Gaia Science Alerts
White Book

Editors: Simon Hodgkin, Lukasz Wyrzykowski

Date: 21 June, 2010
Revision 1: 22 October, 2010

Contents

1 This Document

This document has been revised following the Gaia Science Alerts Workshop, held at the Institute of Astronomy, Cambridge, June 23-25 2010.

The aim is to provide a repository of knowledge about Gaia to a wider community: to summarise the mission from a transient science point of view.

2 Motivation

Gaia will repeatedly monitor the whole sky in the optical. The Gaia Science Alerts stream therefore represents a unique opportunity, and holds a significant responsibility.Triggers from transient phenomena are the first data that the astronomical community will see from the satellite. It is up the Gaia Science Alerts Team to make sure that the alert stream is accurate, reliable, interesting and free from (or at least acceptably low in) contamination. On the other hand, it is vital to make sure that the astronomical community is ready for these Gaia alerts, and that Gaia starts delivering science quickly. In addition, we will raise public awareness about (and involvement with) the satellite very early on in the mission.

With approximately two years to go until the launch of Gaia we organised a workshop (with support from the GREAT-ESF) at the Institute of Astronomy with two major goals:

1. To focus community attention on the scientific possibilities that will arise from the Gaia Science Alerts data stream, and to make sure that astronomers are prepared and motivated to exploit the data as it arrives.

2. To invite the community to influence the scope of the science alerts processing algorithms and alert strategies, to ensure that returns from the mission are maximized, and that exciting opportunities are not overlooked.

3 Gaia: Key Features

3.1 The Scanning Law

Gaia will perform it's observations from the L2 Lagrange point of the Sun-Earth system. It will operate for 5 years, spinning constantly at 60 arcseconds/sec. The two astrometric fields of view will scan across all objects located along a great circle perpendicular to the spin axis. Both fields of view will be registered on one focal plane, but (almost) the same part of the sky will transit the second FOV 106.5 minutes after the first. The spin axis precesses slowly on the sky resulting in multiple observations of the whole sky over the lifetime of the spacecraft.

For a spin rate of 60 arcseconds/sec and a solar aspect angle of 45 degrees, the precession speed is such that 5 years of operation corresponds to 29 revolutions of the spin axis around the solar direction; the precessional period thus equals 63 days. On average, each object on the sky is observed about 80 times (two astrometric fields combined and 20% total dead time assumed). The Scanning Law is illustrated in Figure 1 and the resulting coverage (in Galactic Coordinates) in Figure 2. Note that the ecliptic plane is under-observed (with around 70 observations over 5 years) and ecliptic latitudes ± 45 degrees are over-observed (with up to 200 observations).

3.2 Focal Plane

The Focal Plane Assembly (FPA) is shared by both telescopes and comprises five distinct systems, shown in Fig. 3:

1. The Wavefront Sensor and basic angle monitor.
2. The Sky Mappers (SM) are used to identify which telescope is viewing the object, and to alocate a readout window in subsequent CCDs for the source.
3. The Astrometric Field (AF) is an array of 9 (along scan) by 7 CCDs and are read out in time-delayed integration mode sychronised to the scanning motion of the satellite.
4. The Blue and Red Photometer (BP/RP) CCDs are fed by two low dispersion prisms and cover 330-680nm (blue) and 650-1050nm (red).
5. The Radial Velocity Spectrograph (RVS) measures spectra of all objects brighter than about 17th magnitude. Note, only upper part of the focal plane is covered by the RVS CCDs hence not all objects from a given FoV will have their spectra taken.

The average density of stars on the sky to V=20 is around 25000 per square degree, but with a large concentration near the Galactic plane. Most of the time, this translates to about 350 stars per CCD (or 23000 stars in the entire Astrometric Field).

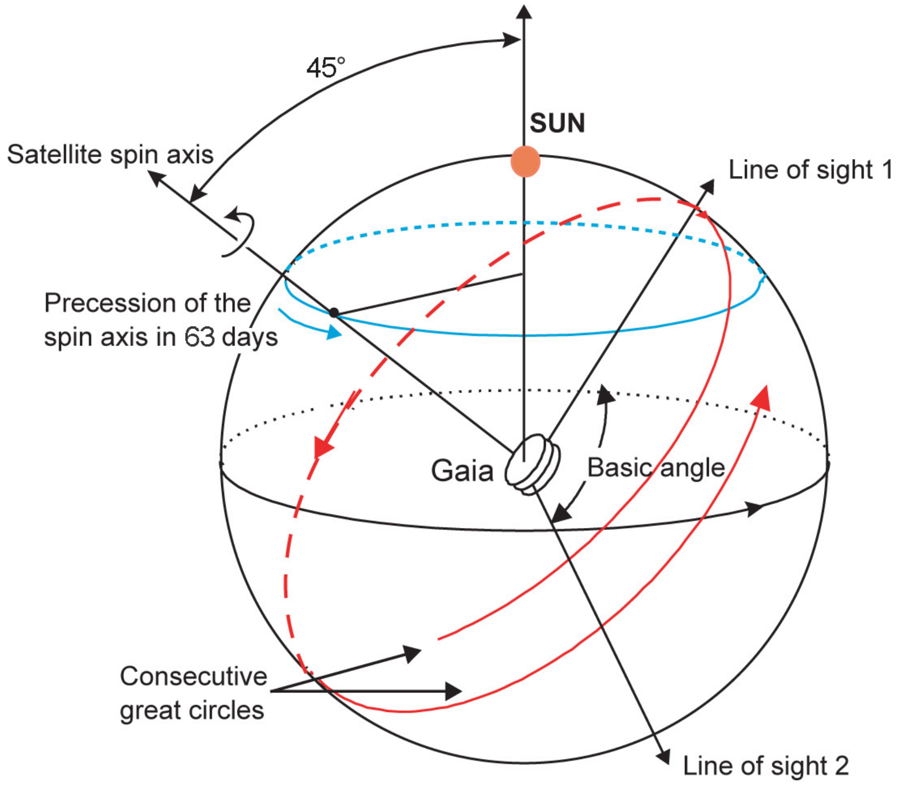

Figure 1: Gaia's two astrometric fields of view scan the sky according to a carefully prescribed 'revolving scanning law'. The constant spin rate of 60 arcsec s^{-1} corresponds to 6-hour great-circle scans. The angle between the slowly precessing spin axis and the Sun is maintained at 45°. The basic angle is 106.5°. Figure courtesy of ESA - J. de Bruijne.

