



Protopopescu, Dan, Britton, David, and Martyniak, Janusz (2013) *NA62 Grid Monte Carlo Production Tools*. Project Report. CERN.

Copyright © 2013 The Authors

<http://eprints.gla.ac.uk/93052/>

Deposited on: 24 April 2014

NA62 Grid Monte Carlo Production Tools

D. Protopopescu* and D. Britton

School of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, United Kingdom

J. Martyniak

Blackett Laboratory, Imperial College, London SW7 2BW, United Kingdom

(Dated: November 26, 2013)

The NA62 Grid Interface for Monte Carlo production and its related system components are presented. This note is intended as a detailed description of the system for administration purposes, and as a user's manual for Grid production management.

Keywords: NA62; Monte Carlo; Grid

Contents

I. Introduction	1
II. The Experiment	1
III. Grid Simulations	1
IV. The NA62 Grid Interface	2
A. Data storage	2
B. File Transfer Controller	2
C. Authentication	2
D. Monitoring jobs	2
E. Monitoring files	4
F. Production list	4
G. Job submission	6
H. Smartphone App	6
I. Miscellany	6
V. Grid Shifts	8
VI. Maintenance and Support	8
A. Participating sites	8
B. Documentation	8
C. Feedback and bug reports	10
VII. Summary	10
Acknowledgments	10
VIII. Appendix	10
A. Database Tables	10
B. List of scripts	10
C. How does it all work	11
D. Site security	11
E. FAQ	11
References	12

I. INTRODUCTION

The NA62 Grid Interface presented in this note is a web interface that provides job submission and monitoring tools for large-scale NA62 Monte Carlo simulations on the Grid. The NA62 Grid Interface is located at

<http://na62.gla.ac.uk>

While trying to avoid very technical details, this note is meant to serve both as a description of the NA62 Grid Interface features and an offline user's manual.

The material contained herein will be updated if changes are made or new features are added to the NA62 Grid Interface, so please make sure that you have the latest version of this note.

II. THE EXPERIMENT

The NA62 experiment [1] aims to measure the very rare kaon decay $K^+ \rightarrow \pi^+\nu\bar{\nu}$ in order to extract a 10% measurement of the Cabibbo-Kobayashi-Maskawa (CKM) parameter $|V_{td}|$. The experiment will use 400 GeV/c protons from the SPS and aims to collect about 80 to 100 $K^+ \rightarrow \pi^+\nu\bar{\nu}$ events at the Standard Model prediction with a signal to background ratio of 10:1 in two years of data taking starting in the fall of 2014.

III. GRID SIMULATIONS

There are many challenging aspects to the NA26 experiment, and to prepare for those both in terms of hardware and software, extensive Monte Carlo simulations in preparation for the actual measurement are crucial. Large scale simulations take time and running on the Grid offers multiple advantages, among which:

- Distributed CPU resources, hundreds of parallel jobs

*Electronic address: protopop@cern.ch

- Redundant data storage (multiple copies of the output files)

In this context, we think that good production management is of essence, and has to offer:

- Run database (simulation specs, production tags, stats)
- Job tracking capabilities (debugging opportunities)
- File catalogue (run specs, event totals, stats)
- Transparency for the NA62 collaboration

While the Grid infrastructure is in place and there are numerous related tools available, there are no high-level user-friendly Grid production management interfaces that can be used out-of-the-box. It is quite common that experiments write their own software for this.

IV. THE NA62 GRID INTERFACE

For NA62, the NA62 Grid Interface components were written 90% from scratch (on top of the standard Grid infrastructure). The disadvantages of such an approach - development effort, need to document, 'reinventing the wheel' - are far outweighed by the advantages: cohesive system, good understanding of how it works, easy to improve, all needed features.

The NA62 Grid Interface runs on the `na62.gla.ac.uk` server. The displayed content is generated dynamically from a database maintained locally. Most of the code is written in PHP and executed via the embedded PHP module of the Apache server. The rest is composed of JavaScript and shell functions [2]. The source code and configuration files are located in the `/www/html/na62` web directory, except for an access library that resides in a secure directory. Data are stored partly on the filesystem and partly in a dedicated MySQL database.

A. Data storage

The backend of the NA62 Grid Interface is a MySQL database running on `hughnot.physics.gla.ac.uk` in Glasgow and maintained by PPE IT. The database structure and tables are listed in appendix VIII A.

The File Transfer Controller (FTC) makes use of another database located at IC [3]. However, after the file replication is successfully completed, all relevant information is copied to the Glasgow DB.

The scripts used for every jobs submission (wrapper, mac and JDL) are stored on a filesystem that is accessible from the web server (WSF). After a job completes, the `stdout` and `stderr` are retrieved by the system and stored on the WSF. This ensures that *full information about the job* is available in case a problem is detected at a later stage.

B. File Transfer Controller

One of the requirements of the NA62 MC production process is to ship data samples back to CERN and/or to RAL. A possible scenario would be to copy the output at the end of a job using `lcg tools` or ideally a more robust File Transfer Service (FTS). This approach however leads to unnecessary complications when handling the transfers (submitting, monitoring, registering with the LFC, resubmitting) and, more importantly, would create another possible point of failure with no easy recovery.

