Ganga: a tool for computational-task management and easy access to Grid resources

J.T. Mościcki^{f,*}, F. Brochu^a, J. Ebke^c U. Egede^b,
J. Elmsheuser^c, K. Harrison^a, R.W.L. Jones^d, H.C. Lee, ^{e,1},
D. Liko^f, A. Maier^f, A. Muraru^f, G.N. Patrick^g, K. Pajchel^j,
W. Reece^b, B.H. Samset^j, M.W. Slaterⁱ, A. Soroko^h,
C.L. Tanⁱ, D.C. Vanderster^f M. Williams^b

^a University of Cambridge, Cambridge, United Kingdom
 ^b Imperial College London, London, United Kingdom
 ^c Ludwig-Maximilians-Universität, Munich, Germany
 ^d Lancaster University, Lancaster, United Kingdom
 ^e NIKHEF, Amsterdam, The Netherlands
 ^f CERN, Geneva, Switzerland
 ^g STFC Rutherford Appleton Laboratory, Didcot, United Kingdom
 ^h University of Oxford, Oxford, United Kingdom
 ⁱ University of Birmingham, Birmingham, United Kingdom
 ^j University of Oslo, Oslo, Norway

Abstract

In this paper, we present the computational task-management tool GANGA, which allows for the specification, submission, bookkeeping and post-processing of computational tasks on a wide set of distributed resources. GANGA has been developed to solve a problem increasingly common in scientific projects, which is that researchers must regularly switch between different processing systems, each with its own command set, to complete their computational tasks. GANGA provides a homogeneous environment for processing data on heterogeneous resources. We give examples from High Energy Physics, demonstrating how an analysis can be developed on a local system and then transparently moved to a Grid system for processing of all available data. GANGA has an API that can be used via an interactive interface, in scripts, or through a GUI. Specific knowledge about types of tasks or computational resources is provided at run-time through a plugin system, making new developments easy to integrate. We give an overview of the GANGA architecture, give examples of current use, and demonstrate how GANGA can be used in many different areas of science. Key words: Grid computing, Data mining, Task management, User interface, Interoperability, System integration, Application configuration *PACS:* 07.05.Kf, 07.05.Wr, 29.50.+v, 29.85.+c, 87.18.Bb, 89.20.Ff

1 1 Introduction

Scientific communities are using a growing number of distributed systems, from local batch systems and community-specific services to generic, global 3 Grid infrastructures. Users may debug applications using a desktop computer. Δ then perform small-scale application testing using local resources and finally 5 run at full-scale using globally distributed Grids. Sometimes new resources 6 are made available to the users through systems previously unknown to them, 7 and significant effort may be required to gain familiarity with these systems 8 interfaces and idiosyncracies. The time cost of mastering application configç uration, tracking of computational tasks, archival and access to the results is 10 prohibitive for the end-users if they are not supported by appropriate tools. 11

GANGA is an easy-to-use frontend for the configuration, execution, and management of computational tasks. The implementation uses an object-oriented design in PYTHON [1]. It started as a project to serve as a Grid user interface for data analysis within the ATLAS [3] and LHCb [4] experiments in High Energy Physics where large communities of physicists need access to Grid resources for data mining and simulation tasks. A list of projects which supported the development of GANGA may be found in section 10.

GANGA provides a simple but flexible programming interface that can be used 19 either interactively at the PYTHON prompt, through a Graphical User Inter-20 face (GUI) or programmatically in scripts. The concept of a *job* component is 21 essential as it contains the full description of a computational task, including: 22 the code to execute; input data for processing; data produced by the applica-23 tion; the specification of the required processing environment; post-processing 24 tasks; and metadata for bookkeeping. The purpose of GANGA can then be 25 seen as making it easy for a user to create, submit and monitor the progress 26 of jobs. GANGA keeps track of all jobs and their status through a repository 27 that archives all information between independent GANGA sessions. It is pos-28 sible to switch between executing a job on a local computer and executing on 29 the Grid by changing a single parameter of a job object. This simplifies the 30

^{*} Corresponding author

Email address: jakub.moscicki@cern.ch (J.T. Mościcki).

¹ On leave from University of Insbruck, Austria.

progression from rapid prototyping on a local computer, to small-scale tests
 on a local batch system, to the analysis of a large dataset using Grid resources.

In GANGA, the user has programmatic access through an Application Pro gramming Interface (API), and has access to applications locally for quick
 turnaround during development.

GANGA is a user- and application-oriented layer above existing job submission and management technologies, such as Globus [5], Condor [6], Unicore [7] or gLite [8]. Rather than replacing the existing technologies, GANGA allows them to be used interchangeably, using a common interface as the interoperability layer.

It is possible to make GANGA available to a user community with a high 41 level of customisation. For example, an expert within a field can implement a 42 custom application class describing the specific computational task. The class 43 will encapsulate all low-level setup of the application, which is always the 44 same, and only expose a few parameters for configuration of a particular task. 45 The plugin system provided in GANGA means that this expert customisation 46 will be integrated seamlessly with the core of GANGA at runtime, and can be 47 used by an end-user to process tasks in a way that requires little knowledge 48 about the interfaces of Grid or batch systems. Issues such as differences in data 49 access between jobs executing locally and on the Grid are similarly hidden. 50

GANGA may be used as a job management system integrated into a larger system. In this case GANGA acts as a library for job submission and control. In particular, GANGA may be used as a building block for the implementation of Grid Portals which allow users access to Grid functionality through their web browsers in a simplified way. These portals are normally domain specific and allow users of a distributed application to run it on the Grid without needing to know much about Grid tools.

GANGA is licensed under the GNU General Public License² and is available for download from the project website: http://www.cern.ch/ganga. The installation of GANGA is trivial and does not require privileged access or any server configuration. The GANGA installer script provides a self-contained package and most of the external dependencies are resolved automatically. However, GANGA generally does not attempt to install Grid or batch submission tools or the application software³. Typically such software is installed and managed separately by system administrators. Simple configuration files allow customi-

² GANGA is licensed under GPL version 2 or, if preferred by the user, any later version. Details of the GPL are available at http://www.gnu.org/licenses/gpl.html.

 $^{^3\,}$ Some external dependencies, such as NorduGrid submission tools, are automatically installed.

⁶⁶ sation and configuration of GANGA at the level of site, workgroup and user.

⁶⁷ Between January 2007 and December 2008 GANGA was used at 150 sites ⁶⁸ around the world, with 2000 unique users running about 250k GANGA ses-⁶⁹ sions⁴.