3.3 Windows

In order to save on data volume being transmitted to the Earth the whole AF array is not read out. Based on the position of a star on the SM CCDs and knowing the spin rate, a star is assigned a readout window.

The time to cross a single CCD for a source is 4.4167 seconds, and in principle the photometry on those timescales from each individual CCD will be available. At the very least, these data will be used to check for source behaviour and consistency across the transit, allowing for removal of cosmic rays, for instance. The time to cross the entire field-of-view of a single telescope is therefore about 40 seconds, and the photometry on transit durations will be combined into a single measurement. Resources permitting, we will store photometry at the higher sampling rate.

3.4 Data Volume and Transmission Timescales

The data volume captured each day will be between 150 and 800 Gbyte/day. There is an onboard 850 Gbyte solid-state mass memory.

Data are downlinked to a single antenna at Cerebros in Spain every day for about 8 hours (when L2 is visible from the ground station). A second antenna in Australia will be used during scans of the galactic plane to assure the large data volume produced there will get downlinked.

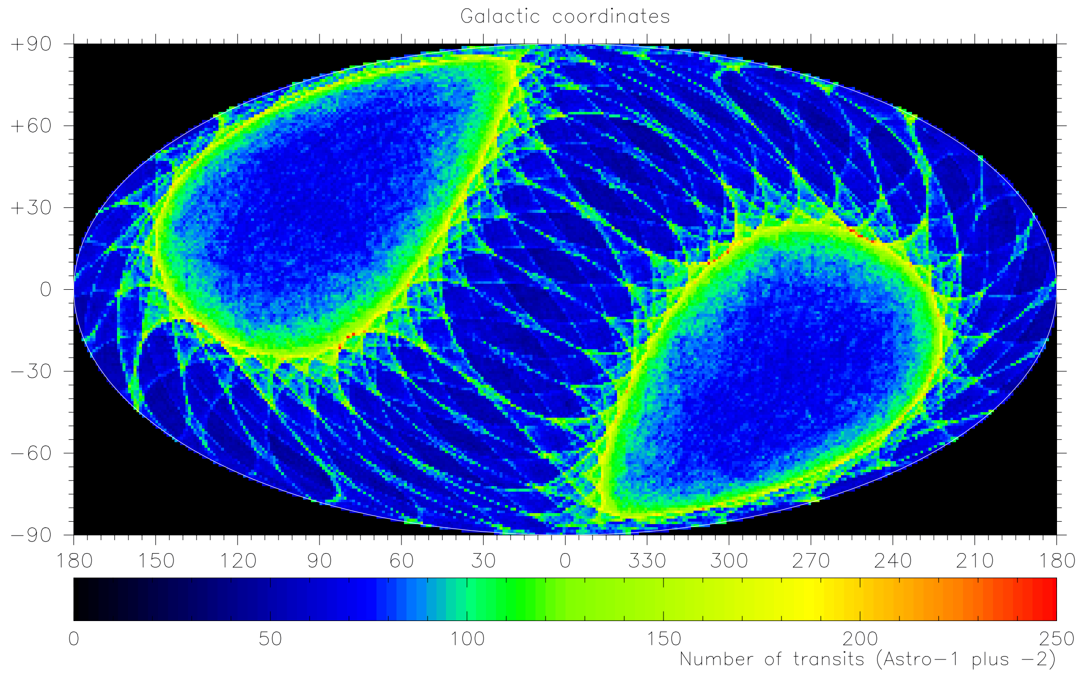

Figure 2: Predicted astrometric transits during the Gaia mission in Galactic coordinates (Hammer projection). Figure courtesy of ESA - J. de Bruijne.

Data are then sent quickly to the Science Operations Centre in Madrid for immediate data processing. Initial stages are: data organisation, cross-matching of transit measurements against the Gaia Source list, and application of a first astrometric solution. The data will then be copied to Cambridge for photometric and Science Alerts processing. The timescales are illustrated in Figure 4. Note that the Initial Data Treatment described here does not wait for transmission to complete before it starts, thus data will be received in Cambridge with a range of times since observation, anywhere between couple of hours and 24 hours (in the worst case). (This is not true for the astrometric solution which is a daily operation - thus astrometry follows photometry with a variable Δt.)

This is a rather simplistic view - and transmission gets more complex for observations when the source density is consistently high - for example great circle scans which spend a large fraction of time in the Galactic plane, or Baade's window. When there is more data than available bandwidth and on board storage can handle, on-board software will handle the selection of objects for downlink. The Gaia Science Team currently have the issue of how to select objects on their agenda (e.g. magnitude limited versus random).

