Leibniz-Rechenzentrum
Leibniz Supercomputing Centre

Template for Project Proposal for Tier 0/Tier1 HPC access at LRZ

Logo of your institution (if available):

Project name (as in our online proposal):
Magneticum

Project-ID at LRZ (if already known):

Principal Investigator / Project Manager (name, affiliation, address):
Klaus Dolag,
Ludwig-Maximilians-Universitaet Muenchen, Fakultaet fuer Physik, Universitaets Sternwarte
Scheinerstr. 1, 81679 Muenchen

Research field:
Cosmology

Confidentiality:

Is any part of the project covered by (enhanced) confidentiality? YES: ☐ NO: X

Collaborators: Joseph Mohr, Andreas Burkert, Alexandro Saro, Rhea-Silvia Remus, Ludwig-Maximilians-Universitaet Muenchen, Fakultaet fuer Physik, USM, Germany; Hans Boehringer, Max-Planck Institute for Extraterestische Physik (MPE), Germany; Margarita Petkova, C^2PAP eployee, Germany, Nicolay J. Hammer, LRZ, Germany; Stefano Borgani, Madhura Killedar, Astronomy Unit, University of Trieste, Italy; Michaela Hirschmann, INAF-Osservatorio Astronomico, Trieste, Italy; Giuseppe Murante, INAF – Osservatorio di Torino, Italy; Veronica Biffi, SISSA, Italy.

1 Abstract

We propose to perform a set of cosmological, hydrodynamical simulation, where the computational most intensive of them covers a cosmic volume of (1 GPc)^3, sampled with 2x3032^3 particles. This covers a large enough volume, with high enough spatial/mass resolution to follow in detail the evolution of galaxies and especially the AGN population. This will be accomplished by one low resolution but large volume (40x larger) simulation needed to reach a simulated volume comparable to the current and ongoing survey sizes. A final element of the simulation campaign will be a set of intermediate size simulations to explore the impact induced by the current uncertainty of the most important cosmological parameters. All simulations will follow cooling, star-formation, a detailed treatment of chemical enrichment by SN Ia, SN II and AGB winds and will include spitzer thermal conduction and magnetic fields. Additionally they will include the evolution of black holes and the corresponding AGN feedback. This will allow us for the very first time to study self consistently galaxy clusters, galaxy groups, galaxies and AGNs within a enormously large volume of the Universe. This will produce a theoretical counterpart to interpret the data coming from current and forthcoming astronomical surveys and instruments like PLANCK, SPT, DES and eROSITA.

2 Description of the Project

To fully exploit the potential of the upcoming large sky surveys for cosmology and for the study of the effects of Dark Energy, we need to improve the predictions of how the large-scale structure is traced by luminous matter. This requires a detailed description in simulation codes of various, complex, non-gravitational, physical processes, which determine the evolution of the cosmic baryons and impact their observational properties. Among them star formation and related feedback and chemical pollution by supernovae, the evolution of black holes and their related AGN feedback and magnetic fields, which all must be self consistently coupled with the underlying hydrodynamics. Being able to perform such large volume, high-resolution cosmological simulations which follow in detail the various physical processes is needed to produce a theoretical counterpart to interpret the data coming from current and forthcoming astronomical surveys and instruments like PLANCK, SPT, DES and eROSITA.

First, pioneering simulations, handling almost 10^{10} particles (e.g. *Magneticum Pathfinder*, Dolag et al., in preparation), have been performed which are used among various other projects for interpreting observations by PLANCK (Planck Collaboration et al., 2013, Dolag et al, in preparation) and SPT (see Figure 1 and Saro et al., in preparation) as well as for studying cluster properties in the X-ray band for Future, high energy resolution instruments like Athena/Astro-H (Biffi et al 2013) and to produce first mock observations for the forthcoming eROSITA satellite (used within the eROSITA cluster working group, see Figure 1). The previous set of simulations covered either very large volumes (order of Gpc^3) with a relatively low spatial/mass resolution, or high spatial/mass resolution, which then could only be reached in much smaller cosmological volume. The former allowed studying the ICM properties of galaxy clusters, whereas the later ones allowed study in detail properties of galaxies (Remus et al. 2013) or of the AGN component (Hirschmann et al. 2013). Here, for the first time, previous simulations were able to predict the multi wavelength properties of AGNs formed within the simulations in a very detailed way (see Figure 1 and Hirschmann et al. 2013).