Such considerations led to the implementation of a custom-designed File Transfer Controller (FTC) that allows us to decouple file transfers from the actual MC simulation task. The FTC is technically speaking a secure FTS client, implemented as a secure Axis Web Service deployed on Tomcat. File transfer requests are internally stored in a MySQL database. A simple cron job manages transfer submissions, status checks and possible resubmissions. A dedicated FTC robot certificate is used for the transfers. The FTC is called from within the Grid job after the MC output is validated. The job simply registers the file transfer request the rest is done asynchronously by the FTC. On success the file is registered with the LFC [4].

C. Authentication

Parts of the website are open to general (read) access. Job management functionality (submission, re-submission, deletion) is password-protected via a dedicated user id and password combination. CERN passwords or Grid certificates are **not** required for using the NA62 Grid Interface. Each user [5] has his/her own password for authentication via the NA62 Grid Interface, which he/she sets during registration. Passwords set via the registration page are encrypted.

D. Monitoring jobs

The most important view is the Jobs table, shown in Fig. 1. Some common characteristics of all views present here are:

1. Explanations/hints above the table
2. Tasks/Options bar at the top of the table
3. Stats/Navigation bar at the bottom of the table

The Jobs and Files tables have in addition an 'Actions' drop down menu below the table and a 'Search' option in the Tasks menu.

1. The explanations/hints above the table are relevant and should be read, they explain the content and/or layout of the table below.
2. The Tasks menu provides switching between the Jobs, Files, Production and Job submission ('Scripter')



NA62 MC Jobs List

This list is dynamically created from the information stored in the GridJobs database. Click the 'Date/Time' cell for job details. Status is automatically updated every 15 minutes. By default, only the last 25 entries are shown. Use the search function for older entries.

Search | Jobs | Files | Production | Misc | Scripter | Refresh | Show all | Expand

[x]	Run number	Decay type	Number of events	MC s/w version	Submission Date/Time	Submitted by	Site	Files					Job status	Status URL	MC Output
								JDL	Mac	Exe	Out	Err			
<input type="checkbox"/>	23753	140	6000	v9/r261	16 May 22:10:02	protopop	UCL	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23752	140	6000	v9/r261	16 May 22:09:54	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-
<input type="checkbox"/>	23750	140	6000	v9/r261	16 May 22:09:44	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23749	140	6000	v9/r261	16 May 22:09:36	protopop	UCL	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23748	140	6000	v9/r261	16 May 22:09:28	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-
<input type="checkbox"/>	23747	140	6000	v9/r261	16 May 22:09:20	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-
<input type="checkbox"/>	23746	140	6000	v9/r261	16 May 22:09:12	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-
<input type="checkbox"/>	23745	140	6000	v9/r261	16 May 22:09:03	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-
<input type="checkbox"/>	23744	140	6000	v9/r261	16 May 22:08:55	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23743	140	6000	v9/r261	16 May 22:08:47	protopop	UCL	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23742	140	6000	v9/r261	16 May 22:08:39	protopop	RL	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23741	140	6000	v9/r261	16 May 22:08:31	protopop	RL	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23740	140	6000	v9/r261	16 May 22:08:23	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23739	140	6000	v9/r261	16 May 22:08:14	protopop	UCL	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23738	140	6000	v9/r261	16 May 22:08:06	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23737	140	6000	v9/r261	16 May 22:07:58	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-
<input type="checkbox"/>	23736	140	6000	v9/r261	16 May 22:07:50	protopop	UCL	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23735	140	6000	v9/r261	16 May 22:07:42	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23734	140	6000	v9/r261	16 May 22:07:35	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-
<input type="checkbox"/>	23733	140	6000	v9/r261	16 May 22:07:26	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-
<input type="checkbox"/>	23732	140	6000	v9/r261	16 May 22:07:18	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-
<input type="checkbox"/>	23731	140	6000	v9/r261	16 May 22:07:10	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-
<input type="checkbox"/>	23730	140	6000	v9/r261	16 May 21:21:06	protopop	UCL	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23729	140	6000	v9/r261	16 May 21:20:58	protopop	UCL	JDL	Mac	Exe	Out	Err	CLEARED	LINK	CC
<input type="checkbox"/>	23728	140	6000	v9/r261	16 May 21:20:50	protopop	GLA	JDL	Mac	Exe	Out	Err	CLEARED	LINK	-

Total: 25 runs, 150000 events. Average MC runtime*: 4 h 47 min. Average time per event*: 3 s.

< 1 2 3 4 5 6 7 8 9 ... >

* these numbers become meaningful once certain selection rules are applied. Legend: CC - CERN Castor, LS - local storage

With selected entries: Choose action

- Resubmit
- Kill
- Get status

Useful commands:

glite-wms-job-status JOB-STATUS-URL - to retrieve the status of a particular job
 glite-wms-job-output JOB-STATUS-URL - to retrieve the output of the job
 glite-wms-job-cancel JOB-STATUS-URL - to kill a specific job

FIG. 1: Screen shot of the Jobs List (full view). The page features a short explanation text at the top, the main navigation/menu bar, the runs table, a stats/navigation bar at the bottom, and an action menu and useful hints below the table. The green boxes in columns 9-13 are links to the relevant run files. The "LINK" box provides quick access to the jobs status (provided you have the correct Grid certificate uploaded in your browser).

views. The Options menu provides display options for the current table (see Fig.2):

- Refresh - manual refresh (otherwise auto refresh is done every 15 minutes)
- Show all - reset all filters (see below)
- Expand - show expanded info for all rows

3. The Stats bar displays rough statistics based on the displayed rows: e.g. total number of items, average MC run time, average MC time per event, total output size etc. On the right hand side of the bottom bar minimal navigation links are provided.