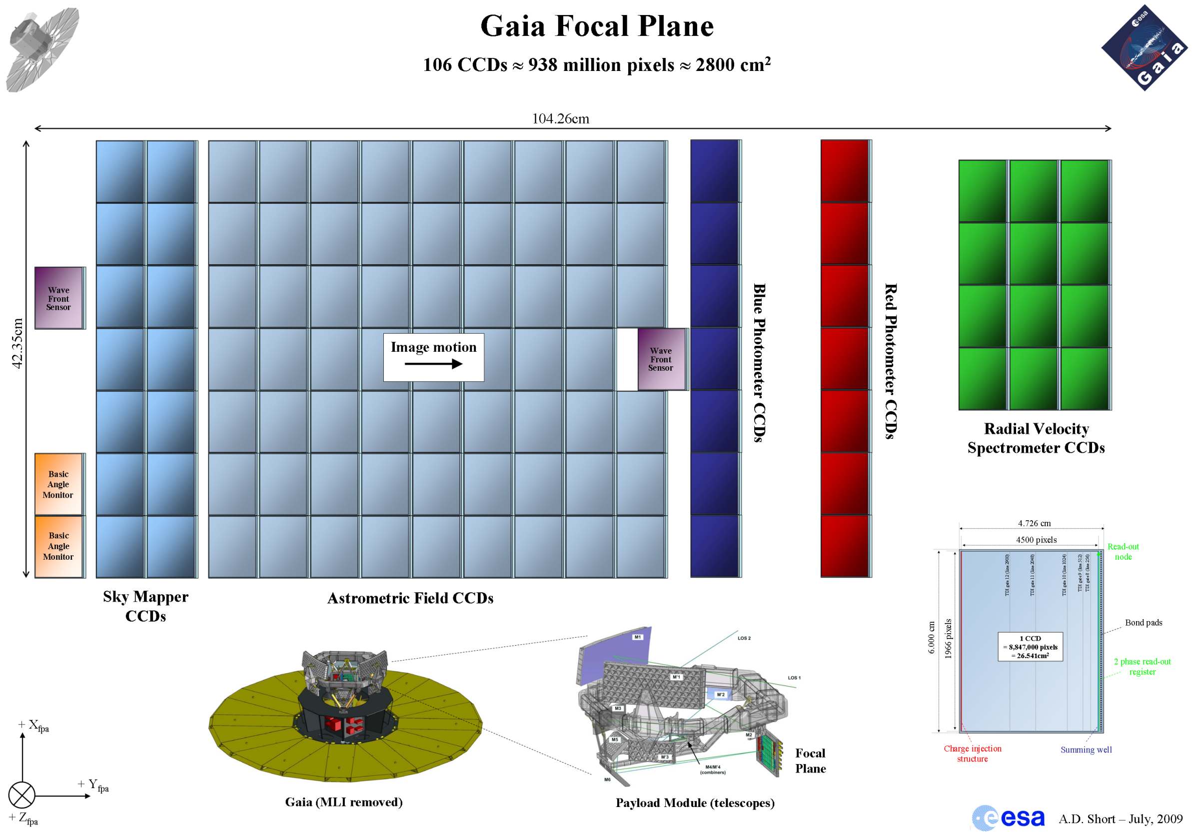

Figure 3: The Gaia focal plane. The viewing directions of both telescopes are superimposed on this common focal plane which features 7 CCD rows, 17 CCD strips, and 106 large-format CCDs, each with 4500 TDI lines, 1966 pixel columns, and pixels of size 10 μm along scan $\times$ 30 μm across scan (59 mas $\times$ 177 mas). Star images cross the focal plane in the direction indicated by the arrow. Figure courtesy of ESA - A. Short.

3.5 Gaia Processing: AlertPipe

The Gaia Science Alerts Development Unit (DU17) falls under the responsibility of the Gaia DPAC within CU5 (Coordination Unit 5: Photometric Processing). CU5 is coordinated by Floor Van Leeuwen (IoA). Our software is an end-to-end pipeline which will digest data released from ESAC, and spit out alerts at the end. It is called *AlertPipe* and is being developed principally by Łukasz Wyrzykowski (IoA), Simon Hodgkin (IoA, workpackage manager) and Ross Burgon (OU), although with significant input and consultation with key people at the IoA and elsewhere (Sergey Koposov and Chris Peltzer in particular).

3.6 Alert Reporting

The alerts are required to conform to readability standards and dissemination should be autonomous, fast and through many different means giving subscribers the most useful, and timely, information possible. A number of alert dissemination methods and alert formats are discussed in a separate document[1]. Each method introduced has its merits and issues and these are discussed fully. The proposed alert dissemination and format for the Gaia Science Alerts publication system

