Gaia on the Ground

M. Davidson, N. Hambly & N. Rowell, Wide Field Astronomy Unit, Institute for Astronomy, University of Edinburgh

Introduction

Much work is required to make sense of the observations sent by Gaia, so that the string of binary 1s and 0s eventually leads to a catalogue of stars. Teams of hundreds of scientists, engineers and programmers work together on each stage of the process.

The first step is to **download the data** from Gaia to Earth using three large antennae in Australia, Spain and Argentina. From here the information is assembled at the Mission Operations Centre in Darmstadt, Germany before being sent to the Science Operations Centre near Madrid, Spain for a first look.

Later, the observations are **distributed to processing centres** across Europe, including Cambridge and Barcelona, for more intensive analysis. As the data grow and become better understood catalogues are prepared and released to the community of astronomers.

Daily Processing - Health checks and alerts

New observations from Gaia arrive every few hours and are eagerly awaited! These data are processed as soon as possible for two main reasons. Firstly, it is important to check that the spacecraft remains **healthy** and is performing well. Secondly, Gaia can detect changes in known objects to find events such as supernovae. Gaia can then **alert** other teams to train their telescopes on the star to watch the action unfold, almost in real time.

We contribute algorithms and software to this real-time system, which runs in Spain, to help tune the models used to extract the key position and brightness measurements of each star. Automated checks have been developed to show the spacecraft and processing status at a glance.

Cyclic Processing - Ever better

The daily processing allows us a quick look at the data but it is limited by the hardware and time available. Once a year the entire set of observations is reprocessed at several locations (for instance, the Mare Nostrum supercomputer in Barcelona). Such **powerful computers** allow us to perform more measurements and calibrations with the latest and greatest techniques. We can even use information from previous processing cycles to further **refine the results**. After each cycle our measurements should be ever closer to their true values.

Our software and methods are used in processing centres in Barcelona, Madrid, Cambridge and Paris. Our team has also worked on the portal through which the Gaia catalogue is made available to the world.

Figure 1: European Space Agency tracking station at Malargüe, Argentina. Credit: ESA

Figure 2: The Mare Nostrum computer where Gaia observations are processed. It is the most powerful machine in Spain and is housed in a former chapel. Credit: Barcelona Supercomputing Center.

The technical bit ...

The Gaia processing faces some challenges which may be familiar to anybody who has worked on a large software project, and some which may not!

- Gaia transmits with the power of three 100W bulbs from 1 million miles away at a data rate up to 10 Mbps.
- At the end of the mission the raw telemetry volume is estimated to be 200 terabytes and the complete data products around 1.68 petabytes.
- Once the telemetry is unpacked the processing almost exclusively uses Java, with data stored and transferred as serialised objects.
- The Gaia Data Processing and Analysis Consortium (DPAC) has over 400 members in 24 countries.
- DPAC uses Subversion for software version control, Jira for issue tracking, and Jenkins/Hudson for continuous integration.
- The Mare Nostrum supercomputer offers a peak performance of 1.1 petaflops, 100.8 terabytes of main memory and 3056 processing nodes.

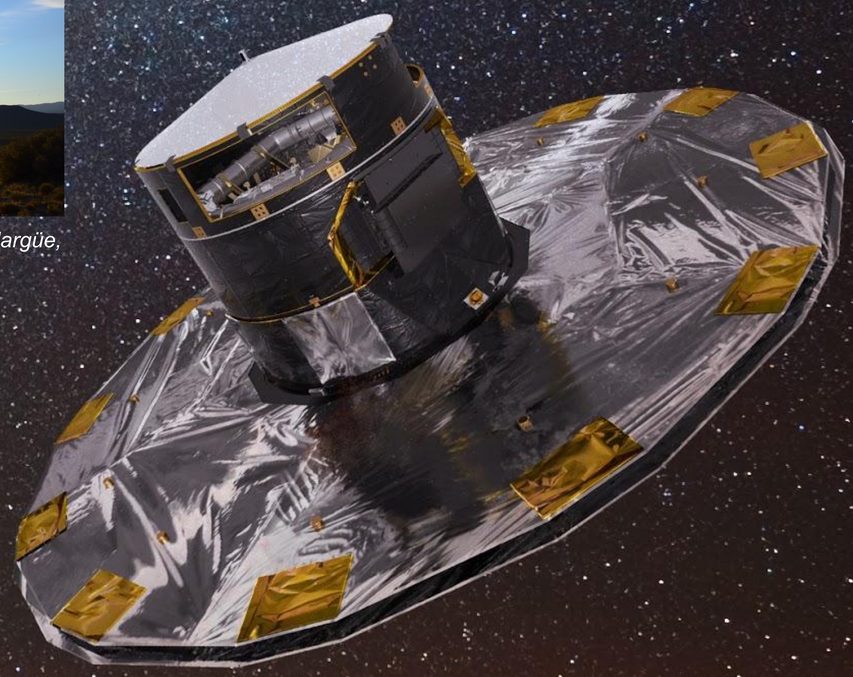

Gaia Artist's impression - credits: ESA/ATG medialab; background image: ESO/S. Brunier