By default, 25 rows are shown, and simple page-by-page navigation links are provided in the bottom right corner of the table. The number of rows displayed can be modified by performing a search (see Fig.4) or by appending to the page URL the option `&entries=N`, where N is the desired number of rows to be displayed.

The first column in the table contains checkboxes. This enables the user to select a subset of runs and apply an Action to them. Available actions here are: 'Resubmit', 'Kill' and 'Get status'. Once the 'Go' button is clicked, the interface displays in plain text a set of commands that the user can copy and paste into a Grid UI terminal for execution [6].

The columns with brown text enable filtering by channel ID, number of events, MC software version, production user, destination site, job status, or Castor output status. The filtering conditions are combined with AND. For example (see Fig.1) clicking '140' on column 3, will select only jobs having this channel ID. If then the user clicks 'RL', only jobs with channel ID equal to 140 and destination RAL will be displayed, etc. Click 'Show all' to reset all filters.

Extended job information can be viewed by clicking the date column, or via the 'Expand' option from the Table menu (this will expand all rows). Once accessed, the detail box will show the job description, reaction channel info, quote of the Grid status command, the destination CE, and, for successfully completed jobs, the MC runtime.

An important feature of the NA62 Grid Interface is that for each and every job submitted the relevant files are archived for later reference and direct links are given. Linked in columns 9-13 of the jobs table are the JDL of the Grid job, the `.mac` configuration file for the MC, the wrapper script ('`exe`') and, after job completion, the `stdout` and `stderr` for that particular job. The job and run settings and parameters stored in the database are thus duplicated for redundancy. And making all these files easily available for every NA62 job run greatly simplifies tracing and debugging for all software and middleware components, during production, and later on during the analysis stage.

The last column in the Jobs table provides a quick status of the MC output. If it is empty or "-", the job has produced no output (something went wrong), LS means the output is on local storage on the site where the job

ran, and CC means there is a replica on CERN Castor. A successful job is a job with CC status. **How are those replicas managed:** The aim is to have two replicas of the MC output, of which one on CERN Castor for analysis. The wrapper script runs the MC and then makes two calls to the FTC to deposit two copies of the `.root` output file: one at RAL and one at CERN. The FTC handles these calls asynchronously; the final status is checked periodically and the DB updated accordingly. It might take up to 12 hours after the completion of a run for the CC replica to be in place.

E. Monitoring files

Monte Carlo output files are monitored via the 'Files' view. This view is similar in many respects to the jobs view and is accessible via a link in the Tasks menu. It provides a list of the MC output files with: a link back to the corresponding run (job), full file name, channel ID and number of events, file size, GUID and Castor replica status. Like in the case of the Jobs list, columns 4, 5 and 8 provide simple filtering. Click the GUID (or 'Expand' from the Table menu) to access extended file information like: full run description, file size in bytes and Grid UI commands to list replicas or download the file to your machine. Files that have a CERN Castor replica will be tick-marked.

By default, the table shows 25 rows, and simple page-by-page navigation links are provided in the bottom right corner of the table. The number of rows displayed can be modified by performing a search (see Fig.4) or by appending to the page URL the option `&entries=N`, where N is the desired number of rows to be displayed.

The first column in the Files table contains checkboxes. The user can select a subset of files and choose an Action from the drop down menu below the table. Available actions here are: 'Get/Download', 'Replicate on CERN Castor', 'List' and 'Delete'. Once the 'Go' button is clicked, the interface displays in plain text a set of commands that the user can copy and paste into a Grid UI terminal for execution [7].

F. Production list

Once the NA62 Physics Group has decided a production plan (typically for the next 6-12 months), this is entered into the database and can be browsed via the 'Production' view (see Fig.5).

The production table shows for each round: its tag (unique for each round), the full description, numbers of runs and files, and the total size of the output. Expanded information can be viewed by clicking the date column, or via the 'Expand' option from the Table menu (this will expand all rows). Once accessed, the detail box will show: round progress (in percents), destination directory on CERN Castor (for the MC outputs), start and