In this paper, we describe in section 2 the overall functionality, in section 3 details of the implementation, and in section 4 how the progress of jobs is monitored. Section 5 gives an overview of the Graphical User Interface. In sections 6 and 7 we discuss how GANGA is customised for specific user communities. Interfacing and embedding GANGA in other frameworks is presented in section 8. In appendix A we provide some examples of how the API in GANGA can be used.

77 2 Functionality

GANGA is a user-centric tool that allows easy interaction with heterogeneous 78 computational environments, configuration of the applications and coherent 79 organisation of jobs. GANGA functionality may be accessed by a user through 80 any of several interfaces: a text-based command line in PYTHON, a file-based 81 scripting interface and a graphical user interface (GUI). This reflects the differ-82 ent working styles in different user communities, and addresses various usage 83 scenarios such as using the GUI for training new users, the command line to 84 exploit advanced use-cases, and scripting for automation of repetitive tasks. 85 For GANGA sessions the current usage fractions are 55%, 40% and 5% respec-86 tively for interactive prompt, scripts and GUI. As shown in Fig. 1, the three 87 user interfaces are built on top of the GANGA Public Interface (GPI) which 88 in turn provides access to the GANGA core implementation. 89

A job in GANGA is constructed from a set of components. All jobs are required 90 to have an application component and a backend component, which define 91 respectively the software to be run and the processing system to be used. 92 Many jobs also have input and output dataset components, specifying data 93 to be read and produced. Finally, computationally intensive jobs may have a 94 splitter component, which provides a mechanism for dividing into independent 95 subjobs, and a merger component, which allows for the aggregation of subjob 96 outputs. The overall component structure of a job is illustrated in Fig. 2. 97

⁹⁸ By default, the GPI exposes a simplified, top-level view suitable for most ⁹⁹ users in their everyday work, but at the same time allows for the details of

 $^{^4\,}$ The usage information was collected from a voluntary usage reporting system implemented in GANGA.



Fig. 1. The overall architecture of GANGA. The user interacts with the GANGA Public Interface (GPI) via the Graphical User Interface (GUI), the Command-Line Interface in Python (CLIP), or scripts. Plugins are provided for different application types and backends. All jobs are stored in the repository.

underlying systems to be exposed if needed. An example interactive GANGA
 session using the GPI is given in Appendix A.

GANGA prevents modification by the user of a submitted job. However, a copy of the job may easily be created and the copy can be modified. GANGA monitors the evolution of submitted jobs and categorises them into the simplified states *submitted*, *running*, *completed*, *failed* or *killed*.

All job objects are stored in a job repository database, and the input and output files associated with the jobs are stored in a file workspace. Both the repository and the workspace may be in a local filesystem or on a remote server.

A large computational task may be split into a number of subjobs automati-110 cally according to user-defined criteria and the output merged at a later stage. 111 Each subjob will execute on its own and the merging of the output will take 112 place when all have finalised. The submission of subjobs is automatically op-113 timised if the backend component supports bulk job submission. For exam-114 ple, when submitting to the gLite workload management system [8] the job 115 collection mechanism is used transparently to the user. Job splitting func-116 tionality provides a flat list of subjobs suitable for parallel processing of fully 117 independent workloads. However, certain backends allow users to make use of 118



Fig. 2. A set of components in GANGA can be combined to form a complete job. The application to run and the backend where it will run are mandatory while all other components are optional.

¹¹⁹ more-sophisticated parallelisation schemes, for example the Message Passing
¹²⁰ Interface (MPI) [8]. In this case, Ganga may be used to manage collections of
¹²¹ subjobs corresponding to MPI processes.

The GPI allows frequently used job configurations to be stored as templates, so that they may easily be reused, and allows jobs to be labelled and organised in a hierarchical jobtree.

GANGA has built-in support for handling user credentials, including classic Grid proxies, proxies with extensions for a Virtual Organisation Management Service (VOMS) [12], and Kerberos [13] tokens for access to an Andrew filesystem (AFS) [14]. A user may renew and destroy the credentials directly using the GPI. GANGA gives an early warning to a user if the credentials are about to expire. The minimum credential validity and other aspects of the credential management are fully configurable.

GANGA supports multiple security models. For local and batch backends, the authentication and authorisation of the users is based on the local security infrastructure including user name and network authentication protocols such as Kerberos. Grid security infrastructure (GSI) [15] provides for security across organizational boundaries for the Grid backends. Different security models are encapsulated in pluggable components, which may be simultaneously used in 138 the same GANGA session.

A Robot has been implemented for repetitive use-cases. It is a GPI script that 139 periodically executes a series of actions in the context of a GANGA session. 140 These actions are defined by implementations of an action interface. Without 141 programming, the driver can be configured using existing action implementa-142 tions to submit saved jobs, wait for the jobs to complete, extract data about 143 the jobs to an XML file, generate plain text or HTML summary reports, and 144 email the reports to interested parties. Custom actions can easily be added by 145 either extending or aggregating the existing implementations or implement-146 ing the action interface directly, allowing for a diverse variety of repetitive 147 use-cases. An example is given in section 6.1. 148

Details of the different kinds of GANGA component are given below, along
with generic examples. More specialised components, designed for a particular
problem domain, are considered in sections 6 and 7.

152 2.1 Application components

The application component describes the type of computational task to be 153 performed. It allows the characteristics and settings of some piece of software 154 to be defined, and provides methods specifying actions to be taken before and 155 after a job is processed. The pre-processing (configuration) step typically in-156 volves examination of the application properties, and may derive secondary 157 information. For example, intermediate configuration files for the application 158 may be created automatically. The post-processing step can be useful for val-159 idation tasks such as determining the validity of the application output. 160

¹⁶¹ The simplest application component (Executable) has three properties:

¹⁶² exe : the path to an executable binary or script;

¹⁶³ args: a list of arguments to be passed to the executable;

¹⁶⁴ env : a dictionary of environment variables and the values they should be ¹⁶⁵ assigned before the executable is run.

The configuration method carries out integrity checks – for example ensuring
that a value has been assigned to the exe property.

168 2.2 Backend components

A backend component contains parameters describing the behaviour of a pro cessing system. The list of parameters can vary significantly from one system

to another, but can include, for example, a queue name, a list of requested
sites, the minimum memory needed and the processing time required. In addition, some parameters hold information that the system reports back to the
user, for example the system-specific job identifier and status, and the machine
where a job executed.

A backend component provides methods for submitting jobs, and for cancelling jobs after submission, when this is needed. It also provides methods for updating information on job status, for retrieving output of completed jobs and for examining files produced while a job is running.

