
7. On-board detection

GAIA WAS ACCEPTED by ESA's Science Programme Committee in 2000. This followed a 3-year feasibility study, led by me (as Project Scientist) and Oscar Pace (as Study Manager), supported by a Science Advisory Group and various industrial teams. Findings appeared as a 380-page report, ESA-SCI(2000)4, in July 2000.

We identified 15 preparatory technologies needed to ensure that the satellite could be developed on schedule, and within budget. Five of these related to advanced CCD performances not available at the time, including 3-side buttable, small-pixel, high-performance chips; the large-area highly-integrated focal plane assembly; and high-speed, low-noise detection chains.

Associated developments were required for efficient on-board compression algorithms for the science data; optimisation of the payload data-handling electronics; and a (non-moving) phased-array antenna suitable for the high data-rate transmission from Gaia's L2 orbit.

Other technology advances were required for the large silicon-carbide mirrors and ultra-stable payload structure; the 10-m deployable solar array sunshield; the micro-Newton reaction system for the fine attitude control; and inch-worm actuators for telescope refocusing.

IN NORMAL CCD imaging, individual pixels accumulate photoelectrons during an exposure, after which the entire CCD is 'read out': pixel columns are successively advanced electronically towards a readout register.

In contrast, and central to the use of the Gaia CCDs, is the method of 'time-delayed integration', or TDI. Here, star images move slowly across the detector, and all CCD columns are stepped continuously towards the readout register, at a rate precisely synchronised with the satellite rotation. Exposures build up as the satellite scans the sky. There is no pausing for discrete image read outs.

The whole challenge of onboard source detection is intimately tied to the detailed CCD technical performance, and the entire system function depends on the adopted readout rate of the columns and rows, the readout noise in the electronics detection chain, and a host of other effects such as full-well capacity and charge overflow, and radiation damage and charge-trapping.

WHILE THERE would be some scientific merit in transmitting the entire data stream from the focal plane CCDs to the ground, certain issues conspire to make this approach unrealistic. Full focal plane read-out in TDI mode would result in data rates of gigabits per second, compared with a realistic limit on the sustained telemetry rate from the L2 orbit of a few megabits per second. Even in the absence of telemetry rate limits, the rapid read-out rates would result in a high 'read-out noise', and a degraded signal-to-noise ratio per pixel.

Detailed analysis showed that full focal plane read-out would, in any case, be of little value: for most of the sky, especially out of the Galactic plane, the fractional area covered by stars is very small, even at 20 mag.

The down-link of small patches of the detector field, centred on each detected object, and extending a little beyond the Airy disk, would instead be a satisfactory solution. Data rates would then drop to manageable levels, and the read-out noise per pixel could be brought down to adequate levels, even for the faintest stars.

Given that the focal plane data on all objects down to about 20–21 mag could be read out and telemetered to ground, the central problem was then how to identify the patches of sky containing these objects. Two options were evaluated: the use of an input catalogue; and the on-board detection of targets.

SEVERAL PROBLEMS made the use of an input catalogue unattractive if not unfeasible. In the first place, there was no existing catalogue which would suit Gaia's needs. And any plausible attempt to create one could never replicate the instrument's broad, red spectral response, and its sub-arcsec angular resolution, with strong implications for biases in the resulting measurements. Another major argument against this approach was the inevitable exclusion of transients, such as variable stars, burst sources, and solar system objects.

Scientifically, the issue was not only an operational one: a clear definition of the selection function used to decide which targets to observe (for example, catalogue completeness as a function of magnitude) would also be a key issue affecting the mission's scientific value.

THE SOLUTION ADOPTED was to detect every prospective target on-board, by means of a dedicated ‘sky mapper’. This would also have the merit that the detection is carried out in the same wavelength band, and at the same angular resolution, as the final observations.

Converging on the finally chosen design was the result of more than 100 technical notes from the scientific side, many advanced by Erik Høg (Copenhagen), and many more from the industrial prime contractor, EADS Astrium (now Airbus Group), Toulouse.