Jobs Table

DB search capabilities →

tasks menu

table menu

Search Jobs Files Production Scripter Refresh Show all Expand															
[x]	Run number	Decay type	Number of events	MC S/w Version	Submission Date/Time	Submitted by	Site	Files					Job Status	Status URL	MC Output
								JDL	Mac	Exe	Out	Err			
<input type="checkbox"/>	338	60	1000	5/r170	17 Aug 09:44:40	protopop	GLA	↗	↗	↗	↗	↗	CLEARED	LINK ▶	-
<input type="checkbox"/>	337	60	1000	5/r170	17 Aug 09:44:35	protopop	IC	↗	↗	↗			RUNNING	LINK ▶	-
<input type="checkbox"/>	336	60	1000	5/r170	17 Aug 09:44:11	protopop	GLA	↗	↗	↗	↗	↗	CLEARED	LINK ▶	-
<input type="checkbox"/>	335	60	1000	5/r170	17 Aug 09:42:58	protopop	RL	↗	↗	↗			RUNNING	LINK ▶	-
Description: Test jobs Decay channel: 60 (radiative corrections applied), random_seed=331 Status command: glite-wms-job-status https://wmslb01.grid.hep.ph.ic.ac.uk:9000/gJtU7V2lc5kDU8cPephVSW CE: lcgce05.gridpp.rl.ac.uk:8443/cream-pbs-grid2000M															
<input checked="" type="checkbox"/>	334	60	1000	5/r170	17 Aug 09:42:30	protopop	RL	↗	↗	↗			CANCELLED	LINK ▶	-
<input type="checkbox"/>	333	60	1000	5/r170	17 Aug 09:42:25	protopop	GLA	↗	↗	↗	↗	↗	CLEARED	LINK ▶	-
<input type="checkbox"/>	332	60	1000	5/r170	17 Aug 09:42:21	protopop	RL	↗	↗	↗			RUNNING	LINK ▶	-
<input type="checkbox"/>	331	60	1000	5/r170	17 Aug 09:41:50	protopop	LIV	↗	↗	↗			RUNNING	LINK ▶	-
<input type="checkbox"/>	330	43	100	5/r170	15 Aug 13:11:56	protopop	LIV	↓	↓	↗	↗	↗	CLEARED	LINK ▶	OK
<input type="checkbox"/>	329	43	100	5/r170	15 Aug 13:11:25	protopop	RL	↓	↓	↗	↗	↗	FAILED	LINK ▶	-
<input type="checkbox"/>	328	43	100	5/r170	15 Aug 13:10:46	protopop	GLA	↓	↓	↗	↗	↗	CLEARED	LINK ▶	-

filters

links to the corresponding text files

← expanded info

action select →

FIG. 2: Annotated screenshot of the jobs table. Detailed functionality is explained in the text.

Files Table

table menu

Search Jobs Files Production Misc Scripter Refresh Show all Expand							
[x]	Run number	File name	Decay type	Number of events	File size	GUID	CC*
<input type="checkbox"/>	23704	pluto_v9_r23704.root	140	6000	586 MB	2130c752-8040-4610-b47a-2e620763afd8	✓
<input type="checkbox"/>	23703	pluto_v9_r23703.root	140	6000	578 MB	9421a13e-292a-49e4-b648-b7ff996b9777	✓
<input type="checkbox"/>	23702	pluto_v9_r23702.root	140	6000	577 MB	8a13eb07-0c15-4cfa-a1f2-45176ce69a18	✓
<input type="checkbox"/>	23700	pluto_v9_r23700.root	140	6000	577 MB	65c5e212-73f7-4512-891c-93b7d9b02da9	✓
<input type="checkbox"/>	23696	pluto_v9_r23696.root	140	6000	583 MB	fe0744eb-0f8a-4699-91a3-c92af623393b	✓
<input type="checkbox"/>	23695	pluto_v9_r23695.root	140	6000	578 MB	3704aafa-1aa8-4f0d-83fa-47b97795f5c2	✓
<input type="checkbox"/>	23694	pluto_v9_r23694.root	140	6000	585 MB	675c8be8-7743-41a2-9b31-fce39380a504	✓
<input type="checkbox"/>	23693	pluto_v9_r23693.root	140	6000	588 MB	f6a99f46-45b0-4910-a639-08dc4532064c	✓
<input checked="" type="checkbox"/>	23691	pluto_v9_r23691.root	140	6000	580 MB	1dbb6e82-3754-40fc-96eb-48a3ac4c8552	✓
Description: Kch2emnumumu-1 production job File size: 607483156 B List replicas: lcg-lr lfn:/grid/na62.vo.gridpp.ac.uk/prod/r261/Kch2emnumumu/pluto_v9_r23691.root Get command: lcg-cp -v lfn:/grid/na62.vo.gridpp.ac.uk/prod/r261/Kch2emnumumu/pluto_v9_r23691.root pluto_v9_r23691.root							
<input type="checkbox"/>	23690	pluto_v9_r23690.root	140	6000	583 MB	eed160c1-3353-40ae-8d28-b078ddac3425	✓
<input type="checkbox"/>	23687	pluto_v9_r23687.root	140	6000	588 MB	5b9ca986-d209-4dde-9950-6bb537bd7d03	✓

← expanded info

action select →

FIG. 3: Annotated screenshot of the files table. Content is self-explanatory. Detailed functionality is explained in the text.

end date of the round (the end date will be fixed once the round is finished), run numbers interval, cumulated run-time, and the job success rate. Ongoing runs are shown in expanded view by default and a two-colour progress bar is displayed: light green - jobs, dark green - Castor replicas.

