourceRA=10 SourceDEC=10 \
Verbose=yes mode=h OutputDir=point \
marx2fits point chandra_acis-s_evt.fits

```

Chandra(ACIS-I)

```

marx S-SIMPUT-Source="phlist.fits" \
ExposureTime=1.0e4 TStart=2012.5 \
GratingType="NONE" DetectorType="ACIS-I" \
DitherModel="INTERNAL" RA_Nom=10 Dec_Nom=10 \
Roll_Nom=50 \
SourceRA=10 SourceDEC=10 \
Verbose=yes mode=h OutputDir=point \
marx2fits point chandra_acis-i_evt.fits

```