Backend components have been implemented for a range of widely used processing systems, including: local host, batch systems (Portable Batch System
(PBS) [16], Load Sharing Facility (LSF) [17], Sun Grid Engine (SGE) [18],
and Condor [19]), and Grid systems, for example based on gLite [8], ARC [20]
and OSG [21]. Remote backend component allows jobs to be launched directly
on remote machines using ssh.

As an example, the batch backend component defines a single property thatmay be set by the user:

queue : name of queue to which job should be submitted, the system
 default queue being used if this left unspecified,

¹⁹⁰ and defines three properties for storing system information:

191 id : job identifier;

¹⁹² status : status as reported by batch system;

¹⁹³ actualqueue: name of queue to which job has been submitted.

In addition, a remote-backend component allows a job defined in a GANGA session running on one machine to be submitted to a processing system known to a remote machine to which the user has access. For example, a user who has accounts on two clusters may submit jobs to the batch system of each from a single machine.

199 2.3 Dataset components

Dataset components generally define properties that uniquely identify a particular collection of data, and provide methods for obtaining information about it, for example its location and size. The details of how data collections are described can vary significantly from one problem domain to another, and the only generic dataset component in GANGA represents a null (empty) dataset. Other dataset components are specialised for use with a particular application, 206 and so are discussed later.

A strict distinction is made between the datasets and the sandbox (job) files. The former are the files or databases which are stored externally. The sandbox consists of files which are transferred from the user's filesystem together with the job. The sandbox mechanism is designed to handle small files (typically up to 10MB) while the datasets may be arbitrarily large.

212 2.4 Splitter components

Splitter components allow the user to specify the number of subjobs to be created, and the way in which subjobs differ from one another. As an example, one splitter component (ArgSplitter) deals with executing the same task many times over, but changing the arguments of the application executable each time. It defines a single property:

218 args: list of sets of arguments to be passed to an application.

219 Specialised splitters deal with creating subjobs that process different parts of 220 a dataset.

221 2.5 Merger components

Merger components deal with combining the output of subjobs. Typical out-222 put includes files containing data in a particular format, for example text 223 strings or data representing histograms. As examples, one merger component 224 (TextMerger) concatenates the files of standard output and error returned 225 by a set of subjobs, and another (RootMerger) sums histograms produced 226 in ROOT format [22]. Merging may be automatically performed in the back-227 ground when GANGA retrieves the job output or it may be controlled manually 228 by the user. 229

230 3 Implementation

In this section we provide details of the actual implementation of some of themost important parts of GANGA.



Fig. 3. A component class implements one of the abstract interfaces corresponding to the different parts of a job.

233 3.1 Components

Job components are implemented as plugin classes, imported by GANGA at 234 start-up if enabled in a user configuration file. This means that users only see 235 the components relevant to their specific area of work. Plugins developed and 236 maintained by the GANGA team are included in the main GANGA distribution 237 and are upgraded automatically when a user installs a newer GANGA version. 238 Currently, the list includes around 15 generic plugins and around 20 plugins 239 specific to ATLAS and LHCb. Plugins specific to other user communities need 240 to be installed separately but could easily be integrated into the main GANGA 241 distribution. 242

Plugin development is simplified by having a set of internal interfaces and a mechanism for generating proxy classes [23]. Component classes inherit from an interface class, as seen in Fig. 3. Each plugin class defines a schema, which describes the plugin attributes, specifying type (read-only, read-write, internal), visibility, associated user-convenience filters and syntax shortcuts.

The user does not interact with the plugin class directly but rather with an automatically generated proxy class, visible in the GPI. The proxy class only includes attributes defined as visible in the schema and methods selected for
export in the plugin class. This separation of the plugin and proxy levels is very
flexible. At the GPI level, the plugin implementation details are not visible;
all proxy classes follow the same design logic (for example, copy-by-value);
persistence is automatic, session-level locking is transparent. In this way the
low-level, internal API is separated from the user-level GPI.

The framework does not force developers to support all combinations of appli-256 cations and backends, but only the ones that are meaningful or interesting. To 257 manage this, the concept of a *submission handler* is introduced. The submis-258 sion handler is a connector between the application and backend components. 259 At submission time, it translates the internal representation of the application 260 into a representation accepted by a specific backend. This strategy allows in-261 tegration of inherently different backends and applications without forcing a 262 lowest-common-denominator interface. 263

Most of the plugins interact with the underlying backends using shell commands. This down-to-earth approach is particularly useful for encapsulating the environments of different subsystems and avoiding environment clashes. In verbose mode, GANGA prints each command executed so that a user may reproduce the commands externally if needed. Higher-level abstractions such as JSDL [24], OGSA-BES [25] or SAGA API [26] are not currently used, but specific backends that support these standards could readily be added.

271 3.2 Job persistence

The *job repository* provides job persistence in a simple database, so that any subsequent GANGA session has access to all previously defined jobs. Once a job is defined in a GANGA session it is automatically saved in the database. The repository provides a bookkeeping system that can be used to select particular jobs according to job metadata. The metadata includes such parameters as job name, type of application, type of submission backend, and job status. It can readily be extended as required.

GANGA supports both a local and a remote repository. In the case of the 279 former, the database is stored in the local file system, providing a standalone 280 solution. In the case of the latter, the client accesses an AMGA [28] metadata 281 server. The remote server supports secure connections with user authentication 282 and authorisation based on Grid certificates. Performance tests of both the 283 local and remote repositories show good scalability for up to 10 thousand 284 jobs per user, with the average time of individual job creation being about 285 0.2 seconds. There is scope for further optimisation in this area by taking 286 advantage of bulk operations and job loading on demand. 287

The job repository also includes a mechanism to support schema migration, allowing for evolution in the schema of plugin components.

290 3.3 Input and output files

GANGA stores job input and output files in a *job workspace*. The current implementation uses the local file system, and has a simple interface that allows transparent access to job files within the GANGA framework. These files are stored for each job in a separate directory, with sub-directories for input and output and for each subjob.

Users may access the job files directly in the file-system or using GANGA commands such as job.peek(). Internally, GANGA handles the input and output files using a simple abstraction layer which allows for trivial integration of additional workspace implementations. Tests with a prototype using a WebDav [30] server have shown that all workspace data related to a job can be accessed from different locations. In this case, a workspace cache remains available on the local file system.

The combination of a remote workspace and a remote job repository effectively creates a roaming profile, where the same GANGA session can be accessed at multiple locations, similar to the situation for accessing e-mail messages on an IMAP [31] server.

307 4 Monitoring

GANGA provides two types of monitoring: the internal monitoring updates the user with information on the progress of jobs, and the external monitoring deals with information from third-party services.