In addition, thanks to the support obtained by successfully applying for the KONWHIR-III program (a program of the Bavarian Ministry of Science for high performance computing) we were able to further optimize P-Gadget3(XXL) for multi-core systems like SuperMUC. Thereby we are now able to efficiently perform simulations, which follow up to 10^{11} particles and run on hundreds of thousands of cores on HPC facilities like SuperMUC (see scaling section). This, for the first time, will allow us to perform a cosmological simulation, which covers large enough volume, e.g. (Gpc)^3 with high enough spatial/mass resolution to follow in detail the evolution of galaxies and especially the AGN population.

Therefore we propose to perform several simulations of cosmological boxes, whereas the computational most intensive covers (1 Gpc)^3, sampled with 2x3032^3 particles (corresponding to the high resolution of the previous performed simulations). Using our highly advanced SPH implementation we will, as before, follow cooling and star-formation, with a detailed treatment of chemical enrichment by SN Ia, SN II and AGB winds, magnetic fields and additional Spitzer thermal conduction. We also will include the evolution of black holes and the corresponding AGN feedback. This will allow us for the very first time to study self consistently galaxy clusters, galaxy groups, galaxies and AGNs within a enormously large volume of the Universe. Especially, we plan to study:

<u>a) Clusters as tracers of the large scale structures</u>

These simulations will allow us to predict X-ray and Sunyaev-Zel'dovich (SZ) signals from massive galaxy clusters and groups with new levels of fidelity and for a very large set of simulated clusters. Individual clusters will be used as templates to improve extraction of cluster properties from joint X-ray, SZ and optical observations. Deep light-cones can be used to predict the performance of cluster finders applied to mock observations and especially shed new light on contamination from the AGN component. Both will help to theoretically predict the significance with which cosmological parameters (especially dark energy parameters) can be extracted from current X-ray and SZ observations. Additionally the set of low resolution simulations with varying cosmological parameters will help to quantify the detailed influence of the cosmological parameters on the observational signatures of clusters and especially the SZ power-spectrum of deep light-cones.

<u>b) Cluster X-Ray and SZ scaling relations</u>

A major advantage of the large volume of these simulations is that X-Ray and SZ scaling relations can be extracted for a fair sample of massive clusters. Making use of the high resolution, it will be possible to self-consistently study the observed transition of the scaling relations from clusters to groups within this simulation and characterize especially the effect of AGN feedback. Moreover, studying the biases affecting the lensing measurements will allow us to verify the degree to which lensing can be used as a powerful complement to X-ray and SZ observations to calibrate cluster masses.

<u>c) Galaxy formation and evolution</u>

The high-resolution achieved for the simulations will allow us to investigate the formation and evolution of the different galaxy types and the history of star formation and chemical enrichment in great details. We will focus especially on the high-redshift (cosmological redshifts z=1-3) regime