[1]Gaia-C5-TN-OU-RBG-001

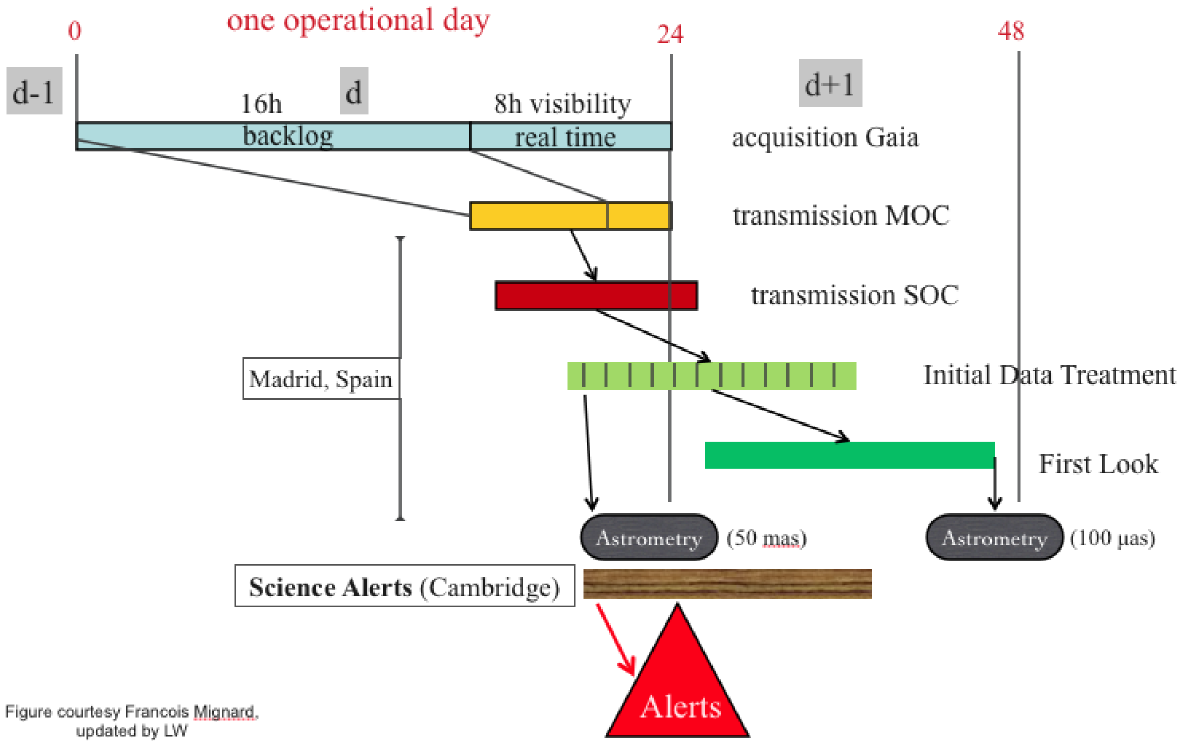

Figure 4: Scheme of the data flow and its timings. Gaia data will be transmitted only during 8h visibility window. Initial Data Treatment (IDT) will process data in chunks, each of them will be send to Cambridge. Figure courtesy of F. Mignard.

includes the use of e-mail, a VOEventNet-like[2] service and a dedicated website for alert dissemination and the use of the VOEvent standard as the alert format. This combination should give the alerts publication system the speed, flexibility, reach and robustness it requires.

4 Science Performance

from: http://www.rssd.esa.int/index.php?project=GAIA&page=Science_Performance

4.1 Astrometric Performance

	B1V	G2V	M6V
V < 10	$< 7\mu$as	$< 7\mu$as	$< 7\mu$as
V = 15	$< 25\mu$as	$< 24\mu$as	$< 12\mu$as
V < 20	$< 300\mu$as	$< 300\mu$as	$< 100\mu$as

Table 1: End-of-mission parallax standard errors for unreddened stars averaged over the sky shall comply with the above requirements.

In principle, in a single transit the astrometric performance (across scan) is in the 20-30 μas range for G=10-13 falling to 200 μas at G=17 and 650 μas at G=19 (see Figure 5). However there is a

[2] www.skyalert.org

significant delay for any given source observation to be mapped onto the global astrometric solution (assuming 6 monthly processing windows in the intial design). The IDT includes an astrometric solution (OGA1) with precision around 10-50 mas with a systematic error of around 50mas. OGA2 (produced by First Look) reduces the error to around 100μas and should be available within 24h after IDT (TBC).

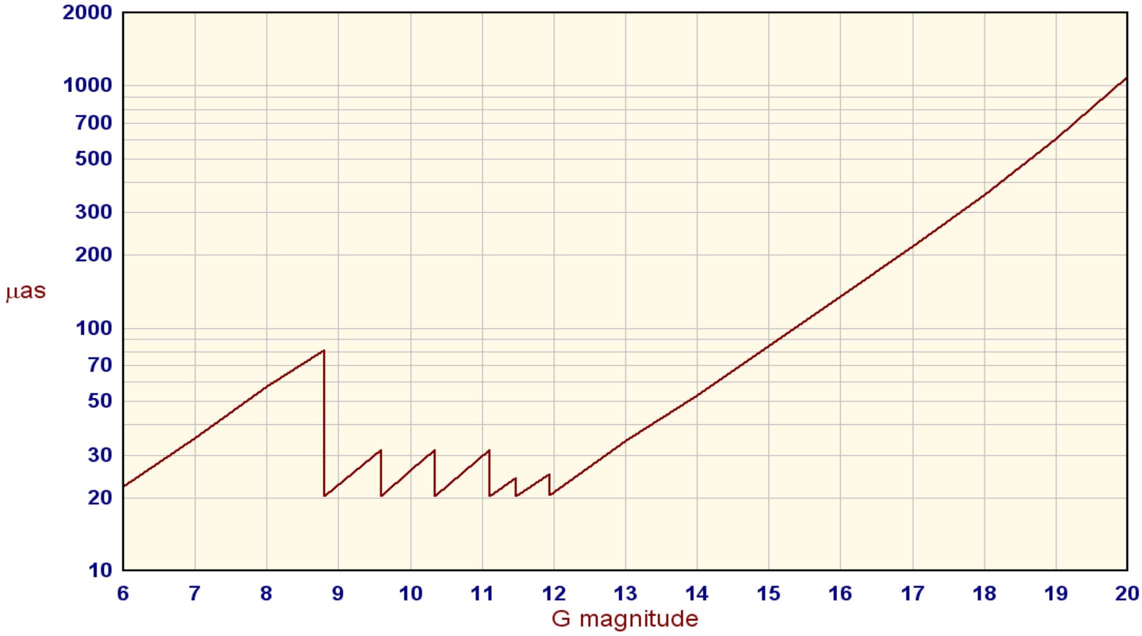

Figure 5: Astrometric performance across-scan for a point source from a single transit as a function of G magnitude.

4.2 Photometric Performance

The main photometry stream from Gaia is obtained from AF CCDs and is a broad filter photometry called G. Another source of the photometry is based on low-resolution, dispersive, spectrophotometry using the Blue and Red Photometers (BP and RP), from which we derive G_{BP}, G_{RP} magnitudes. DU17 will receive both streams of the photometry in uncalibrated form. Figure 6 shows the expected error as a function of magnitude for each field-of-view transit.

4.3 Spectroscopic Performance

The underlying low-resolution epoch spectra from BP and RP will also be available from the IDT in a raw form.

The BP/RP spectrograph comprises two low-resolution fused-silica prisms. The BP disperser covers 330-680 nm with resolution 4-32 nm/pixel, while the RP covers 640-1000 nm with resolution 7-15 nm/pixel.

