

Minutes of the Offline upgrade forum

CERN, June 27, 2012

Status: final

Present: F.Carminati, M.Niculescu, A.Gheata, L.Betev, P.Hristov, M.Litmaath, R.Grosso, C.Cheshkov

Action list

No previous action list.

Discussion (F.Carminati)

F.Carminati discusses the plans for the ALICE Offline upgrade. The time of “free increase” of the computing resources available thanks to the increase in the clock cycle is over. Moore’s law still holds and the available computing power still doubles every 18 months, however to harness it the code has to be transformed in order to exploit the available parallelism at different levels.

F.Carminati presents the opportunities to increase program performance:

- Vectorisation (SIMD);
- Instruction pilelining (IP);
- Multiple Level parallelism (ILP);
- GPUs, MICs and similar devices;
- HyperThreading (HT);
- Multiple Core parallelism (MC);
- Multiple Socket parallelism (MS);
- Multiple Node parallelism (MN);
- Grid – level parallelism (Grid);

F.Carminati says that all these levels have to be explored and several of them will possibly be combined.

F.Carminati exposed the timescale of the upgrade. The time till the end of the year will be devoted to R&D in order to prepare a plan for the parallelisation of AliRoot. This will be carried out during the first Long Shutdown in the years 2013-2014. The idea is to arrive at the end of the two years with a new parallel version of AliRoot replacing completely the current one. Failure to do so would mean that the parallel version would not be uptaken, as when the data will flow only a validated version will be used.

The two possible approaches are a complete rewrite and a progressive parallelisation. This can be top-down, starting with multi-threaded, event level parallelisation or bottom-up, optimising algorithms on CPUs and GPUs. The experience gained with the G5 prototype shows that no bottom-up approach alone would provide results unless concurrency is exposed at higher level.

M.Litmaath said that this approach proved very productive for CMS in reducing memory footprint, albeit forking processes rather than spawning threads.

After discussion it is therefore decided to start an investigation of the possibilities to run the AliRoot reconstruction in multi-threaded mode at the event level, sharing geometry and OCDB. Once this done, next step will be to consider parallelisation at the sub-detector level.

For the simulation, A.Gheata will continue to develop the simulation prototype which is based on a complete rewrite of the program.

P.Hristov proposed to involve more people in this forum, and in particular Peter proposed to involve people such as M.Magner (experience with GPUs: ITS vertexing and analysis code) and other people who might have experience and involvement in AliRoot & parallelization.

It was decided that this meeting would be held every week for the moment only for invited people. We will see later how to generalise or expand it. In case of absence of F.Carminati, A.Gheata will chair the meeting. M.Niculescu kindly accepted to take the minutes.

The meeting was closed and it was decided to postpone A.Gheata's presentation for next time.

AOB

None.

New action list

1. None for the moment