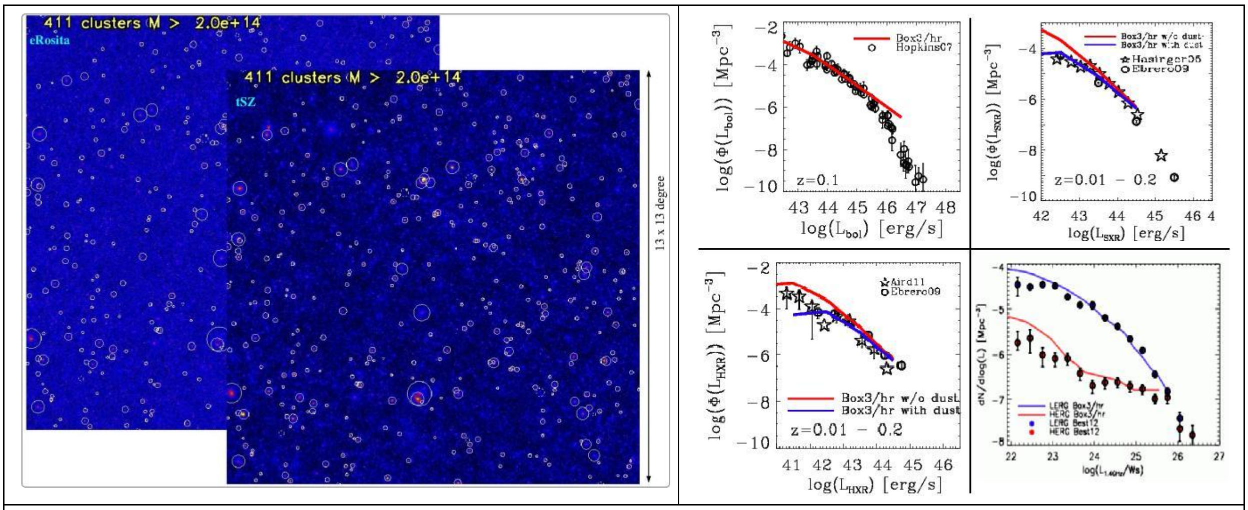

Figure 1: Left: Deep light-cone of the X-ray emission as would be seen by eROSITA as well as thermal Sunyaev Zel'dovich signal as template for SPT observations, obtained by a low resolution simulation of a large cosmological volume (1.3Gpc)^3. Right: Predicted AGN luminosities taken from the high resolution simulation of a smaller cosmological volume (182 Mpc)^3, compared to observations. Shown are luminosities obtained for different wavelength regime, from upper left to lower right the bolometric, soft X-ray, hard X-ray and radio luminosities are shown. Taken from Hirschmann et al., 2013.

where new, pioneering observations (e.g. with the ESO Very Large Telescopes) have revealed young galaxies with structures that are very different compared with present-day galaxies. The simulations will provide insight into how these galaxies formed, what the origin of their surprising substructures might be and how high-redshift galaxies evolve into the present-day galaxy population. They will allow us to predict how metal-rich galactic winds and central super-massive black holes shape galaxies and limit their star formation and how the evolution of galaxies is linked to their intergalactic surroundings.

d) BH growth and AGN evolution

The high-resolution achieved for the simulations together with the large volume covered will allow us to investigate the environmental dependencies as well as the AGN trigger mechanisms. It also allows us to study in detail the high-z formation of luminous quasars. Additional, we plan to study the spatial clustering of AGNs and the possibility to verify the degree to which AGNs can be used as a powerful probe of structure formation and cosmology, especially for the high redshift universe.

3 Numerical Methods and Algorithms

Gadget3 is a cosmological, highly optimized and fully MPI parallelized TreePM-MHD-SPH code. In its current version it also allows for an effective OpenMP parallelization within each MPI task. It makes use of the FFTW library to perform the PM part of the gravity solver. Here the gravity is split into a long range part, which is computed by sorting the particles onto a mesh (PM) and then solving the Poisson equation via FFT methods and a short range part, where a direct sum of the forces between particles is performed. For the short range part, forces at intermediate distances can be approximated by grouping particles and taking the moments of these groups (Tree method). The MHD solver is based on smoothed particle method (SPH) which allows to utilize various modern versions of the kernel functions and also makes use of a special treatment of the numerical viscosity. Additionally, a conjugant gradient solver allows to solve transport equations like thermal conduction. Stellar and sink particles are spawned as part of the treatment of star-formation and black hole evolution. Both make use of different sub-grid models to properly treat physical processes, which are far below the resolution limit.