The 'Runs' and 'Files' columns provide direct links to the corresponding lists of runs and files for the chosen production round. The filtering is done via the round tag.

To add production rounds to the database, click 'Production' in the Tasks menu: a form will be displayed at the top of the table. One has to enter the channel ID, a round number, a total events target and admin credentials to write to the DB. The round tag is created from the channel name (a standard denomination is in place) and the round id. Make sure this combination is unique, else your DB entry will be rejected. Once a new production round is inserted in the DB, it will show up in the corresponding drop down menu from the production mode job submission interface, and submissions can be started.

G. Job submission

Jobs submission is the most important functionality provided by the NA62 Grid Interface. Multiple interfaces and mechanisms are implemented for the actual job submission:

1. production mode
2. benchmarking mode
3. cloning mode
4. Ganga round completion
5. autopilot

1. The *production mode* submission interface, shown in Fig.7, combines flexibility with ease of use. It is accessed simply via the 'Scripter' link in the Tasks menu, and should be used for regular MC production. The form is pre-filled based on the current production, and to continue a production round, the Grid shift taker should only change the run interval and click 'Prepare'. If a new production round should be started, then the 'description' (channel ID), number of events should be changed etc. **How is the number of events chosen:** The ideal run time for a job is between 4 and 8 hours. This will depend on the reaction channel simulated, so a few test jobs should be submitted via the benchmarking interface (2) to obtain an average time per event that can be used to calculate an optimum number of events per run for a given production run.

2. A link to the *benchmarking mode* interface is provided in the header text of interface (1). This interface offers more options, and should be used for tests. **All test submissions should contain the keyword 'test' in the job description**, such that they are not counted in the overall production round statistics.

Both (1) and (2) submission modes offer the option to choose the destination sites. This allows the temporary exclusion of a Grid site that has technical problem, or submission of test jobs to a target site.

3. The benchmarking interface can be run on *Cloning mode*. The difference here is that while submission modes (1) and (2) will archive JDL and .mac pairs for all runs to be submitted, in cloning mode only one copy of this pair is archived and linked in the table. Cloning mode should be used only for long, smooth running production rounds.

Cloning mode is also available via the mobile interface, which will be described in the next section.

4. The *Ganga round completion* option is provided directly from the 'Production table'. All scheduled but unfinished production rounds will show a 'Run via Ganga' link. Once started, Ganga will run jobs until the corresponding production round is completed. Ganga submission mode is only available to advanced users (designated as such in the `shifters` DB table).

5. *Autopilot production* can only be activated by the NA62 Grid Interface admin. This can be done for long, smooth running production periods that require minimal supervision. When autopilot production is enabled

All the above production modes use the same set of scripts and PHP functions. Hence, the various submission interfaces act as plugins for the NA62 Grid Interface. Besides being easier to maintain and customise, this setup provides excellent cohesiveness and robustness to the system.

H. Smartphone App

There is also an iPhone webapp for monitoring NA62 Grid production. The app provides a pie chart showing the distribution of running, scheduled, cleared and 'on Castor' jobs (over the last 2 days), and alternate views of simplified run, file and production tables.

The iPhone app also provides quick links to NA62 MC documentation and the NA62 Grid Interface developer email, plus a simple job cloning interface only available to advanced users. An Android webapp might be made available if there is demand.

I. Miscellany

There is a 'Misc' section of the NA62 Grid Interface, linked in the Tasks menu, providing a short description of the experiment and:

- links to Howto wikis
- a link to the Grid Shifts logbook/blog
- a site performance table

The Howto wikis are very important sources of information on the components of the NA62 Grid production