311 4.1 Internal monitoring

GANGA automatically keeps track of changes in job status, using a monitoring procedure designed to cope with varying backend response times and load capabilities. As seen in Fig. 4, each backend is polled in a different thread taken from a pool, and there is an efficient mechanism to avoid deadlocks from backends that respond slowly. The poll rate may be set separately for each backend.



Backend system (Grid, Batch,Condor,...)

Fig. 4. The internal monitoring updates the status of jobs using a pool of threads running in the GANGA core. Additional monitoring thread runs in a job wrapper and sends the monitoring information to external services.

The monitoring sub-system also keeps track of the remaining validity of authentication credentials, such as Grid proxies and Kerberos tokens. The user is notified that renewal is required, and if no action is taken then GANGA is placed in a state where operations requiring valid credentials are disabled.

322 4.2 External Monitoring

GANGA's external monitoring provides a mechanism for dynamically adding third-party monitoring sensors, to allow reporting of different metrics for running jobs.

The monitoring sensors can be inserted both on the client side - where GANGA runs - and on the remote environment (worker node) where the application runs, allowing the user to follow the entire execution flow. Monitoring events are generated at job submission time, at startup, periodically during execution, and at completion.

Individual application and backend components in GANGA can be configured
to use different monitoring sensors, allowing collection of both generic execution information and application-specific data.

Use is currently made of two implementations of external monitoring sensors. One is the ATLAS Dashboard application monitoring [32]. Another is a custom service that allows the GANGA user to examine job output in real-time on the Grid. This streaming service is not enabled by default, but must be set up for each user community separately, and may then be requested by a user for specific jobs.

340 5 Graphical User Interface

The GANGA Graphical User Interface (GUI), shown in Fig. 5 and built using PyQt3 [33], makes available all of the job-management functionality provided at the level of the GANGA Public Interface. The GUI incorporates various convenience features, and its multi-threaded nature results in a degree of parallelism not possible at the command line: job monitoring and most jobmanagement actions run concurrently, ensuring a good response time for the user.

	[[[[
V	status	name	application	exe filename	backend	ba	Job (
	failed	G4Brachy	Executable	File	LCG	No	status = 'completed',
	completed		Executable	echo	LCG	htt	name = ",
	completed	DIANEWorkerAgent	Executable	File	LCG	htt	outputdir = //afs/cern.ch/user/m/moscicki
	completed	DIANEWorkerAgent	Executable	File	LCG	htt	outputair = /ais/cern.ch/user/m/moscicki
	completed	DIANEWorkerAgent	Executable	File	LCG	htt	id = 1
1	completed	DIANEWorkerAgent	Executable	File	LCG	htt	info = 1,
2	completed	DIANEWorkerAgent	Executable	File	Local	25	submit counter = 1
3	completed		Executable	echo	Local	12).
4	new		Executable	echo	Local	-1	inputdata = None ,
5	new		Executable	echo	Local	-1	merger = None ,
6	failed		Executable	echo	Local	93	inputsandbox = [] ,
	failed		Executable	echo	Local	15	application = Executable (
8	completed		Executable	echo	Local	19	exe = 'echo' ,
						1	env = {} ,
					Fin	1 🧶 🖸 🔰	4
ving job	5						
ving job	6						
ving job	57						
ving job	8						
ving job	9						
ving job	0 10						

Fig. 5. GANGA graphical user interface (GUI). The overview of jobs can be seen to the left, and the details of an individual job are to the right.

The job monitoring window takes centre stage, with job status and other monitored attributes displayed in table format. Other features include subjob monitoring, subjob folding/hiding, a job-details display drawer, a logicalcollections drawer, and a text-based job-search facility. Many characteristics of
the monitoring window can be customised, allowing, for example, selection of
the job attributes to be monitored, and of the colours used to denote different
job states.

The construction of a job, entailing selection of the required plugins and the 355 entry of attribute values, is achieved from a job-builder window. This displays 356 a foldable tree of job attributes, and associated data-entry widgets. The tree 357 and widgets are generated dynamically based on plugin schemas, ensuring 358 that the GUI automatically supports user-defined plugins without any change 359 being needed to the GUI code. To assist with data entry, drop-down menus list 360 allowed values, wherever these are defined; and tool tips provide explanations 361 of individual job attributes. The job-builder window also features tool buttons 362 for performing a wide range of job-related actions, including creation, saving, 363 copying, submission, termination and removal. Finally, a multifunction Extras 364 tool button provides access to arbitrary additional functionality implemented 365 in the plugins. 366

The GUI also has a scriptor window, providing a favourite-scripts collection, a job-script editor and an embedded PYTHON session. The favourite-scripts collection allows frequently used GANGA scripts to be created, imported, exported and cloned; the job-script editor facilitates quick modification and execution of scripts; and the embedded PYTHON session allows interactive use of GANGA commands.

Finally, a scrollable log window collects and displays all messages generatedby GANGA.

³⁷⁵ 6 Use in experiments at the Large Hadron Collider

The ATLAS and LHCb experiments aim to make discoveries about the fun-376 damental nature of the Universe by detecting new particles at high energies, 377 and by performing high-precision measurements of particle decays. The ex-378 periments are located at the Large Hadron Collider (LHC) at the European 379 Laboratory for Particle Physics (CERN), Geneva, with first particle colli-380 sions (events) expected in 2009. Both experiments require processing of data 381 volumes of the order of petabytes per year, rely on computing resources dis-382 tributed across multiple locations, and exploit several Grid implementations. 383 The data-processing applications, including simulation, reconstruction and fi-384 nal analysis for the experiments, are based on the C++ GAUDI/ATHENA [34] 385 framework. This provides core services, such as message logging, data access, 386 histogramming, and a run-time configuration system. 387

The data from the experiments will be distributed at computing facilities around the world. Users performing data analysis need an on-demand access mechanism to allow rapid pre-filtering of data based on certain selection criteria so as to identify data of specific interest.

The role of GANGA within ATLAS and LHCb is to act as the interface for data analysis by a large number of individual physicists. GANGA also allows for the easy exchange of jobs between users, something that can otherwise be difficult because of the complex configuration of analysis jobs.

396 6.1 The LHCb experiment

The LHC*b* experiment is dedicated to studying the properties of *B* mesons (particles containing the *b* quark) and in this section we describe the way in which GANGA interacts with the application and backend plugins specific to LHC*b*.

In a typical analysis, users supply their own shared libraries, containing userwritten classes, and these are loaded at run-time. The LHC*b* applications are driven by a configuration file, which includes definitions of the libraries to load, non-default values for object parameters, the input data to be read, and the output to be created.