Specifically, the proposed simulations will combine the following modules:

- low-viscosity SPH to allow the development of turbulence within the ICM (Dolag et al. 2005). Together with the usage of modern SPH kernels and the module for thermal conduction (see below) this represents a highly advanced SPH implementation;
- star formation and a detailed model of chemical enrichment (Tornatore et al. 2004/2007) which self-consistently includes the dependence of the gas cooling on the local metalicity (Wiersma et al. 2009), see below for more details;
- AGN feedback by sink particles (Springel et al. 2005), including radio mode feedback (Fabjan et al. 2010) as well as additional improvements (Hirschmann et al. 2013), see below for more details;
- magnetic fields (Dolag & Stasyszyn 2009, Stasyszyn et al 2013) using non-ideal MHD (Bonafede et al 2011);
- On the fly Sub-Find (Springel et al. 2001/2010, Dolag et al. 2009);
- On the fly Photometric code to assign optical/near-IR luminosities to galaxies (Saro et al. 2006, Nuza et al 2010);
- Novel sub-data access scheme allowing an efficient read-out of particles belonging to a galaxy cluster;
- thermal conduction (Dolag et al. 2004), but now based on a conjugant gradient solver;

Currently, we are testing following modules if they can be included:

- magnetic seeding by star-formation (Beck et al 2013);

- addaptive gravitational softening to extend the dynamical range of the simulation (Iannuzzi & Dolag 2011);
- Non isotropic thermal conduction (Arth et al., in preparation);

Treating star-formation:

We are using the multiphase model for star-formation (Springel & Hernquist 2003), where the ISM is treated as a two-phase medium where clouds of cold gas form from cooling of hot gas and are embedded in the hot gas phase assuming pressure equilibrium whenever gas particles are above a given threshold density. The hot gas within the multiphase model is heated by supernovae and can evaporate the cold clouds. A certain fraction of massive stars (10 per cent) is assumed to explode as supernovae type II (SNII). The released energy by SNII (10^{51} erg) is modelled to trigger galactic winds with a mass loading rate being proportional to the star formation rate (SFR) to obtain a resulting wind velocity of v_{wind} = 350 km/s. Our simulations also include a detailed model of chemical evolution according to Tornatore et al. (2007). Metals are produced by SNII, by supernovae type Ia (SNIa) and by intermediate and low-mass stars in the asymptotic giant branch (AGB). Metals and energy are released by stars of different mass to properly account for massdependent life-times (with a lifetime function according to Padovani & Matteucci 1993), the metallicity-dependent stellar yields by Woosley & Weaver (1995) for SNII, the yields by van den Hoek & Groenewegen (1997) for AGB stars and the yields by Thielemann et al. (2003) for SNIa. Stars of different mass are initially distributed according to a Chabrier initial mass function (IMF; Chabrier 2003). Radiative cooling rates are computed by following the same procedure presented by Wiersma et al. (2009). We account for the presence of the cosmic microwave background (CMB) and of ultraviolet (UV)/X-ray background radiation from quasars and galaxies, as computed by Haardt & Madau (2001). The contributions to cooling from each one of 11 elements (H, He, C, N, O, Ne, Mg, Si, S, Ca, Fe) have been pre-computed using the publicly available CLOUDY photoionisation code (Ferland et al. 1998) for an optically thin gas in (photo-)ionisation equilibrium.

Treating Black Hole growth and AGN feedback:

The model for AGN feedback is based on the original implementation presented by Springel et al. 2005, with feedback energy released as a result of gas accretion onto supermassive black holes. In this AGN model, BHs are described as sink particles, which grow their mass by gas accretion and merging with other BHs. Gas accretion proceeds at a Bondi rate, and is limited by the Eddington rate. Once the accretion rate is computed for each BH particle, a stochastic criterion is used to select the surrounding gas particles to be accreted. Unlike in Springel et al. 2005, in which a selected gas particle contributes to accretion with all its mass, we included the possibility for a gas particle to accrete only with a slice of its mass, which corresponds to 1/4 of its original mass. In this way, each gas particle can contribute with up to four "generations" of BH accretion events, thus providing a more continuous description of the accretion process. Eddington-limited Bondi accretion produces a radiated energy which corresponds to a fraction $\varepsilon_r = 0.1$ of the rest-mass energy of the accreted gas, which is determined by the radiation efficiency parameter ε_r. The BH mass is correspondingly decreased by this amount. A fraction of this radiated energy is thermally coupled to the surrounding gas. We use $\varepsilon_f = 0.1$ for this feedback efficiency, which increases by a factor of 4 when accretion enters in the quiescent "radio" mode and drops below one-hundredth of the limiting Eddington rate (e.g. Sijacki et al. 2007; Fabjan et al. 2010). More detailed description (including other technical changes) of the used BH model can be found in Hirschmann et al. 2013.