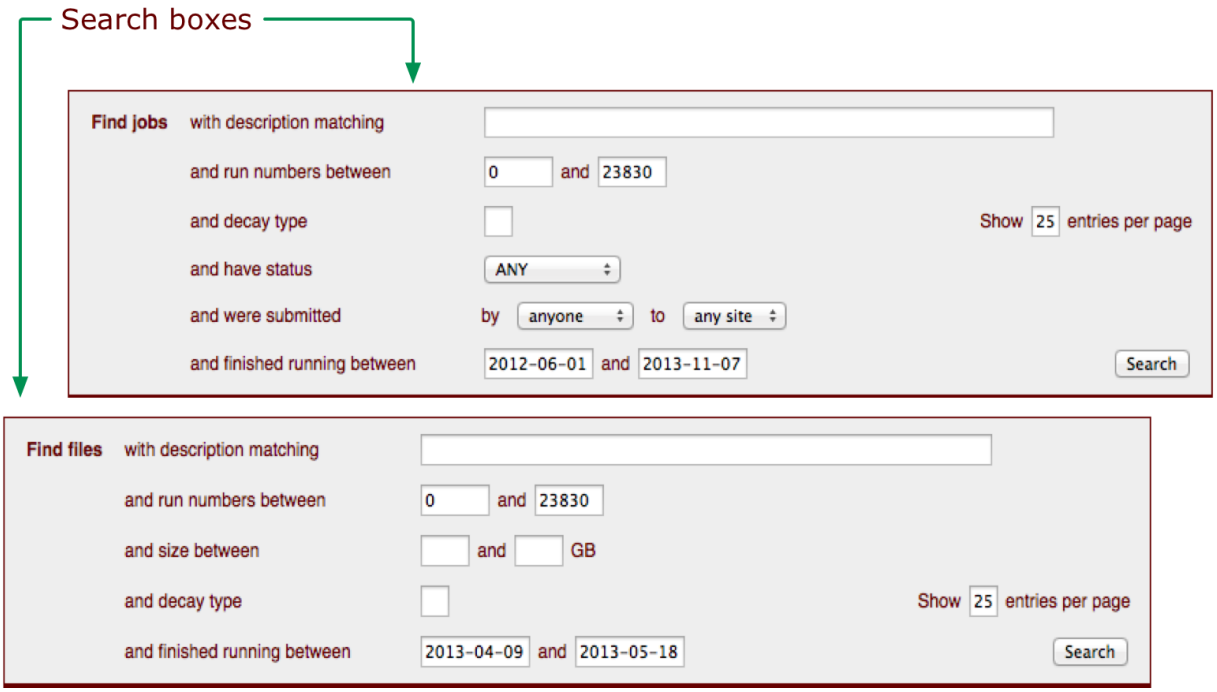


FIG. 4: Cropped screenshots of the run and file search boxes. Options shown are self-explanatory.

Production Table

Jobs | Files | **Production** | Misc | Scripter | Refresh | Show all | Expand

New production round Channel ID (0-134): Round number (1-9): Events target (e.g. 10k, 5M): ← add rounds

Enter user and password to confirm Add

Tag	Description	Runs	Files	Total size
Kch2emnumumu-1	Production round, 500k events decay type 140, $K^+ \rightarrow e^- u \mu^+ \mu^+$ (Kch2emnumumu)	140	93	0.06 TB
Kch2mumnu-1	Production round, 500k events decay type 141, $K^+ \rightarrow \mu^- u e^+ e^+$ (Kch2mumnu-1)	138	115	0.09 TB
Kch2munumumu-2	Production round, 500k events decay type 103, $K^+ \rightarrow \mu^+ u \mu^+ \mu^-$ (Kch2munumumu)	119	93	0.03 TB
<div style="display: flex; align-items: center;"> <div style="width: 100px; text-align: right;">Progress bar →</div> <div style="width: 600px; height: 15px; background-color: #008000; margin-left: 10px;"></div> <div style="margin-left: 10px;">← expanded info</div> </div> <p> Progress: 109% Destination directory: /castor/cern.ch/grid/na62/mc/prod/r261/Kch2munumumu Date interval: 16 May 10:35 - 17 May 07:11 Run numbers: 23421 - 23580 Cumulated runtime: 112 hours Job success rate*: 78% </p>				
Kch2enumumu-2	Production round, 500k events decay type 102, $K^+ \rightarrow e^+ u \mu^+ \mu^-$ (Kch2enumumu)	119	83	0.06 TB
Kch2munuee-2	Production round, 500k events decay type 101, $K^+ \rightarrow \mu^+ u e^+ e^-$ (Kch2munuee)	167	89	0.09 TB

← stats

FIG. 5: Annotated screenshot of the production table. Content is self-explanatory. Detailed functionality is described in the text.

system and MC software, on how the NA62 Grid Interface should be used, how Grid production should be carried out, etc. **The following two wikis should be read by all Grid shifts volunteers:**

1. Production Howto - explains in detail job submission, monitoring, production management, troubleshooting
2. Production Rota - gives the shift volunteers list, a shift checklist - what to do on shift, a link to the production plan etc.

The Grid Shifts Logbook is practically a Wordpress blog located at

<http://na62shifts.wordpress.com>

Shift volunteers should log in there everything that they do, round start and end dates, daily progress, problems encountered, tips or requests for other shift takers etc. Shift takers should obtain the user id and password for the logbook from the current production coordinator or the NA62 Grid Interface admin. The logbook should be a comprehensive and reliable source of information for everyone involved in NA62 Grid production and beyond.

The site performance table provides a quick summary of Grid site performance and could a) help shift takers identify problems with a given site or queue and b) give site admins a measure of how well their site serves the NA62 VO.



FIG. 6: Map of participating sites.

V. GRID SHIFTS

Once a production plan is in place, most of simulation jobs on the Grid are run automatically, via cron jobs that collect the requirements from the database. However, there are periods (testing, benchmarking) when Grid jobs have to be run (semi)manually, and volunteers for Grid Shifts are needed.

To volunteer for Grid Shifts one has to register via <http://na62.gla.ac.uk/register.php>, provide an email address and choose a username and password. Registrants are encouraged to safeguard this password; if lost, they might be asked to repeat the registration process. Bogus registration requests will result in banning your IP address.

Shift assignments are done simply via an ad-hoc Doodle poll [8], announced via email.

Weekly email reminders may be employed during Grid shifts periods to keep production on track and inform participants of the progress. The reminders are triggered via cron jobs running on one of the Glasgow machines (see Table II).

VI. MAINTENANCE AND SUPPORT

The main system components, including the database backend and Apache server, run on physics.gla.ac.uk servers which are maintained by PPE IT admins and/or the Glasgow GridPP team.

