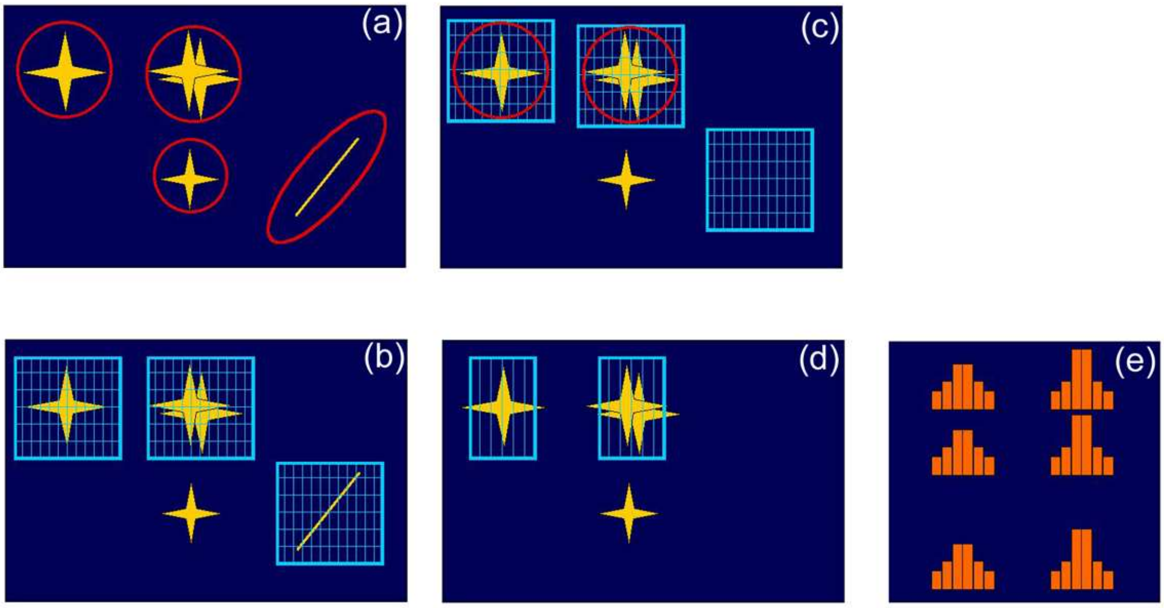

The on-board data handling consists of a sequence of operations starting with (a) object detection, (b) first selection, (c) actual observation in the next CCD strip and confirmation, (d) final selection, (e) observation in the following CCD strips, post-processing with packetisation, and compression for subsequent downloading.

Since Gaia will perform a continuous all-sky survey at high angular resolution, the on-board processing, which must operate autonomously, needs to be able to cope with virtually all the varieties and peculiarities of the real sky, both in terms of the nature of objects (bright stars, multiple stars, nebulosities, solar-system objects, planets, etc.) and object densities (Galactic pole, Baade's Window, cores of globular clusters, etc.). Generic and adaptive algorithms are therefore required.

Given real-time processing constraints, limits on the acceptable CCD read-out noise, and the limited telemetry bandwidth, not all CCD pixel data can be read and subsequently transmitted to the ground. A limited number of 'windows', regions of interest around target objects, are therefore observed in the focal plane, thus effectively removing the 'empty space between the stars' from the data stream. The object detection and confirmation tasks, aimed at distinguishing real objects from prompt-particle events such as cosmic rays and solar protons, thus have to be accompanied by a selection step which decides which sources are tracked in the remaining CCDs and how these objects are observed.

From a scientific point of view, a statistical analysis of the final Gaia Catalogue, and especially of selection effects, requires that the on-board selection process is entirely and exactly reproducible on the ground. For this reason, some detection parameters will need to be downloaded for objects which are detected and confirmed but not observed because of selection criteria.

In general, the on-board data handling should allow the observation of the maximum number of objects in the best possible conditions. Multiple stars in particular pose a challenge in this respect, given the limited number of windows available. Similarly, optimum window placement in dense fields is non-trivial given the fact that the computing requirements should remain modest.

A Payload Data Handling Electronics demonstration study under an industrial contract has been conducted. This was aimed at designing a possible overall implementation architecture, and developing and integrating a representative breadboard, taking constraints in terms of scientific requirements and mass, volume, and power budgets into account. Based on this study, the current implementation architecture distributes the on-board processing between dedicated hardware (field programmable gate arrays) for 'pixel-based operations' and software for 'floating-point operations'.