4 Computer Resources

Explain why this project needs to run on a Tier-0/Tier-1 system,

The most computational intensive simulation (*Box2b/hr*) will follow $2x3032^3$ particles and is required in order to cover a volume of $(1Gpc)^3$, large enough to produce mock observations for Planck, SPT and eROSITA, while simultaneously providing the resolution essential for a self-consistent description of ICM, galaxies and AGN properties. For this we will need 64TB of main memory (including already system overhead on SuperMUC), so we will need to use 32768 cores (e.g. 4 islands of SuperMUC).

Including recent improvement in the code performance, we expect this simulation to take 25 Million CPUh on SuperMUC, based on our various test simulations. As part of a weak scaling test performed on SuperMUC we already performed several time-steps for the according initial conditions.

The second most computational intensive simulation (*Box0/mr*) will be performed in close collaboration with the LRZ staff team. It will be a low resolution (e.g. same resolution of previous simulations) but very large volume simulation, utilizing $2x4172^3$ particles covering a volume of $(3.5Gpc)^3$. As this simulation will need 16 islands of SuperMUC, it therefore has to be performed in several steps, always starting after maintenance periods, in close collaboration with LRZ staff. Note that for this simulation we need much less on the fly post processing (factor 4 less), therefore the total CPU time needed is estimated to be 15 Million CPUh.

Additional, we plan to perform 13 versions of *Box1/mr* from the previous Magneticum Pathfinder simulation set. This simulations will utilize 2x1526^3 particles, covering $(1.3GPc)^3$ volumes and test the influence of the most important cosmological parameters σ_8, Ω_M and Ω_b. These simulations will – in an orthogonal study – explore 5 values for each of these cosmological parameters, varied around their current, bets fit values (Planck team 2013). This will, for the first time, allow us to study the influence of these parameters on observables like scaling relations, mass function and most importantly, the amplitude SZ power spectra obtained from deep light-cones, obtained from these hydro dynamical simulations. Having 5 values for each of these cosmological parameters will allow us to fit simple functional forms to the results, which then can be used for scaling of templates used to interpret/fit the results from different surveys. As each of these simulations will take 1 Million CPUh, this will sum up to 13 Million CPUh in total.

Scaling of Gadget3 on SuperMUC:

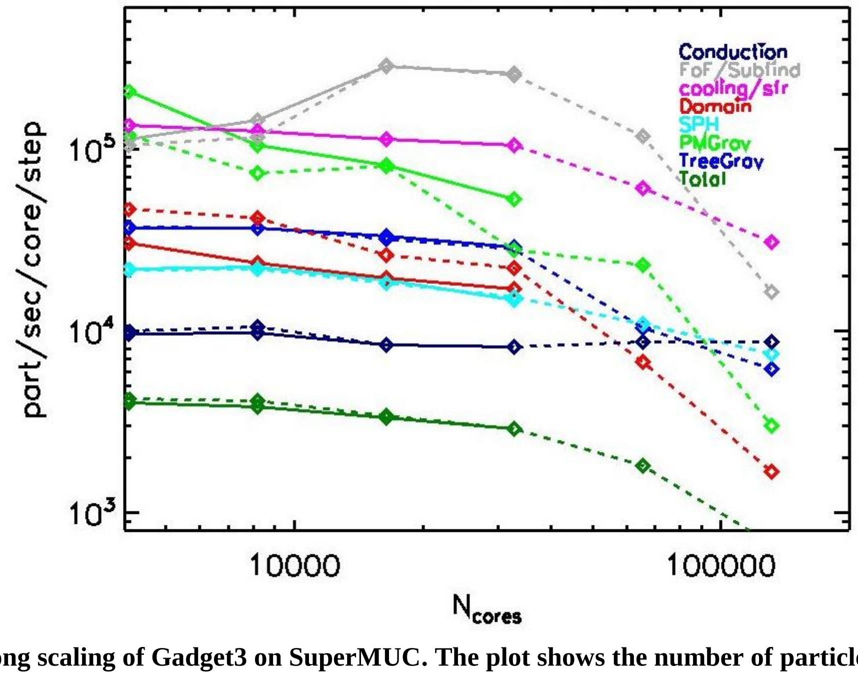

Figure 2 – Strong scaling of Gadget3 on SuperMUC. The plot shows the number of particles processed per core, time step, and wall clock time, measured over the first 10 time-steps, broken down for the most computationally intensive parts of the calculation and spans from 0.5 to 16 islands of SuperMUC. Solid and dashed lines indicate two different communication structures which can be used.

Justify the number of core hours requested.

Long range forces (fftw based PM part) and the conjugate gradient solver are only applied at global time-steps, whereas gravity, hydrodynamics and sub-grid models are based on individual time-steps of the particles. Till now, we have performed several tens of end to end simulations of various cosmological boxes as well as many short term performance tests on SuperMUC, including test performed during the "Extreme scaling workshop" hosted at LRZ, where we could test the scaling of the code up to 16 islands (e.g. 131072 cores) of SuperMUC (see previous section and Figure 2). Based on these various test runs we inferred the expected need of CPU time for the different simulations we are planning (see table).

Type of run	NumPart	Boxsize	M_{DM} [$10^{10}M_o$]	ε_{Soft} [kpc]	Timesteps	Cores used	Total core hours
Box2b/hr	$2x3128^3$	$(1Gpc)^3$	0.099	7	~750000	32768	25 Mio
Box0/mr	$2x4172^3$	$(3.5GPc)^3$	1.9	14	~250000	131072	15 Mio
13 x Box1/mr	$2x1526^3$	$(1.3GPc)^3$	1.9	14	~180000	4096	13 x 1Mio
TOTAL							**53** Mio

Describe how you intend to manage the resources you have requested

The first essential element of our work-flow is the newly founded "Computational Centre for Particle and Astrophysics" (C^2PAP, coordinator J. Mohr) of the excellence cluster "Origin and Structure of the Universe" in Munich. Dedicated computer hardware (computing cluster and storage systems) has become operational in August 2013 and is located at LRZ. This will allow us to store significant parts of the simulation results. Also the computing cluster will be used for parts of the post processing and staff members of the C^2PAP are already involved in improving and optimization of the code. This gives us also support on all technical aspects for preparing and running the simulation by the according personnel (M. Petkova). The second element is the just finished KONWHIR-III project, which was successful in preparing and optimizing of the TreePM-MHD code Gadget3 for SuperMUC (see scaling section). The third element is the comprehensive, international collaboration which will contribute to the scientific mining of the simulations.

Task	Timeframe	Personal	Level
Code verification and remaining optimization	Ongoing	Petkova, Dolag, Murante	0
Initial condition production	ongoing (new PLANCK results)	Dolag	0
First Production runs	month 1-6 (Box2b/hr) month 6-9 (Box0/mr) month 9-12 (13 x Box1/mr)	Dolag, Petkova, Hammer	1
Photon cubes (X-Ray)	parallel to simulation	Biffi, Dolag	1
Galaxy and AGN catalogues	Parallel to simulation	Hirschmann, Remus, Saro, Dolag	1
SZ light-cone	parallel to simulation	Saro, Dolag	1
Lensing light-cone	after simulation finished	Killedar, Dolag	1
High redshift cluster and proto cluster	after month 3	Borgani, Mohr, Boehringer, Saro, Dolag	2
Cluster scaling relations	after simulation finished	Boehringer, Borgani, Dolag, Biffi, Mohr	2
Mock observations at various wavelengths	After full light-cone construction	Mohr, Saro, Biffi, Killedar, Dolag, Murante	2
BCG, galaxy and AGN properties	After month 3	Burkert, Hirschmann, Remus, Dolag, Saro	2

The work proposed here can be characterized in 3 levels. Preoperational work (level 0) is ongoing utilizing preparatory access to SuperMUC. Level 1 regards the technical aspect of running the big simulations and the main post processing of the primary output. With the simulation ongoing, most of the additional post processing can be done snapshot by snapshot (as soon as individual snapshots get available), while the simulations are progressing. The different tasks will be distributed as indicated in the table above. As an exception, the gravitational lensing signal can only be computed when the individual simulations are almost completely finished. Level 2 marks the activities directly related to the individual science goals, which for simplicity are grouped together in 4 classes. Some of the science goals can be already investigated, as soon as the simulation reaches redshifts of z~1-3, especially the projects focusing on proto clusters and groups. Other science goals can be pursued as soon as light cones are available, while others will need further processing on the light-cones to mimic instrumental effects.