GANGA includes an application component for GAUDI-based applications to
simplify the task of performing an analysis. During the configuration stage,
and before job submission, the application component undertakes the following
tasks:

- it locally sets up the environment for the chosen application;
- it determines the user-owned shared libraries required to run the job;
- it parses the configuration file supplied, including all its dependencies;
- it uses information obtained from the configuration file to determine the input data required and the outputs expected;
- it registers the inputs and outputs with the submission backend.

The user, then, only needs to specify the name and version of the application to run, and the configuration file to be used.

⁴¹⁸ Code under development by a user may contain bugs that cause runtime errors ⁴¹⁹ during job execution. The transparent switching between processing systems ⁴²⁰ when using GANGA means that debugging can be performed locally, with ⁴²¹ quick response time, before launching a large-scale analysis on the Grid, where ⁴²² response times tend to be longer. Some studies in LHC*b*, rather than being based on GAUDI, are performed using the ROOFIT [37] framework, most notably studies that make use of simplified event simulations. Jobs for these studies require large amounts of processing power, but do not require input data and produce only small amounts of output. This makes them very easy to deploy on the Grid, with support in GANGA provided by a generic ROOT [22] application component.

In the LHCb computing model [29], Grid jobs are routed through the DIRAC [35] 429 workload management system (WMS). DIRAC is a pilot-based system where 430 user jobs are queued in the WMS server and the server submits generic pi-431 lot scripts to the Grid. Each pilot queries the WMS for a job with resource 432 requirements satisfied by the machine where the pilot script is running. If a 433 compatible job is available, it is pulled from the WMS and started. Otherwise, 434 the pilot terminates and the WMS sends a new pilot to the Grid. This system 435 improves the reliability of the Grid system as seen by the user. GANGA pro-436 vides a DIRAC backend component that supports submission of jobs to the 437 DIRAC WMS, making use internally of DIRAC's PYTHON API [36]. 438

A splitter component implemented specifically for LHCb is able to divide the
analysis of a large dataset into many smaller subjobs. During the splitting, a
file catalogue is queried to ensure that all data associated with an individual
subjob is available in its entirety at a minimum of one location on the Grid.
This gives significant optimisation, as it avoids subjobs having to copy data
across the network before an analysis can start.

In total, above 300k user jobs finished successfully in 2008 with a total CPU consumption of 87 CPU years. The jobs ran at a total of 140 Grid sites across the globe. The system was responsive to a highly irregular usage pattern and spikes of several thousand simultaneous jobs were observed during the year. This usage is expected to rise dramatically after the start of the LHC*b* data taking.

The Robot in GANGA is used within LHCb for *end-to-end* testing of the distributed analysis model. It submits a representative set of analysis jobs on a daily basis, monitors their progress, and checks the results produced. The overall success rate and the time to obtain the results is recorded and published on the web. The Robot monitors this information, producing statistics on the long-term system performance.

457 6.2 The ATLAS experiment

458 ATLAS is a general-purpose experiment, designed to allow observation of new 459 phenomena in high-energy proton-proton collisions. The distributed analysis model is part of the ATLAS computing model [38]
which requires that data are distributed at various computing sites, and user
jobs are sent to the data.

An ATLAS analysis job typically consists of a PYTHON or shell script that configures and runs user algorithms in the ATHENA framework [38], reads and writes event files, and fills histograms/n-tuples. More-interactive analysis may be performed on large datasets stored as n-tuples.

There are several scenarios relevant for a user analysis. Some analyses require a fast response time and a high level of user interaction, for which the parallel ROOT facility PROOF [41] is well suited. Other analyses require a low level of user interaction, with long response times acceptable, and in these cases GANGA and Grid processing are ideal.

Analysis jobs can produce large amounts of data, which may initially be stored 472 at a single Grid site, and may subsequently need to be transferred to other 473 machines. This is supported in ATLAS by the Distributed Data Management 474 system DQ2 [39]. This provides a set of services for moving data between 475 Grid-enabled computing facilities, and maintains a series of databases that 476 track the data movements. The vast amounts of data involved are grouped 477 into datasets, based on various criteria, for example physics characteristics, to 478 make queries and retrievals more efficient. 479

480 6.2.1 ATLAS Grid infrastructures

The ATLAS experiment employs three Grid infrastructures for user analysis and for collaboration-wide event simulation and reconstruction. These are the Grid developed in the context of Enabling Grids for e-Science (EGEE, mainly Europe) [42], accessed using gLite middleware [8], the Open Science Grid (OSG, mainly North America) [21], accessed using the PanDA system [40], and NorduGrid (mainly Nordic countries) [43], accessed using the ARC middleware [20]. GANGA seamlessly submits jobs to all three Grid flavours.

488 6.2.2 ATLAS user analysis

A typical ATLAS user analysis consists of an event-selection algorithm developed in the Athena framework. Large amounts of data are filtered to identify
events that meet certain selection criteria. The events of interest are stored in
files grouped together as datasets in the DQ2 system. The GANGA components
for Athena jobs include the following functionality:

• During job submission, DQ2 is queried for the file content and location of the dataset to be analysed. The number of possible Grid sites is then ⁴⁹⁶ restricted to the dataset locations.

A job can be divided into several subjobs, each processing a given number
 of files from the full dataset.

In a Grid job, after the ATHENA application has completed, the user output is stored on the storage element of the site where the job was run, and is registered in DQ2.

In the second half of 2008, more than 4×10^5 Grid jobs were submitted through GANGA by ATLAS users. Following a procedure similar to that of LHC*b*, the GANGA Robot submits test jobs daily to ATLAS Grid sites. Test results are used to guide users to sites that are performing well, avoiding job failures on temporarily misconfigured sites.

507 6.2.3 ATLAS small-scale event simulations

In addition to data analysis, users sometimes need to simulate event samples 508 of the order of a few tens of thousands of events. The A then a MC applica-500 tion component has been developed to integrate software used in the offi-510 cial ATLAS system for event simulation. This component consists of a set of 511 PYTHON classes that together handle input parameters, input datasets and 512 output datasets for the three production steps: event generation, detector sim-513 ulation, and event reconstruction. As in the case of user analysis, datasets are 514 managed by the DQ2 system. 515

516 7 Other usage areas

GANGA offers a flexible and extensible interface that make it useful beyond the original scope of particle-physics applications in the ATLAS and LHC*b* experiments. Here we provide details of just a few of the other contexts in which GANGA has been used.

521 7.1 Enabling industrial-scale image retrieval

Imense Ltd⁵, a Cambridge-based startup company, has implemented a novel image retrieval-system (Fig. 6), featuring automated analysis and recognition of image content, and an ontological query language. The proprietary image analysis, developed from published research [44], includes recognition of visual properties, such as colour, texture and shape; recognition of materials, such as grass or sky; detection of objects, such as human faces, and determination of

⁵ http://imense.com

their characteristics; and classification of scenes by content, for example beach,
forest or sunset. The system uses semantic and linguistic relationships between
terms to interpret user queries and retrieve relevant images on the basis of the
analysis results. Moreover, the system is extensible, so that additional image
classification modules or image context and metadata can easily be integrated
into the index.