The FTC components are located at IC and maintained by the IC group; however, these components are only necessary at the file replication stage, while all essential run parameters are archived in Glasgow, on a filesystem that is itself backed up on tape every night.

A. Participating sites

At the time of writing, there are five participating sites from the UK and one from Belgium, totalling 350-700 worker nodes. We make use of 100 TB of disk+tape storage space at RAL and have access to extensive GridPP expertise.

B. Documentation

The NA62 Grid Interface is designed to be very easy to navigate and intuitive to use. However, ample documentation is provided via: a) hints and/or relevant explanations on every NA62 Grid Interface webpage b) this technical note (which will be updated if any functionality is added or modified) and c) detailed wikis linked from the NA62 Grid Interface.

Job parameters:

Choose description: production job ← defaults from production plan

Run number(s): to

Number of events (per run):

Random seed: ← for multiple runs this is set automatically

MC version:

Radiative corrections: on off ← restricted set of options

Disable detector(s):

Disable Cedar Cherenkov Effect:

Script parameters:

Destination (any of): GLA IC LIV RAL BIR UCL CNAF (check)

Display commented lines from scripts:

Write scripts to disk: + enter user: and password: to confirm

Have you checked [today's production plan](#) and [the logbook](#) before submitting new jobs? ← quick links to logbook and wiki

FIG. 7: Annotated screenshot of the production-mode job submission interface.



FIG. 8: Annotated screenshots of the NA62 iPhone WebApp for Grid production monitoring. To install this app on your smartphone, navigate to <http://na62.gla.ac.uk/app/> and use the 'Add to Home Screen' option.

C. Feedback and bug reports

We welcome all feedback, suggestions and bug reports regarding the NA62 Grid Interface, the related Wiki, shifts logbook (blog) and occasional Doodle polls. Please address all communication to na62.gridadmin@cern.ch. The admin(s) will reply to you as soon as possible.

VII. SUMMARY

The NA62 Grid Interface is a one-of-its-kind system, created specifically for NA62 and providing the collaboration with significant productivity gains. Excellent Grid job monitoring and management is combined with ease of use and an intuitive interface. Add to this its comprehensive archiving capabilities and its single-point access architecture for all simulation-related data and one has a system unequalled among small Grid VOs.

Acknowledgments

Lots of thanks to the Glasgow, IC and RAL GridPP support groups for providing important middleware expertise. Many thanks to Antonino Sergi (CERN) for liaising with the NA62 software development group. Special thanks to Andrew Pickford (Glasgow) for providing DB and Apache server maintenance and support, and to Mark Slater (Birmingham) for installing and customising the Ganga auto-submission plugin.

Many thanks to our Grid shift volunteers and everyone else who helped improve the NA62 Grid Interface by providing feedback, making bug reports or simply chanced upon a software bug, or providing comments or suggestions for new features.

VIII. APPENDIX

Included here are some technical details regarding the NA62 Grid Interface components and a general FAQ section.

A. Database Tables

Database tables accessed by the NA62 Grid Interface at the time of writing are listed in Table I.

B. List of scripts

Low level functionality is implemented via `bash` shell scripts. These scripts are located and run on the Glas-

Functionality	Database Table	Use
Run list	jobs	O
List of output files	files	O
Production runs list	production	O
List of software versions/revisions	versions	I
Reaction channels	channels	I
Activity monitoring	events	A
Load monitoring	history	A
Shift workers list	shifters	A

TABLE I: MySQL tables in the GridJobs database. These tables are used for displaying information online (O), by the software internally (I) for cross-referencing, or accessible to the admin only (A).

Functionality	Script name	Run via	Alters DB
Job status updates	updatedb.sh	C	✓
Run history updates	history.sh	C	✓
Check Castor replicas	verify-castor.sh	C	✓
Submission robot	action.sh	C	✓
Autopilot production	run-production.sh	C	✓
Email reminders	reminder.sh	C	
Robot proxy renewal	renew.sh	C	
Produce list of missing Castor replicas for the FTC	castor-xlist.sh	M/C	✓
Repair missing or erroneous info in the database	fix_dest.sh fix_nev.sh fix_file_sizes.sh	M M M	✓ ✓ ✓
Recover outputs from zombie jobs	recover_zombies.sh	M	✓
Delete local copies from a given SE	cleanup.sh	M	✓
Job manipulation	resubmit-run.sh split-job.sh glite-wrap.sh	S S S	↔ ↔ ↔
Used by other scripts	jobfiles.sh joboutput.sh jobstatus.sh jobtime.sh on_cern_castor.sh	S S S S S	
Nightly DB backups	backup_db.sh	C	

TABLE II: List of shell scripts used. These scripts can be run as cron jobs (C), manually (M), or are used from within other scripts (S). Scripts marked with ↔ modify the DB indirectly.

gow machines and listed in Table II.

The cron scripts run asynchronously, hence the system behaviour does not depend on the job queue size.