Discuss the routes that you intend to use for dissemination
Besides the scientific outcome from the primary science goals as listed above we expect the obtained simulation results to be base for various future Bachelor and Master thesis projects (as in the past, see Bachelor/Master projects on http://www.usm.uni-muenchen.de/CAST/student_projects.htm). Also, as for the simulations obtained in past calls for CPU time at HPC centers, we expect that the products of the simulation (catalogues, maps, profiles, light-cones) will be used in various scientific collaborations even beyond the current layout.

5 Assistance needed from LRZ / Collaboration with LRZ

Especially for the second simulation we request consultancy by LRZ staff to enable us to use the available "special" que to be able to run the simulation after regular scheduled maintenance periods on 16 islands of SuperMUC. Here we also request guidance on the scheduling of the jobs (e.g. the frequency of writing restart files, length of individual job parts to be scheduled) to optimize the run time against losses due to unavoidable aborts due to individual node failures.

6 Key Personnel and Experiences

- Klaus Dolag (USM, Staff)
  Research interests: Magnetic Fields, Cosmology, Clusters of Galaxies
  Previous proposals: h0073 (HLRB2), pr86re (SuperMUC), DEISA-6, GAUSS

- Margarita Petkova (Computational Center for Particle and Astrophysics)
  Field of expertise: High performance computing, code optimization and parallelization, numerical/computational cosmology

- Nicolay J. Hammer (Applications-Support, LRZ)
  Field of expertise: High performance computing, hydrodynamics, SuperMUC

7 References

Biffi et al. 2013: Investigating the velocity structure and X-ray observable properties of simulated galaxy clusters with PHOX, MNRAS 428

Dolag & Stasyszyn 2009: "An MHD Gadget for cosmological simulations", MNRAS 398

Dolag et al. 2004: "The effects of Thermal Conduction in Simulated Galaxy Clusters", ApJ 606L

Dolag et al. 2005: "Turbulent gas motion in galaxy cluster simulations: The role of SPH viscosity", MNRAS 364

Dolag et al. 2009: "Substructures in hydrodynamical cluster simulations", MNRAS 339

Fabjan et al. 2010: "Simulating the effect of AGN feedback on the metal enrichment of galaxy clusters", MNRAS, 401

Hirschmann et al 2013: "Cosmological simulations of black hole growth: AGN luminosities and downsizing", astro-ph 1308:0333

Planck Collaboration 2013: Planck intermediate results. V. Pressure profiles of galaxy clusters from the Sunyaev-Zeldovich effect, A&A 550

Remus et al. 2013: "The Dark Halo—Spheroid Conspiracy and the Origin of Elliptical Galaxies", ApJ 776

Saro et al. 2006: "Properties of the galaxy population in hydrodynamical simulations of clusters", MNRAS 369

8 Reviewers

Steffan Gottlöber (Leibniz-Institut für Astrophysik Potsdam, Germany, sgottloeber@aip.de)
Gustavo Yepes (Universidad Autónoma de Madrid, Spain, gustavo.yepes@uam.es)
August Evrard (Department of Physics, University of Michigan, USA, evrard@umich.edu)

9 Acknowledgement

After your proposal has been successfully evaluated, please always include an acknowledgement in your publications such as:

Computer resources for this project have been provided by the Gauss Centre for Supercomputing/Leibniz Supercomputing Centre under grant: <project-ID>.

10 Upload

Upload this Project Proposal to: https://www.lrz.de/services/compute/supermuc/upload-report/
Select „Detailed Project Description“ as type of the document.