



Contribution ID: 71

Type: **not specified**

SPARK Demo & Presentation

Tuesday, 13 December 2016 16:25 (30 minutes)

To face the increasing volume of data we will have to manage in the coming years, we are testing and prototyping implementations in the Big Data domain (both data and processing).

The CDS proposes a “X-Match” service which does a cross correlation of sources between very large catalogues (1 billions rows). It is a fuzzy join between two tables of several hundred millions of lines (e.g. 470,992,970 sources for 2MASS). A user can do a cross-match of the (over 10,000) catalogues proposed by the CDS or he can upload his own table (with positions) to cross-match it with these catalogues. It is based on optimized developments implemented on a well-sized server. The area concerned by the cross-match can be the full Sky (which involves all the sources), a cone with only the sources (which are at a certain angular distance from a given position), or a HEALPix cell.

This kind of treatment is potentially “heavy” and requires appropriate techniques (data structure and computing algorithm) to ensure good performances and to enable its use in online services.

Apache Spark seemed very promising and we decided to improve the algorithms, by using this technology in a suitable technical environment and by testing it with large datasets. Compared to Hadoop, Spark is designed to work as much as possible in the memory.

We performed comparative tests with our X-Match service and we reached an execution time better than the X-Match service. We will detail this experiment step by step and show the corresponding metrics.

We will focus on the bottleneck we encountered during the shuffle phase of Spark and especially the difficulty to enable the « data co-location » which is a way to decrease the data exchange between the nodes. An illustration of how Spark works will be done through a quick demo

Presenter: SCHAAFF, Andre (Strasbourg astronomical Data Center (CDS))