



CONTRACT NUMBER 508830

**DEISA**  
**DISTRIBUTED EUROPEAN INFRASTRUCTURE FOR  
SUPERCOMPUTING APPLICATIONS**

**European Community Sixth Framework Programme**  
**RESEARCH INFRASTRUCTURES**  
**Integrated Infrastructure Initiative**

Grid-enabled tools for evaluating  
basic properties of the dark matter

**Deliverable ID: D-JRA2-2.4**  
**Due date: October, 31, 2005**  
**Actual delivery date: November 28, 2005**  
**Lead contractor for this deliverable: EPCC, UK**

**Project start date: May 1<sup>st</sup>, 2004**  
**Duration: 4 years**

<b>Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)</b>		
<b>Dissemination Level</b>		
<b>PU</b>	Public	X
<b>PP</b>	Restricted to other programme participants (including the Commission Services)	
<b>RE</b>	Restricted to a group specified by the consortium (including the Commission Services)	
<b>CO</b>	Confidential, only for members of the consortium (including the Commission Services)	

## Document Keywords and Abstract

<b>Keywords:</b>	Grid-enabling, pre- and post-processing cosmology tools, group finders, Friends-of-Friends, FoF, SubFind, MergerTrees, Power, Correl, Virgo, GADGET, GADGET2
<b>Abstract:</b>	<p>This report describes the processes involved in porting, Grid-enabling and, where necessary, parallelising a number of pre- and post-processing tools employed by the Virgo Consortium. These codes are Power and Correl. The process of Grid-enabling these tools involved the incorporation of a portable binary data format. The format employed is HDF5.</p> <p>Specifically, this report is D-JRA2-2.4, and describes the work undertaken on Power and Correl in terms of the Specifications Document, D-JRA2-2.1</p>

**Table of Content**

**1 INTRODUCTION .....2**

    1.1 Executive Summary ..... 2

    1.2 References and Applicable Documents..... 3

    1.3 List of Acronyms and Abbreviations ..... 3

**2 VIRGO’S COSMOLOGICAL SIMULATION CODES .....4**

    2.1 GADGET..... 4

**3 DESCRIPTION OF POWER AND CORREL .....4**

    3.1 Power..... 4

    3.2 Correl ..... 5

**4 CONCLUSIONS .....5**

# 1 Introduction

## 1.1 Executive Summary

This project is DEISA's JRA2 – Cosmology Applications, and is a joint effort between EPCC and the UK's Virgo Consortium [1]. Funding for JRA2 comes from both DEISA and VirtU [7], where VirtU is Virgo's e-Science Virtual Universe project, funded by the UK Particle Physics and Astronomy Research Council (PPARC). VirtU started on the 1<sup>st</sup> October 2004 and is to run for a 36 months period. It will form the foundations of the Theoretical Virtual Observatory.

All the tasks in WP2 were jointly funded by both DESIA and VirtU and, as such, all deliverables need to be reported to both DEISA and VirtU.

This report is the second deliverable of WP2, D-JRA2-2.2, 'Grid-enabled tools for evaluating basic properties of the dark matter'.

The input and output of one of Virgo's largest cosmological simulations codes, namely GADGET2 [8], are created and manipulated by a number of pre- and post-processing tools. GADGET2 has been ported to the DEISA infrastructure by JRA2's WP1 [1].

The main motivation of this work is that, once GADGET has been run on DEISA, albeit on a single platform [1], the post-processing tools can be run anywhere on the DEISA infrastructure, provided one has access to the data. This is possible due to the MC-GPFS, however, the data files must be stored in a portable data format.

The tools have been ported to the DEISA infrastructure, Grid-enabled and parallelised where necessary. Specifically, the tools ported were Friends-of-Friends (FoF), SubFind, MergerTrees (TraceSubgroups, SplitHalos and BuildTrees), Power and Correl. This report considers only Power and Correl. The Group Finders, namely FoF and SubFind are considered in [3], whilst the MergerTrees suite is considered in [4], The User Guide for all these tools is given in [5].