C. How does it all work

A production plan is drafted by the NA62 Physics Group and communicated to DP or the production coordinator. NA62 MC software packages are provided by the NA62 software group via AS [9]. The software is built by DP and installed on sites (or distributed via CVMFS) by JM, while site technical support is provided by the site administrators. During production, job submission is done by the production coordinator, or the current shift taker, or autonomously using the scripts generated by the online tools. Once a job is submitted, it is automatically inserted in the local DB. From there, job tracking and DB logging are done automatically via cron jobs. At execution, the job wrapper script verifies that all components are in place, executes the MC, and upon successful completion triggers replication of the output file(s) at RAL and CERN via the FTC. Production stats are updated automatically along the way. Validity of MC output is checked weekly by DP and AS. Periodic file cleanups and DB corrections are done manually by DP.

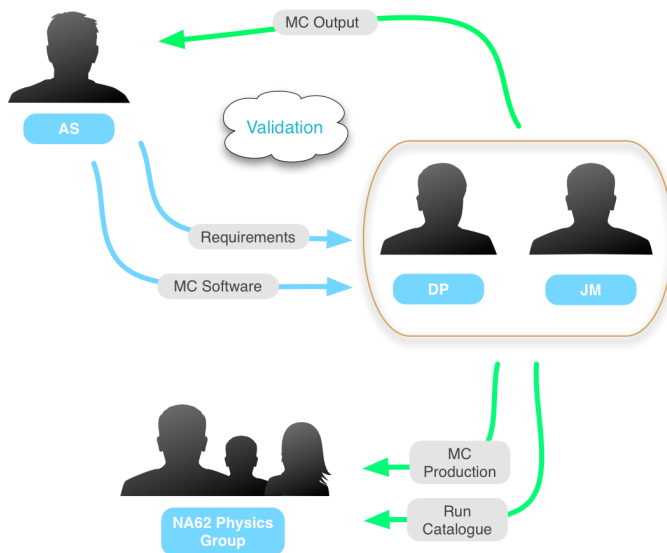


FIG. 9: Organisational diagram.

The final result is that, at the end of a production period, the NA62 Physics Group has all simulations available on local disk at CERN for data analysis.

D. Site security

Data security is accomplished via custom-written authentication functions. Login employs dedicated passwords independent of the regular CERN or Grid credentials. Grid certificates are not required from Grid shift takers or the production coordinator. A robot certificate is used for all Grid job submissions and monitoring. A separate robot certificate is used by the FTC.

All user passwords are encrypted. All login attempts (id, time stamp, originator's IP) are recorded in the `events` DB table and suspect IPs are blocked from using certain parts of the NA62 Grid Interface.

The DB server is firewalled. SQL injection via forms is preempted by replacing text input fields with drop-down menus, or where not possible by pre-parsing `POST` or `GET` data before constructing MySQL queries. Illegal characters in input fields are automatically expunged. In FORM submissions, `GET` method is used on publicly accessible pages, where one might want to directly quote a link, but the more secure `POST` method is used on pages that require login.

E. FAQ

Frequently asked questions are answered in here, with the notice that new items may be added or explanations changed if new features or procedures are implemented.

How do I login ? You do not have to log in in order to access files, jobs and production information, but you have to enter a correct user and password combination to submit jobs.

What's my password ? Each user has his/her own password for logging in to this interface. If you don't know yours, forgot it or never knew you have one, please read the next FAQ item. Passwords set via this interface are encrypted.

What to do if I don't know my password ? You have chosen a username and password when you first registered for NA62 Grid shifts. If you have forgotten your details, or if you have not registered before, please register via <http://na62.gla.ac.uk/register.php> The NA62 Grid admin will notify you once your registration is validated.

How do I change my password ? You can not change your password, unless you elect to re-register. The new password will be stored encrypted in the database.

How do I sign up for shifts ? If Grid shifts need to be filled, the NA62 Grid admin will notify you via email. Follow the instructions provided.

What is an 'advanced user' ? Advanced users are users that have high technical expertise with regard to the NA62 Grid Interface, for example the production coordinator. Advanced user can for example test beta features or edit DB settings.

Can I run my own MC simulations ? The system described in this note is designed for large-scale simulations coordinated by the NA62 Physics Group, and as such, we only run the *official* NA62MC and follow the *approved* schedule. To run your own version of MC software, you must contact the production coordinator.

What's the password for the NA62 logbook ? Shift takers can obtain the user id and password for the logbook from current production coordinator or the NA62

Grid Interface admin.

How often are the job submission requests picked up ? The action bot picks up requests every 7 minutes. The actual submission may take up to 5 minutes if the server load is high.

How do I install the smartphone app ? To install the NA62 webapp on your iPhone, for example, simply navigate to <http://na62.gla.ac.uk/app/> and use the 'Add to Home Screen' option.

[1] NA62 CERN web page: <http://na62.web.cern.ch/na62>

[2] approximately 4000 lines of PHP and 1500 lines of bash

[3] IC stands for Imperial College, London

[4] LFC stands for 'LCG File Catalogue'

[5] by NA62 Grid Interface 'user' we mean production coordinator or Grid jobs shift taker

[6] this requires having the appropriate Grid certificate

[7] some actions require having the appropriate Grid certificates, and may be available to advanced users only

[8] see doodle.com

[9] AS=Antonino Sergi (CERN), JM=Janusz Martyniak (IC), DP=Dan Protopopescu (Glasgow).