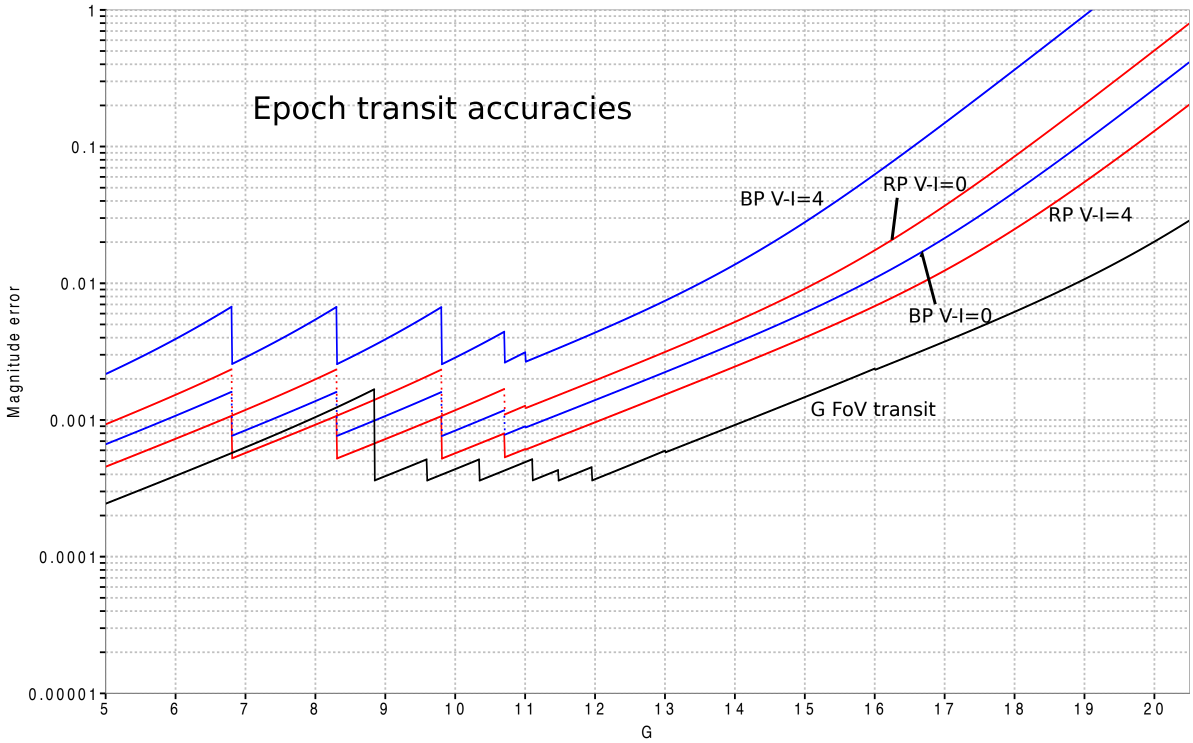

Figure 6:

The onboard high resolution spectrograph (RVS) has a resolution $\sim$ 11500 and radial velocities will be measured for all stars (roughly 150 million) with V< 17. The RVS data is being processed on daily basis by CU6 in Toulouse and in principle could be available for access by the Science Alerts pipeline (TBC).

5 Detection and Classification Methods

Gaia data will provide a unique opportunity for almost-real-time detection of all-sky astrophysical phenomena. Details of the detection and classification of the alerts are still in the process of design and development. Below we outline our current design.

5.1 Detection of anomalies

The detection will be two-fold: (i) known sources exhibiting anomalous behaviour and (ii) the appearance of new sources. Both types of alert detection will rely solely on the fluxes observed in the Astrometric Field. As mentioned above, at that stage for a single observation (transit) all per-CCD AF flux measurements will be available for the analysis, hence the detection can be performed on high temporal resolution data.

Anomaly detection with an "old source" will rely on the history of that source, ideally incorporating both Gaia observations collected so far and historic ground-based data. Within a single processing batch of data we will investigate the most current observed transits (usually one or two) for each object by comparing them to the available historic measurements. The transit will be regarded as anomalous if its fluxes (8 or 9 per-CCD fluxes) deviate significantly from the previous measurements and are consistent with one another (i.e. to avoid triggering an alert because of a cosmic ray hitting one of the CCDs).

One of the simplest detection algorithms calculates the mean and *rms* of all previous measurements of a source. The alert is triggered on a transit if at least 5 of all AF CCDs fluxes deviate by more than 3 sigma from the mean historic flux.

New sources are defined as not existing in the Gaia database despite previous observation. Naturally, triggers on new sources will only be valid after enough data was gathered by the mission and the whole sky was observed at least once or twice (to assure we are not triggering on a source observed previously but not transmitted to the ground due to e.g. crowding).

The preliminary detection algorithm for new sources is fairly simple: we will trigger an alert for a source we did not see before which has its brightness significantly (e.g. 1 mag) above the detection limit for Gaia (around 20 mag). These numbers obviously required tuning from simulation and during the mission.

5.2 Preliminary classification of alerts

Preliminary alerts found during the detection process are filtered using as much available information as we can find. Some examples are given here.

5.2.1 Light curve classification

The first level of classification will be performed on the morphology of the light curve, e.g. amplitude and slope. The more data points near the moment of detection of an anomaly, the better our characterisation, hence the higher the confidence and reliability of the classification.

We are currently exploring the classification technique based on Bayesian Classifier with Gaussian Mixture, developed and presented in Deboscher et al. (2007), which was applied to the classification of periodic variable stars. In our variant of the application we can not use the periodogram (Fourier decomposition of a light curve) and have to come up with other sets of parameters. In general each light curve we will be investigating will be composed of some historic measurements (for old sources only, though) and one or two field-of-view transits, each containing 8-9 flux measurements. Such lightcurves can be described with a combination of amplitudes and slopes (change of flux in time).