In this context, we take Grid-enabling to mean allowing the tools to utilise distributed data sets within the DEISA infrastructure, at least, and by introducing a portable, binary data format, where this format was initially chosen to be HDF5 [10], into the tool's I/O routines. For the purposes of this work package it was decided that what grid-enabling entails in practice on the DEISA infrastructure is that the associated codes need to be able to access binary data files wherever they may reside on the infrastructure. This implies that the binary data files must be readable irrespective of the platform on which they were generated. To this end, we employ HDF5. It was found that HDF5 performed the I/O faster than the other portable binary data formats tested, but also faster than the standard C `fwrite()` routine on a Sun and an IBM cluster [3].

Power, which was already parallelised, has been adapted to read dark matter only data files using either the default binary format or HDF5. Correl was extracted from L-GADGET2, thus the code was already parallelised. The I/O routines were written to read dark matter only data files using either the default binary format or HDF5. In addition, a large look-up table has been included for debugging purposes, which ensures that an array of random numbers is the same on all processors. Both Power and Correl have

been installed, run and checked for correctness on HPCx, IDRIS and RZG. This work was performed by the authors.

## 1.2 *References and Applicable Documents*

- [1] D-JRA2-1.2, Grid-enabled implementation of GADGET for metacomputing environments.
- [2] D-JRA2-2.1, Specification Document for D-JRA2-2.2, D-JRA2-2.3, D-JRA2-2.4.
- [3] D-JRA2-2.2, Grid-enabled “Group finders”.
- [4] D-JRA2-2.3, Grid-enabled tools to build halo merger trees.
- [5] D-JRA2-2.5, Documentation for D-JRA2-2.2, D-JRA2-2.3, D-JRA2-2.4.
- [6] Virgo: <http://www.virgo.dur.ac.uk>.
- [7] VirtU: <http://star-www.dur.ac.uk/~csf/virtU/virtU-final.pdf>.
- [8] GADGET: <http://www.mpa-garching.mpg.de/galform/gadget>.
- [9] DEISA: <http://www.deisa.org>.
- [10] HDF5: <http://hdf.ncsa.uiuc.edu/HDF5>.

## 1.3 *List of Acronyms and Abbreviations*

<b>GADGET</b>	<b>GA</b> laxies with <b>D</b> ark matter and <b>G</b> as intErac <b>T</b>
<b>MPI</b>	<b>M</b> essage <b>P</b> assing <b>I</b> nterface
<b>HDF5</b>	<b>H</b> ierarchical <b>D</b> ata <b>F</b> ormat <b>5</b>
<b>VirtU</b>	<b>V</b> irtual <b>U</b> niverse

## 2 Virgo's Cosmological Simulation Codes

The codes described in this document are all used as pre- and/or post-processing tools to the Virgo [6] Consortium's main Cosmological Simulation Codes, namely: GADGET2 [8].

### 2.1 GADGET

GADGET is a code to simulate the evolution of the universe, modelling the motion of both dark matter and gas and, essentially, is a very large N-body simulation code. For more information, see [8].

As part of JRA2, WP1, GADGET2 was ported to the DEISA infrastructure, where GADGET2 is the most up-to-date, general version of GADGET that currently exists.

## 3 Description of Power and Correl

### 3.1 Power

As described in the Specification Document, [1], Power is used to calculate the power spectrum of a dark matter distribution and is a parallel C/MPI code with around 1600 lines. However, also from [1], GADGET2 actually contains parallel code similar to Power. It was considered, at the time, to be wise to extract the code from GADGET2 itself, as GADGET2 includes an advanced form of Power.

However, extracting Power from GADGET was not straightforward, as the Power spectrum calculation included in GADGET2 depends on the domain decomposition GADGET2 is using.

Therefore, we took the existing parallel version of Power and adapted it for use with HDF5 GADGET2 snapshots.