Fig. 6. Schematic representation of the image-retrieval system developed by Imense Ltd. Image characteristics are determined by applying feature-extraction algorithms, and an ontological query language bridges the semantic gap between terms that might be employed in a user query and terms understood by the processing system.

533

⁵³⁴ By using the GANGA framework for job submission and management, it has ⁵³⁵ been possible to port and deploy a large part of Imense's image-analysis tech-⁵³⁶ nology to the Grid and build a searchable index for more than twenty-million ⁵³⁷ high-resolution photographic images.

The processing stages for the image-search system – image analysis and in-538 dexing – are intrinsically sequential. Analysis has been parallelised at the level 539 of single images or small subsets of images. Each image can therefore be pro-540 cessed in isolation on the Grid, with this processing usually taking a few to 541 ten seconds. In order to minimise overheads, images are grouped in sets of a 542 few hundred per job submitted through GANGA. Results of the image pro-543 cessing and analysis are passed back to the submission server once a job has 544 successfully completed. 545

Support for Imense has been added to GANGA through the implementation of
two specialised components: an application component that deals with running
the image-processing software, and a dataset component for taking care of the

output. As usual with GANGA, the jobs can run both locally and on the Grid,
giving maximum flexibility.

At runtime, images are retrieved and segmented one at a time, all of the images are classified, and finally an archive is created of the output files (several per input image). The archive is returned using the sandbox mechanism in GANGA when using the Local backend, and is uploaded to a storage element when using the Grid LCG backend.

The specialised dataset component provides methods for downloading a results archive from a storage element, and for unpacking an archive to a destination directory. These methods are invoked automatically by GANGA when an image-processing job completes: the effect for the user is that a list of images is submitted for processing and results are placed in the requested output location independently of the backend used.

562 7.2 Smaller collaborations in High Energy Physics

Large user communities, such as ATLAS and LHCb, profit from encapsulating 563 shared use cases as specialised applications in GANGA. In contrast, individual 564 researchers or developers in the context of rapid prototyping activities may 565 opt to use generic application components. In such cases, GANGA still provides 566 the benefits of bookkeeping and a programmatic interface for job submission. 567 As an example of this way of working, a small community of experts in the 568 design of gaseous detectors use GANGA to run the GARFIELD [45] simulation 569 program on the Grid. A GANGA script has been written that generates a chain 570 of simulation jobs using the GARFIELD generator of macro files and GANGA's 571 Executable application component. The GARFIELD executables, and a few 572 small input files, are placed in the input sandbox of each job. Histograms and 573 text output are then returned in the output sandbox. This simple approach 574 allowed integration of GARFIELD jobs in GANGA in just a few hours. 575

576 7.3 GANGA integrated with lightweight Grid middleware

The open-plugin architecture of GANGA allows easy integration of additional Grid middleware, as has been achieved, for example, with the ARC (Advanced Resource Connector) Grid middleware [20]. This is a product of the NorduGrid project [43], and is used by many academic institutions in the Nordic countries and elsewhere.

ARC jobs are accepted and brokered by a Grid manager, running at site level,
 and resource lookup is done through load balancing and runtime environments

advertised by individual sites. File storage and access is 'cloudy', meaning that all files registered in Grid-wide catalogues are accessible to all worker nodes. File transfers are handled by the Grid manager, between job acceptance and execution. ARC-connected resources are used e.g. by researchers in bioinformatics, genomics, meteorology, in addition to high-energy physics.

GANGA has been interfaced to ARC through a backend, which translates GANGA input into ARC-readable xRSL language. The ARC user client is lightweight, and binaries are provided as an external library at GANGA install time. The main user of this integration is the ATLAS experiment (see sec. 6.2), where it is the main user access portal to one of the experiment's three main computing grids. Further collaboration between ARC and GANGA is envisaged, to employ GANGA as a fully featured frontend to ARC.

596 8 Interfacing to other frameworks

The GANGA Public Interface constitutes an API for generic job submission 597 and management. As a result, GANGA may be programmatically interfaced to 598 other frameworks, and used as a convenient abstraction layer for job manage-599 ment. GANGA has been used in combination with DIANE [46], a lightweight 600 agent-based scheduling layer on top of the Grid, in a number of scientific ac-601 tivities. These have included: dosimetry-related simulation studies in medical 602 physics [47]; regression testing of the Geant 4 [48] detector-simulation toolkit; 603 in-silico molecular docking in searches for new drugs against potential variants 604 of an influenza virus [49]; telecommunication applications [50]; and theoreti-605 cal physics [51]. The DIANE worker agents are executed as GANGA jobs, so 606 that resource usage may be controlled by the user from the GANGA interface. 607 This approach allows the efficiency of the DIANE overlay scheduling system 608 to be combined with the well-structured job management offered by GANGA, 609 as well as combining Grid and non-Grid resources under a uniform interface. 610 Also, this allows the efficient implementation of low-latency access to Grid 611 resources and improvements to responsiveness when supporting on-demand 612 computing and interactivity [52]. 613

GANGA may be embedded in web-based services such as the bio-informatics portal developed by ASGC, Taipei. The portal is fully customized for analysis of candidate drugs against avian flu. The portal engine delegates job management to the embedded DIANE/GANGA framework, as shown in Fig. 7. Following this approach, users can switch between different resources, or access heterogeneous computing environments through a single same web interface.



Fig. 7. Ganga as a job management component embedded in DIANE, with an application portal.

620 9 Conclusion

GANGA has been presented as a tool for job management in an environment of heterogeneous resources and is particularly suited to the Grid paradigm that has emerged in large-scale distributed computing. GANGA makes it easy to define a computational task that can be executed locally for debugging, and subsequently be run on the Grid, for large scale data mining. We have shown how GANGA simplifies task specification, takes care of job submission, monitoring and output retrieval, and provides an intuitive bookkeeping system.

We have demonstrated the advantages of having a well-defined API, which can be used interactively at the PYTHON prompt, through a GUI or programmatically in scripts. By virtue of its plugin system, GANGA is readily extended and customised to meet the requirements of new user communities. Examples of GANGA usage have been provided in particle physics, medical physics and image processing.