In order to apply the technique we first have to obtain a set of realistic distributions for each of the parameters. This is performed using simulations of various types of anomalies (e.g. dwarf nova, microlensing event, M-dwarf flare) with realistic Gaia sampling and noise. After that training process we obtain a model for each class of anomalies, which then can be used in classifying the incoming data. As an outcome from the classifier the probabilities are given for given anomaly belonging to each of the trained classes.

The training process in this method relies solely on the simulated data and is subject to how it approximates the real data. During the early stage of the mission the verification of this classification method and the models trained will have to take place to avoid biases and false-alerts.

5.2.2 Spectral classification

One of the main features of Gaia are almost-simultaneous low-resolution spectra. This allows for a significant reduction in false-alarms as the spectrum of an object can, in most cases, narrow down the classification of an anomaly to a single class.

In order to classify the spectra we employed the Self-Organizing Maps (SOMs) method (also called Kohonen Maps). A nice feature of the SOM is its simplicity in coding, and it's speed. As an input to a SOM we use the raw BP and RP spectra simply concatenated to one another. During the training process (to take place well before the mission) numerous simulated BP/RP spectra are being presented to the SOM, which then sorts them in a completely unsupervised mode. After the training we obtain a map in which the spectra are arranged in such a way that the similar spectra are located very "near" to each other and very different spectra are "far". Distance is defined usually by Euclidian difference between spectra.

Then, on such trained SOM, known spectral types, temperatures and absolute magnitudes can be mapped back, by calculating mean value of these parameters for all input patterns which ended up building a given node of the map. Any new incoming spectra can be quickly and easily linked with corresponding spectral type or absolute magnitude.

In case the real data (raw spectra) are very different from the simulated ones for which the spectral type is well known, another option is possible. Because of the natural property of the SOM to group similar patterns together, we can train the SOM during the mission with all incoming spectra. After a period of time, the spectra will be sorted out according to their shapes. Our SOM can then be calibrated, by locating spectra of stars with known spectral type on a map. The classification of new incoming spectra can be performed as for simulated spectra.

SOM can be also used for detecting unusual spectral behaviour. If SOM was trained with a stellar spectra only, then classification of any other spectrum (e.g. of supernova) will have significantly large quantization error (QErr), defined as a "distance" from current spectrum and the best matched one from the SOM. The distance can be also calculated from the mean node of the SOM. Such feature of a SOM can be very useful during the classification process as any flux anomaly accompanied with large QErr is more likely to be scientifically interesting.

5.2.3 Cross-matching classification

At the final stage of the classification process, each potential alert will be cross-matched against existing ground- and space-based catalogues and other alerts from other surveys.

6 Challenges

We went into this meeting having outlined a number of specific challenges that we tried to address. This document has evolved to record progress made against them, and incorporates new issues as they arose. The challenges can be broadly grouped into subject area.

6.1 Science and targets

- What are the most exciting targets of opportunity for the Gaia Science Alerts stream?

  At the meeting many communities and research areas were represented. Some comments on what we should look for include:

  Ultra-luminous Surpernovae: *Peculiar light curves, U band magnitude reaching -23, host galaxies faint, e.g. Quimby et al. (2010). And the relationship between GRBs and SNe.*

  Microlensing and Planets: *High magnitude events (small β) are the most suitable for the detection of planets.*

Astrometry of Microlensing Events: *There is a case for integrating astrometric and photometric data for microlensing events, for example one can measure the angular Einstein Radius (displacement of centroid of light is proportional to size of the Einstein Ring). For some long time-scale events the effect of Gaia orbiting the Sun can cause perturbations, which combined with the angular Einstein Radius can lead to solving for the mass of lensing object and search for Black Holes or Brown Dwarfs. Black Hole events will cause astrometric microlensing signals around 300 microarcseconds (distance dependent). Brown Dwarf events are an order of magnitude smaller.*

GRBs and Orphan Afterglows: *Although the detection efficiency of the afterglows is very low in Gaia due to it sparse sampling, there is still a room for Gaia detecting the brightest (and usually the longest) GRBs and alerting about them.*

Synergies: *There is an overlap possible with high energy missions: XMM-Newton, MAXI/ISS (whole sky in 90 mins 0.5-30 keV), Swift (all-sky 15-150 keV - but has ToO capability with the XRT). Future missions include SVOM (Swift like), JANUS (proposal: NIRT+XCAT), ASTROSAT. See Paul O'Brien's talk.*

- How can we improve our predictions of the expected event rates for the various transient classes.

  Our simulations are currently rather primitive, and much work needs to be done to increase the reality of the lightcurves passing through AlertPipe (we have designed templates to do this), while at the same time developing our detection algorithms and decision tree beyond their essentially placeholder current designs.

  Future simulation work will include analysis of the onboard detection algorithm: specifically to see how it behaves for point sources in the environments of extended sources. We understand that the next release (2010) of GIBIS (The Gaia pixel-level simulator) will include the onboard detection algorithm.

  Apart from extending our simulations, we can see some interesting statistics from other surveys, e.g. for PTF (subject to biases - but it sounds like their follow-up is pretty unbiased - does this mean of supernova type ?). Of 600 transients followed up, they breakdown as: 59% SN Ia, 20% SN II, 5% SN Ib/c, 5% CV, 11% unclassified.

6.2 Publishing Alerts

- How do we minimize contamination, either from astronomically 'uninteresting' events, or spurious onboard detections (e.g. cosmic rays, mis-identified sources, moving objects and so on).

  As far as astronomically 'uninteresting' events go, the pervading view of the meeting could be summarised (Wozniak) as "One man's trash is another man's gold". As long as the data stream is fully and reliably classified, then the need to filter the stream reduces. For spurious detections, it's not clear what an acceptable contamination level is - as long as we can give reliable probabilities then this should not be a problem (subject to the maximum event rates we and our subscribers can reasonably handle).

- What is the minimum information that the follow-up observers require? What would they like in an ideal world ?

- How will we disseminate the information?