Power's I/O statements have been replaced to function with GADGET's snapshot files, instead of the original galaxy coordinates, in the default native binary and HDF5.

Power's I/O also had to be altered to read dark matter distributions.

The new Power code is parallelised in the same manner as FoF and SubFind, [3], and the MergerTrees suite, [4]. The three dimensional volume is split into two dimensional slabs. A single processor reads in the snapshot files in blocks and these blocks are distributed around the remaining processors using MPI\_Ssend. This method allows any size of simulation to be analysed by not forcing the I/O processor to read the entire snapshot into memory.

This new Power code (L-Power) reads all of the particles of one species from a snapshot, calculates the density field on a mesh, then does a fast Fourier transform to get the power spectrum. It writes the power spectrum as an ASCII table, 'powerspec.txt'.

### 3.2 Correl

Correl will be used to calculate the correlation function of a dark matter distribution. The Correl code is quite sophisticated, in that the correlation function can be measured accurately, even for large scales, in a relatively fast time when compared to more traditional methods.

A serial version exists, written in C with about 1200 lines of code. GADGET2 also contains a parallel code, which is more advanced than Correl. It was stated in the Specifications Document, [1], that it would be *this* code that will be extracted, ported to the DEISA infrastructure and Grid-enabled [1] to form the DEISA version of Correl.

A new parallel Correl code was produced by extracting the necessary routines from L-GADGET2, which is a specialised version of GADGET2 intended for very large, dark matter only simulations. The I/O statements in this code were then modified to read the dark matter particles from GADGET snapshots (which may also contain gas and star particles) stored in the HDF5 format.

This new version of Correl (L-Correl) read the dark matter particles from a snapshot and found the two point correlation function by counting the number of pairs of particles as a function of separation. To do this, a random sample of particles was chosen – using all the particles in the simulation would have taken too long. A sphere was placed around each of these selected particles and the distances to all the other particles in the sphere were calculated. The radius of each of the spheres was chosen at random. This procedure gives an estimate of the number of pairs of particles as a function of separation and hence the correlation function, which is simply the excess probability of finding a neighbouring particle at a given separation.

Under testing, it was found that the output from the correlation function code had a dependency on the number of processors used. This happened because the random numbers used to determine which particles were in the sample changed when the number of processors was varied. As a result, there were small variations in the correlation function.

To solve this problem, we replicated a long look-up table of random numbers to every processor, using the particle ID as an index into this lookup table, which makes the result of the code independent of the number of processors. This lookup table is included, or not, using pre-processing tags.

## 4 Conclusions

Power, which was already parallelised, has been adapted to read dark matter only data files using either the default binary format or HDF5.

Correl was extracted from L-GADGET2, thus the code was already parallelised. The I/O routines were written to read dark matter only data files using either the default binary format or HDF5. In addition, a large look-up table has been included for debugging purposes, which ensures that an array of random numbers is the same on all processors.

Both Power and Correl has been installed, run and checked for correctness on HPCx and two of the core DEISA supercomputers, namely the platforms at IDRIS and RZG. This work was performed by the authors. This was a straightforward process, as the codes adhere to standard language features. The fact that the I/O routines now employ HDF5 means that the binary datasets did not have to be manipulated when porting the datafiles between platforms.

A User Guide to both Power and Correl can be found in [5].

The purpose of the DEISA research infrastructure is to enable scientific discovery across a broad spectrum of science and technology. The Virgo Consortium are typical of new users of the infrastructure requiring support and modest development of their software applications to take advantage of the new infrastructure. This work has been done under the auspices of the DEISA and VirtU projects and as a result a new community of scientific researchers have been brought to the infrastructure. The software applications are available for other users under the existing terms of the Virgo Consortium who are the application owners.

The objective of this work has been achieved, in that the output from GADGET cosmological simulations can now be post-processed from any of the DEISA sites, given that the Power and Correl codes have been ported to those sites.