Existing command-line submission interfaces, such as gLite, tend to include only limited usability features. Some higher level tools, for example GridWay[53], present jobs as if they were Unix processes and corresponding command line utilities. Interfaces based on Condor job-submission scripts have also been developed [54]. A distinctive feature of GANGA is that it may easily be adapted to different styles of working, allowing simultaneous use of three different interfaces. GANGA also provides a higher level of abstraction than most jobmanagement tools, and allows a user to focus on solving the domain-specific problems, rather than changing their way of working each time they switch to a new processing system.

GANGA has a large user base and is in active development. GANGA is a tool which may easily be used to support new scientific or commercial projects on a wide range of distributed infrastructures.

647 10 Acknowledgements

The development of GANGA has been supported by the GridPP project in 648 the United Kingdom [2], with funding from the Science and Technology Facil-649 ities Council (STFC) and its predecessor, the Particle Physics and Astronomy 650 Research Council (PPARC); by the D-Grid project in Germany, with fund-651 ing from the Bundesministerium für Bildung und Forschung (BMBF); and by 652 the project for Enabling Grids for E-science (EGEE), co-funded by the Eu-653 ropean Commission (contract number INFSO-RI-031688) through the Sixth 654 Framework Programme. 655

GANGA has received contributions over the years from a number of individuals.
Particular thanks are due to David Adams, Marcello Barisonzi, Mike Kenyon,
Wim Lavrijsen, Janusz Martyniak, Pere Mato, Caitriana Nicholson, Rebecca
Ronke, Vladimir Romanovski, David Tuckett, Ruth Dixon del Tufo, Craig
Tull.

The developers would also like to thank the large number of users, from both within and outside particle physics, for their valuable suggestions for improving GANGA, and for their help in debugging problems.

664 A Examples

Below we give a set of examples of working with GANGA. For ease of reading, PYTHON keywords are in bold. First we look at a complete GANGA session.

```
~ % ganga
668
       *** Welcome to Ganga ***
669
       Version: Ganga-5-1-0
670
671
       Documentation and support: http://cern.ch/ganga
672
       Type help() or help('index') for online help.
673
       This is free software (GPL), and you are welcome to redistribute
674
       it under certain conditions; type license() for details.
[1]: j=Job(name='MyJob')  # Create a default job
[2]: j.submit()  # Submit the job
675
676
677
678
       # wait for the monitoring
```

680 : j.peek('stdout') # Look at the output : j=j.copy(name='GridJob') # Make a copy of the job : j.backend=LCG() # Change backend to the Grid 681 [4] [5]682 683 684 6 j.submit() # Submit the job 685 : jobs List jobs[7] # 686 ...job listing... [8]: Exit 687 # Quit Ganaa 688

In the next example, we create a job for analysis of LHCb data. A splitter is used to divide the analysis between subjobs. Data are assigned using logical identifiers, and the DIRAC WMS ensures that subjobs are sent to locations where the required data are available.

```
j=Job(application=DaVinci(), backend=Dirac())
693
     [1].
     [2]: j.inputdata=LHCbDataset(files=[ # Data to read
694
695
                'LFN:/foo.dst'
                'LFN:/bar.dst'
696
     . . .
697
               many more data files ])
698
     [3]
          j.splitter=DiracSplitter()
                                                # We want subjobs
699
     [4]: j.submit()
700
701
     Job submission output
```

Here, we use the fact that standard PYTHON commands are available at the
 GANGA prompt, and print information on subjobs.

```
704 # Status of jobs and where they ran
705 [5]: for subjob in j.subjobs:
706 ... print subjob.status, subjob.actualCE
707
708 42
709 # Find backend identifier of all failed jobs
710 [6]: for j in jobs.select(status='failed'):
711 ... print j.backend.id
713 42
```

⁷¹⁴ Groups of jobs may be accessed and manipulated using simple methods:

```
715 [1]: jobs.select(status='failed').resubmit()
716 [2]: jobs.select(name='testjob').kill()
717 [3]: newjobs = jobs.select(status='new')
718 [4]: newjobs.select(name='urgent').submit()
```

719 References

- [1] G. van Rossum and F.L. Drake Jr., *The Python language reference manual: revised and updated for version 2.5* (Network Theory Limited, Bristol, 2006).
- P.J.W. Faulkner et al. [GridPP Collaboration], GridPP: development of the UK
 computing Grid for particle physics, J. Phys. G: Nucl. Part. Phys. 32 N1.
- [3] G. Aad et al. [ATLAS Collaboration], The ATLAS Experiment at the CERN
 Large Hadron Collider, JINST 3 (2008) S08003.
- [4] A.A. Alves Jr. et al. [LHCb Collaboration], The LHCb Detector at the LHC,
 JINST 3 (2008) S08005.
- [5] I. Foster, C. Kesselman and S. Tuecke, *The Anatomy of the Grid: Enabling Scalable Virtual Organizations*, International J. Supercomputer Applications, **15** (3), 2001.

- [6] D. Thain, T. Tannenbaum, and M. Livny, *Distributed Computing in Practice: The Condor Experience* Concurrency and Computation: Practice and Experience, **17** (2-4), pp 323-356, 2005.
- A. Streit et al., UNICORE From Project Results to Production Grids Grid
 Computing: The New Frontiers of High Performance Processing Advances in
 Parallel Computing 14, Elsevier, 2005, pp. 357-376
- [8] P. Andreett et al., The gLite workload management system, J. Phys.: Conf. Ser.
 119 (2008) 062007.
- [9] M.P. Thomas et al., Grid portal achitectures for scientific applications, J. Phys.:
 Conf. Ser. 16 (2005) 596.
- [10] M. Li and M. Baker, A review of Grid Portal technology, pp. 126-156 of:
 J.C. Cunha and O.F. Rana (Eds.), Grid computing: software environments and tools (Springer-Verlag London Ltd, 2006).
- [11] M. Snir and S. Otto, MPI-The Complete Reference: The MPI Core, MIT Press
 (1998) ISBN: 0262692155 .
- [12] R. Alfieri et al., From gridmap-file to VOMS: managing authorization in a Grid
 environment, Future Generation Computer Systems 21 (2005) 549.
- [13] B.C. Neumann and T. Ts'o, Kerberos: an authentication service for computer
 networks, IEEE Communications Magazine 32-9 (1994) 33.
- [14] J.H. Morris et al., Andrew: a distributed personal computing environment,
 Commun. ACM 29-3 (1986) 184.
- [15] I. Foster *et al.*, A Security Architecture for Computational Grids Proc. 5th ACM
 Conference on Computer and Communications Security Conference, pp. 83-92,
 1998.
- [16] R.L. Henderson, Job scheduling under the Portable Batch System, pp. 279-294
 of: D.G. Feitelson and L. Rudolph (Eds.), Job scheduling strategies for parallel
 processing [Lecture Notes in Computer Science 949] (Springer, Berlin, 1995).
- [17] U. Schwickerath and V. Lefebure, Usage of LSF for batch farms at CERN,
 J. Phys.: Conf. Ser. 119 (2008) 042025.
- [18] W. Gentzsch, Sun Grid Engine: towards creating a compute power Grid, pp. 3536 of: R. Buyya, G. Mohay and P. Roe (Eds.), Proc. First IEEE/ACM
 International Symposium on Cluster Computing and the Grid (IEEE Computer
 Society, Los Alamitos, CA, 2001).
- [19] D. Thain, T. Tannenbaum and M. Livny, *Distributed computing in practice: the Condor experience*, Concurrency Computat.: Pract. Exper. 17 (2005) 323.
- [20] M. Ellert et al., Advanced Resource Connector middleware for lightweight
 computational Grids, Future Generation Computer Systems 23 (2007) 219.
- ⁷⁶⁸ [21] R. Pordes et al., The Open Science Grid, J. Phys.: Conf. Ser. 78 (2007) 012057.