This is now documented in a techincal note[3] *(not yet public). Basically via VOEvent and Email. See Section 3.6.*

- Will we need to run a local server to store all the additional images/lightcurves and so on? Or can this be outsourced?

  As detailed in the above document, we will run a local server for ancilliary/more detailed information.

- We'd like feedback on our proposed detection algorithms, and input on possible improvements and enhancements to their design. For example, how important are Gaia BP/RP spectra (both accumulated and concurrent)?

  It's becoming clear that the BP/RP spectra will actually prove to be very useful in terms of classification - this is perhaps a previously underrated ability of Gaia by the Flux Alerts methodology. The problem will be the calibration of these spectra.

  We saw some interesting thresholding analaysis from the Raptor project based on Parzen windows to construct density estimates for variability as a function of magnitude. It will handle structures in the rms diagram for example (e.g. from gates)

- What's the best way to handle disappearing objects?

- Do we have a good handle on how to avoid contamination from moving objects.

- For alert classification, what lessons can learn from ongoing transient surveys? What order should we cross-match against ground-based catalogues/databases? Which are the most useful resources?

  The Panstarrs SN search uses a large and complex decision tree. The biggest headache for them is the large numbers of false positives. For example: 1.8 million objects are initially detected as transients. 0.5 million of these are rejected immediately based on DQC flags. 0.6 million recurrent objects are also rejected, leaving 0.6 million unique objects. 99% of these are then rejected using image quality flags to leave 6500 'genuine' unique objects which are passed into the decision tree for classification. Then they are eyeballed, and only 280 survive.

  Also from Panstarrs: Remote VO crossmatching was too time consuming, and they resorted to a local VO database (a few months of effort). Of interest, their team comprises a total of 6 people in Hawaii and 3 in QUB.

  Raptor use a Bayesian Belief Network: depending on such source properties as the presence of a catalogue progenitor, the Galactic latitude, the presence of a host galaxy in SDSS, and statistics of the lightcurve itself. One advantage is that it does not require all the data to solve the network - one or two inputs should be enough.

  PTF use a Real/Bogus decision based on 'expert' human eyeballing of difference images (aka Group-Think). They have a million candidates per night. Poznanski cites the example of Supernova Zoo (Galaxy Zoo) where there are people with 20,000 classifications.

- How do we estimate the human resources required to operate the pipeline (to the point of releasing alerts)? How much development should we expect during early operations (for how long)? Who is responsible for keeping house: i.e. tracking ongoing follow-up and maintaining a history.. evolving the classification and priority of any given alert.

  Preliminary estimates have gone in to a draft of the Gaia Operations grant bid. It will be illustrative to examine other project personpower requirements, e.g. Smartt indicated that PS1

[3] Gaia-C5-TN-OU-RBG-001

has a total of about 9 FTE, 3 of which are based in QUB. They plan to publish their alerts. We are currently aiming for a minimum of 2.5 FTE in routine operations mode, though with more during ramp up and development in the first year of the mission.

- Are we doing anything that other elements of the Gaia processing or science teams might be interested in (e.g. health monitoring of the arrays, day-to-day photometric calibration).
- When should the alert stream be turned on? Early (with high risk) or later when we understand things better?
- What kind of methods could we use for detecting anomalous photometric and spectrometric behaviour of objects given the Gaia's unique sky scanning law?

6.3 Follow-up Observations

One of the primary goals of the meeting is to develop a roadmap for the coordination and preparation of ground based observing campaigns.

- Who out there is keen to follow-up Gaia Alerts ? Is our perception that we are providing an interesting data stream merited?

  Significant interest was expressed from the SNe, Microlensing and CV communities. It seems clear that we need further follow-up discussions with these groups to discuss in more detail their requirements on the Alert Stream to be gauged against resources: both operational and developmental.

  An important point is that all ongoing transient surveys are severly limited in the amount of follow-up they can actually achieve. Therefore prioritisation of candidates is essential.

- What resources are required to follow-up the GSA sources? Which of these are/will-be in place?

  *Wozniak summarised the available ground based facilities as a function of telescope size, though I suspect this is northern hemisphere only: ****

  Ofek listed 13 telescopes involved in the follow-up of PTF, ranging from 1m (wise) to 10m (Keck).

- What needs to be done in order to prepare the way for succesful follow-up? And on what timescales ?

  *** input from Gerry ?**. Basically we should start writing LoI's prety soon for the verification phase.. see next section*

- Who is planning to coordinate follow-up of the GSA data stream ? What aspects of GSA most interests these individuals ?

  At the moment, no coordination of the follow-up of the GSA data stream has been planned or discussed. In fact to be clear, this is not something that the DPAC will coordinate - the GSA data stream is public to the world - and it is the responsibility of the various scientific communities to coordinate any such activity themselves. For Science Verification however, we (DPAC: CU5: DU17) have received a letter of intent from the Gaia Variable Star team (DPAC: CU7) to help with the process. We plan to organise/coordinate follow-up during Gaia Cycle 0.

The PTF statistics are very interesting, even with access to a large range of telescopes and significant time allocation, they are still only following up a very small fraction of their transients (13%). One reason could be the limited size of their collaboration. But another could be the restrictions on access to telescope time. The key for Gaia will be a large, coordinated and well motivated collaborative effort with significant telescope time. One bonus for Gaia is the relatively bright nature of the transients ($G < 19$ mag c.f. PTF which goes down to 21st mag), which means 4m follow-up is viable.

- How do we go about verifying our alert stream in the early phases of the mission using ground-based facilities. How long should such a programme operate for .. how is success measured?

  The meeting was essentially divided on this issue. On the one hand, a desire was expressed to see the data as early as possible and to involve the community in the verification of the Alert Stream. The other point of view (and this came principally from DPAC personnel) is that we should understand our Alert Stream and calibrate the probabilities/contamination in advance of making it live.

6.4 Outreach

Science alerts can be an attractive and interesting product of the Gaia satellite also for the general public, students and amateur astronomers. If they are provided with the right tools for receiving and simple processing of the alerts stream, they could follow the real-time science done by the professionals or even contribute to it.

