Paraver hands-on for HARMONIE

This document describes the steps to be followed so that users can get the same Paraver windows showed during the «BSC-HIRLAM collaboration: HARMONIE code profiling current status» presentation. It is intended to be interactively followed by everyone at the same time with the guidance of the lecturer.

This hands-on is thought to be performed on CCA, but you can alternatively do it on your personal laptop. You will find the paraver_harmonie_hands-on.tar.gz file on the following path of CCA:

/scratch/ms/spesiccf/c3xy

Copy it into your personal scratch and uncompress it using “tar zxfv”. The resulting folder contains the following elements:

- metcoop25c_harmonie_detail_285-MPI_1-OpenMP_no_IO_server: contains the trace that we will use for this hands-on. It doesn’t have OpenMP events, since it is not the objective of this introductory hands-on.
- metcoop25c_harmonie_detail_143-MPI_2-OpenMP_no_IO_server: contains a trace with OpenMP events. It is not expected to be used for this hands-on, but someone might feel encouraged to do so.
- cfgs: contains the basic configuration files to get the Paraver windows of the presentation.
- wxparaver-4.8.2-Linux_x86_64.tar.bz2: compressed file with Paraver.
- wxparaver-4.8.2-Linux_x86_64_legacy.tar.bz2: compressed file with Paraver for legacy operating systems.

If you are using you personal laptop with Mac or Windows operating systems, you should download the corresponding Paraver version from tools.bsc.es/downloads. If you work on CCA, you should uncompress the wxparaver-4.8.2-Linux_x86_64_legacy.tar.bz2 file using “tar jxfv”.

Now you are ready to execute Paraver by executing its binary from the command line (if you work on Mac or Windows, just do double click):

> ./wxparaver-4.8.2-Linux_x86_64_legacy/bin/wxparaver

The Paraver window will pop up and you should load the trace by clicking File > Load Trace… and search for the MASTERODB_regular_time-step.prv trace file inside the metcoop25c_harmonie_detail_285-MPI_1-OpenMP_no_IO_server folder. If you list the contents of this folder, you will see the three following files:

- MASTERODB_regular_time-step.prv: the trace file itself. It contains the performance events.
- MASTERODB_regular_time-step.pcf: it describes the type of performance events contained in the .prv file.
- MASTERODB_regular_time-step.row: contains the ids of the MPI processes and the OpenMP threads.

You are ready to generate the first window of the presentation, which shows the MPI call functions. You need to load the corresponding configuration file (.cfg) from the cfgs folder. Click File > Load Configuration… and search for the “MPI_call.cfg” file. After a period of processing time a new window will pop up. What you see it is a complete regular time step of HARMONIE. At this point of the hands-on, the lecturer will interactively show you the basics actions that you can perform,

such as zooming both time scale and processes, print communication lines, visualize basic info, what are the colors, etc.

The rest of configuration files that the cfgs folder contains are the following (ordered by appearance in the presentation):

- useful_duration.cfg: shows the duration of the computational chunks.
- mpi_stats.cfg: shows the MPI profile to know the load balance, parallel efficiency, etc.
- instructions.cfg: shows the executed number of useful instructions.
- ipc.cfg: shows the useful Instructions per Cycle (IPC).
- Instr_correlated_with_IPC.cfg: shows a histogram of the useful instructions correlated with the IPC. Useful to detect workload imbalances.
- L3_Cache_misses_per_1000_instr.cfg: shows an aggregated line of all MPI processes with the ratio of L3 cache misses per 1000 useful instructions.

All the previous configurations will be explored and explained during the hands-on so that everyone is more or less synchronized and the lecturer can give details of each case.

So far we have seen the very basics of Paraver so we can already have an idea of the potential of both Extrae and Paraver. Extrae is also a key tool for the BSC tools workflow, but we will explain its basic usage in an other hands-on. In case of remaining free time for this hands-on, you can move to the next Bonus session where you can experiment by yourself different configuration files not attached in this hands-on.

Bonus

If you feel encouraged to explore more options of Paraver, feel free to “play” with any functionality. We suggest you to explore different configuration files included in the Paraver folder:

wxparaver-4.8.2-Linux_x86_64_legacy/cfgs

There are many different configuration files, from very basic aspects to more advanced options, and for MPI, OpenMP, hardware counters, etc. If you want to explore OpenMP, remind to change the loaded trace by the other one included in this hands-on:

metcoop25c_harmonie_detail_143-MPI_2-OpenMP_no_IO_server

Note that not all configuration files might be useful for our two traces since it depends on the available events in the trace. You can check the available performance events in the .pcf file. To do so, just search for “EVENT_TYPE”.