- [22] R. Brun and F. Rademakers, ROOT an object oriented data analysis
 framework, Nucl. Instrum. Methods A389 (1997) 81.
- [23] E. Gamma *et al.*, *Design patterns: elements of reusable object-orientated software* (Addison-Wesley, 1995).
- [24] Job Submission Description Language (JSDL) Specification, Version 1.0
 http://www.gridforum.org
- 775 [25] OGSA Basic Execution Service http://www.ogf.org
- ⁷⁷⁶ [26] A Simple API for Grid Applications (SAGA) http://www.ogf.org
- ⁷⁷⁷ [27] F. Perez and B.E. Granger, *IPython: a system for interactive scientific* ⁷⁷⁸ *computing*, Computing in Science and Engineering **9-3** (2007) 21.
- [28] B. Koblitz, N. Santos and V. Pose, *The AMGA metadata service*, J. Grid
 Computing 6 (2008) 61.
- [29] R. Antunes-Nobrega *et al.* [LHCb Collaboration], *LHCb computing*, Technical
 Design Report CERN/LHCC 2005-019 LHCb TDR-11 (2005).
- [30] E.J. Whitehead Jr., World Wide Web Distributed Authoring and Versioning
 (WebDAV): an introduction, StandardView 5 (1997) 3.
- [31] P. Heinlein and P. Hartleban, *The book of IMAP: building a mail server with Courier and Cyrus* (No Startch Press, San Francisco, CA, 2008).
- [32] J. Andreeva et al., Dashboard for the LHC experiments, J. Phys. Conf. Ser. 119
 (2008) 062008.
- [33] B. Rempt, *GUI programming with Python: QT edition* (Command Prompt Inc,
 White Salmon, WA, 2001).
- [34] G. Barrand et al., GAUDI a software architecture and framework for building
 HEP data processing applications, Computer Physics Communications 140
 (2001) 45.
- [35] A. Tsaregorodtsev et al., DIRAC: a community grid solution, J. Phys. Conf.
 Ser. 119 (2008) 062048.
- [36] S. Paterson, *LHCb distributed data analysis on the computing Grid*, PhD Thesis,
 University of Glasglow (2006) [CERN-THESIS-2006-053].
- [37] W. Verkerke and D. Kirkby, *The RooFit toolkit for data modeling*, Contribution
 MOLT007 in: Proc. 2003 Conference for Computing in High Energy and Nuclear
 Physics, La Jolla, CA [SLAC eConf C0303241].
- ⁸⁰¹ [38] G. Duckeck *et al.* (Eds.), *ATLAS computing*, Technical Design Report ⁸⁰² CERN/LHCC 2005-022 ATLAS TDR-017 (2005).
- [39] M. Branco *et al. Managing ATLAS data on a petabyte-scale with DQ2*, J. Phys.
 Conf. Ser. **119** (2008) 062017.

- [40] T. Maeno, PanDA: distributed production and distributed analysis system for
 ATLAS, J. Phys. Conf. Ser. 119 (2008) 062036.
- [41] M. Ballintijn *et al.*, *Parallel interactive data analysis with PROOF*, Nucl.
 Instrum. Methods 559 (2006) 13.
- [42] R. Jones, An overview of the EGEE project, pp. 1-8 of: C. Türker, M. Agosti and
 H.-J. Schek (Eds.), Peer-to-peer, Grid, and service-orientation in digital library
 architectures [Lecture Notes in Computer Science 3664] (Springer, Berlin,
 2005).
- [43] M. Ellert et al., The NorduGrid project: using Globus toolkit for building Grid
 infrastructure Nucl. Instrum. Methods A502 (2003) 407.
- ⁸¹⁵ [44] C. Town and D. Sinclair, Language-based querying of image collections on the ⁸¹⁶ basis of an extensible ontology, Image and Vision Computing **22** (2004) 251.
- [45] R. Veenhof, Garfield simulation of gaseous detectors, CERN Program Library
 User Guide W5050 (1984 et seq.).
- ⁸¹⁹ [46] J.T. Mościcki, Distributed analysis environment for HEP and interdisciplinary ⁸²⁰ applications, Nucl. Instrum. Methods **A502** (2003) 426.
- [47] J.T. Mościcki et al., Distributed Geant4 Simulation in Medical and Space
 Science Applications using DIANE framework and the GRID, Nucl. Phys. B
 (Proc. Suppl.) 125 (2003) 327-331
- [48] J.Allison et al., Geant 4 a simulation toolkit, Nucl. Instrum. Methods A506
 (2003) 250.
- ⁸²⁶ [49] H.-C. Lee et al., Grid-enabled high-throughput in silico screening against ⁸²⁷ influenza A neuraminidase, IEEE Trans. NanoBioscience **5-4** (2006) 288.
- ⁸²⁸ [50] J.T. Mościcki *et al.*, *ITU RRC06 on the Grid*, in prep. for Journal of Grid ⁸²⁹ Computing
- [51] J.T. Mościcki *et al.*, *Lattice QCD on the Grid* in prep. for Computer Physics
 Communications
- [52] C. Germain-Renaud, C. Loomis , J. T. Mościcki and R. Texier, Scheduling for
 Responsive Grids, J. Grid Computing 6, (2008) 15-27
- [53] J. Herrera *et al.*, Porting of Scientific Applications to Grid Computing on
 GridWay Scientific Programming 13 4, pp. 317-331, 2005
- [54] R.P. Bruin *et al.*, Job submission to grid computing environments Concurrency
 and Computation: Pract. and Exper. 20 11, pp. 1329-1340 (2008)