In practice, existing catalogues were used to facilitate early attitude determination and source matching.

SUPPORTING THE many design choices were two very detailed accuracy assessment budgets: one maintained on the scientific side, under my responsibility, by Jos de Bruijne, and one on the industrial side by the prime contractor’s chief system engineer, Frédéric Safa.


Every possible effect that the scientific and engineering teams could identify that might have an influence on the final astrometric accuracies were assessed, tabulated, and fed into these global accuracy budgets.

IN ITS FINAL configuration, Gaia consists of two telescopes, separated by $106^{\circ}5$, feeding a common focal plane. The spacecraft rotates around an axis perpendicular to the two viewing directions, once every 6 hours.

The focal plane comprises 106 CCDs in seven rows, each with its own autonomous control unit. It takes a bit more than a minute for star images to drift across the entire 1.5° focal plane in the along-scan direction. Operating in TDI mode, the scan rate is adjusted to coincide with the clock rate of the CCD readout. Each line of pixels is shifted in a little less than 1 ms.

The first two vertical CCD strips, SM1 and SM2, are the sky mappers, and perform the onboard object detection. Optical baffles ensure that each sees only one telescope field. In the absence of prior star knowledge, all pixels are read out, at the expense of a higher read noise.

Fast onboard analysis then produces a list of point-like objects for observation in the following CCDs (in which both fields are superimposed), rejects isolated cosmic ray events, and determines the scan rate about both axes. This information is provided to the attitude control subsystem, which allows the star positions to be predicted for the remaining focal plane crossing.




BEYOND THE astrometric field, star images cross the blue and red photometers, where prisms provide low-resolution spectra used to derive star colours. Finally, images reach the radial velocity spectrometer, providing high-resolution spectra for the brighter stars.

Various other tasks rely on measurements carried out in the focal plane. CCD ‘gating’ restricts the integration time for bright stars, and allows the measurement of objects that would otherwise cause detector saturation. Charge-injection mitigates some of the problems of radiation-induced charge trapping. Two CCDs are dedicated to wave-front sensors which allow telescope focussing. Two further CCDs are dedicated to a specific laser metrology system which monitors any tiny changes in the angle between the two viewing directions.

ONLY SMALL WINDOWS around the predicted positions of each object are read out. These windows are 12 pixels across scan (2.1 arcsec). Along scan, they are either 18 pixels (1.1 arcsec) for $G = 13 - 16$, or 12 pixels (0.7 arcsec) for $G > 16$ mag.

For $G > 13$ mag, windows are binned across scan to reduce the CCD readout noise. Rules determine which windows are used in very dense sky regions, and in conflicts due to close binary stars.

Images typically have a different *transverse* speed for each field, due to the satellite’s precessional scan motion. With an integration time of 4.42 seconds per CCD, the transverse motion can reach 4.5 pixels over a single CCD, and can therefore result in significant across-scan smearing. As most samples are binned in this direction, the net effect on the observations is small.



THE ELONGATED CCD pixels, I should emphasise, are a consequence of the telescope’s rectangular primary mirrors ($1.45 \times 0.5 \text{ m}^2$), which result in an asymmetric point spread function with aspect ratio 3:1. The image width in the scan direction is about 0.1 arcsec, determining Gaia’s ultimate resolution as an imaging system.

The CCDs have 4500 pixels in the along-scan direction, and 1966 pixels in the across-scan direction. Each pixel is rectangular in the same 3:1 ratio, the 10×30 micron pixels corresponding to $59 \times 177 \text{ mas}^2$ on the sky.

Since each source is observed some 70 times with different scan directions during the nominal 5-year mission, the one-dimensional scans can eventually be used to create a two-dimensional reconstructed image.

MANY MORE details and examples of the onboard detection and windowing are given by Fabricius et al. (2016) and Arenou et al. (2018).