- What form of alerts would be the most comprehensible for the general audience?
- What kind of tools or programs should be provided? E.g. supernova follow-up with the backyard/school telescope.

  Two obvious tools: (1) A watch list of interesting known erruptive variables - the lightcurves of all these stars will be made publically available, (2) A **Where is Gaia** *tool (i.e. where is it looking)*

- What kind of feedback from amateur observers would be useful for the scientific verification of the alerts?

6.5 Publicity

7 Planning Ahead

With the Gaia launch drawing ever nearer it's very important to put some milestones in place. Identifying these key milestones, their due dates, and associating them with responsible people is one of the main goals of the meeting

- Letters of Intent to relevant observatories concerning the Science Verification phase of Alert-Pipe
- Proposals for SV phase follow-up
- Letters of Intent to relevant observatories planning scale of planned GSA follow-up work

- Proposals for GSA scientific follow-up
- Finish pre-launch development of AlertPipe, tested on simulated data and fully documented.
- Secure funding for continued AlertPipe development and operation in post-launch era.

A Additional Questions (and answers) and Suggestions

- **Can alerts be done against a list of sources ?** Yes - but note that the results will be released to the public.
- **Can we have access to the raw data flow from DPAC ?** Politically there is no problem - but technically this does not fit into the funded scheme.
- **On a related note: Please don't try to filter things out of the alert stream: let others handle their own filtering.** This comes at a price - if we alert on everything at 2sigma - that's a lot of alerts for a billion sources.
- **Where is Gaia ?** An excellent suggestion - a tool/algorithm to display where and when Gaia is pointing prior to and during the mission.
- **Will we output the results from all CCDs or just per field-of-view?** The current model is to present per field-of-view photometry. If cost allows, and a requirement can be identified, then per-ccd photometry could be distributed.
- **How many alerts will there be - and what can we cope with?**
- **How will Gaia cope with observations in the Galactic Centre?** The spacecraft has two downlink modes which affect which data we receive, and on what timescales, in regions of high source density: e.g. the Galactic Centre.

B Unresolved Questions

1. What is the astrometric precision for a source straight out of IDT, and how does the error vary with time.

C Resources

1. Gaia Science Alerts Working Group WIKI: http://www.ast.cam.ac.uk/research/gsawg
2. GSAW2010 meeting website: http://www.ast.cam.ac.uk/research/gsawg/index.php/Workshop2010
3. Movies of GSAW2010 talks (also available from the workshop's agenda page) http://www.ast.cam.ac.uk/∼w

D Registered Participants

1. Dr Giuseppe Altavilla, Bologna, Italy
2. Dr Martin Altmann, Heidelberg, Germany

3. Dr David Bersier, Liverpool, UK
4. Dr Ross Burgon, Open University, UK
5. Dr Simon Clark, Open University, UK
6. Prof. Michel Dennefeld, Paris, France
7. Dr Martin Dominik, St.Andrews, UK
8. Prof. Wyn Evans, Cambridge, UK
9. Dr Laurent Eyer, Geneva, Switzerland
10. Dr Roger Ferlet, IAP, Paris, France
11. Dr Boris Gaensicke, Warwick, UK
12. Dr Avishay Gal-Yam, Weizmann, Israel
13. Prof. Gerry Gilmore, Cambridge, UK
14. Dr Andreja Gomboc, Ljubljana, Slovenia
15. Prof. Andy Gould, Ohio State University, US
16. Dr Simon Hodgkin, IoA, UK
17. Dr Rene Hudec, Ondrejov, Czech Republic
18. Dr Sergey Koposov, IoA, UK
19. Dr Agnes Kospal, Leiden, The Netherlands
20. Dr Pavel Koubsky, Ondrejov, Czech Republic
21. Dr Andrew Levan, Warwick, UK
22. Fraser Lewis, Cardiff, UK
23. Mr Mark ter Linden, ESAC, Spain
24. Dr Ashish Mahabal, Caltech, US
25. Prof. Lech Mankiewicz, PAN Warsaw, Poland
26. Dr Francois Mignard, OCA, Nice, France
27. Dr Nami Mowlavi, Geneva, Switzerland
28. Professor Tim Naylor, Exeter, UK
29. Dr Ludmila S. Nazarova, Moscow, Russia
30. Dr Yael Naze, University of Liege, Belgium
31. Prof. Paul O'Brien, University of Leicester, UK
32. Dr Eran Ofek, Caltech, US
33. Dr Patricio F. Ortiz, Leicester, UK

34. Prof. Don Pollacco, Belfast, UK
35. Dr Dovi Poznanski, Berkeley, US
36. Dr Timo Prusti, ESA
37. Dr Jenny Richardson, Cambridge, UK
38. Dr Sarah Roberts, Cardiff, UK
39. Dr George Seabroke, UCL, UK
40. Prof. Stephen Smartt, Queen's University Belfast, UK
41. Dr Danny Steeghs, Warwick, UK
42. Prof. Iain Steele, Liverpool, UK
43. Dr Mark Sullivan, Oxford, UK
44. Dr Damien Segransan, University of Geneva, Switzerland
45. Dr Paolo Tanga, Nice, France
46. Prof Nial Tanvir, Leicester, UK
47. Dr Patrick Tisserand, Mount Stromlo Observatory, Australia
48. Dr Massimo Turatto, Catania, Italy
49. Dr Nick Walton, Cambridge, UK
50. Mr Marcin Wardak, Warsaw, Poland
51. Dr Peter Wheatley, Warwick, UK
52. Dr Przemek Wozniak, Los Alamos, US
53. Dr Lukasz Wyrzykowski, IoA, UK
54. Mr Kamil Zloczewski, Warsaw, Poland

E References

Quimby et al. 2010, Nature, submitted (http://arxiv.org/abs/0910.0059)