Developping XIOS with multithread : to accelerate the IO of climate models

June 22, 2018

1 Context

The simulation models of climate systems, running on a large number of computing resources can produce an important volume of data. At this scale, the IO and the post-treatement of data becomes a bottle-neck for the performance. In order to manage efficiently the data flux generated by the simulations, we use XIOS developped by the Institut Pierre Simon Laplace and Maison de la simulation.

XIOS, a libarary dedicated to intense calculates, allows us to easily and efficiently manage the parallel IO on the storage systems. XIOS uses the client/server scheme in which computing resources (server) are reserved exclusively for IO in order to minimize their impact on the performance of the climate models (client). The clients and servers are executed in parallel and communicate asynchronuously. In this way, the IO peaks can be smoothed out as data fluxes are send to server constantly throughout the simulation and the time spent on data writing on the server side can be overlapped completely by calculates on the client side.

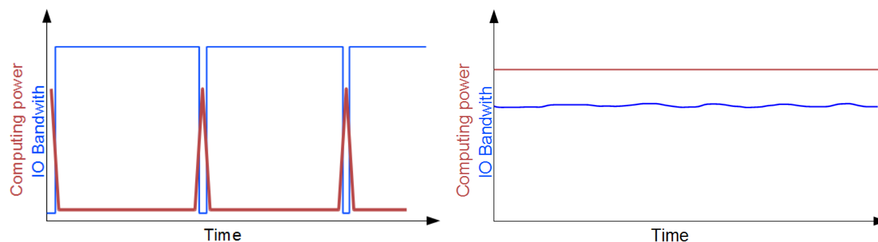

Figure 1: On the left, each peak of computing power corresponds to the vallay of memory bandwidth which shows that the computing resources are alternating between calculates and IO. ON the right, both curves are smooth which means that the computing resources have a stable charge of work, either calculates or IO.

XIOS works well with many climate simulation codes. For example, LMDZ[1],

[1] LMDZ is a general circulation model (or global climate model) developped since the 70s at the "Laboratoire de Météorologie Dynamique", which includes various variants for the Earth and other planets (Mars, Titan, Venus, Exoplanets). The 'Z' in LMDZ stands for "zoom" (and the 'LMD' is for 'Laboratoire de Météorologie Dynamique"). http://lmdz.lmd.jussieu.fr

NENO[2], ORCHIDEE[3], and DYNAMICO[4] all use XIOS as the output backend. MétéoFrance and MetOffice also choose XIOS to manege the IO for their models.

2 Developpement of a thread-friendly MPI for XIOS

XIOS is a library dedicated to IO management of climate code. It has a client-server pattern in which clients are in charge of computations and servers manage the reading and writing of files. The communication between clients and servers are handled by MPI. However, some of the climate models (*e.g.* LMDZ) nowadays use an hybrid programming policy. Within a shared memory node, OpenMP directives are used to manage message exchanges. In such configuration, XIOS can not take full advantages of the computing resources to maximize the performance. This is because XIOS can only work with MPI processes. Before each call of XIOS routines, threads of one MPI process must gather their information to the master thread who works as an MPI process. After the call, the master thread distributes the updated information among its slave threads. As result, all slave threads have to wait while the master thread calls the XIOS routines. This introduce extra synchronization into the model and leads to not optimized performance. Aware of this situation, we need to develop a new version of XIOS (EP_XIOS) which can work with threads, or in other words, can consider threads as they were processes. To do so, we introduce the MPI endpoints.

The MPI endpoints (EP) is a layer on top of an existing MPI Implementation. All MPI function, or in our work the functions used in XIOS, will be reimplemented in order to cope with OpenMP threads. The idea is that, in the MPI endpoints environment, each OpenMP thread will be associated with a unique rank and with an endpoint communicator. This rank (EP rank) will replace the role of the classic MPI rank and will be used in MPI communications. In order to successfully execute an MPI communication, for example `MPI_Send`, we know already which endpoints to be the receiver but not sufficient. We also need to know which MPI process should be involved in such communication. To identify the MPI rank, we added a "map" in the EP communicator in which the relation of all EP and MPI ranks can be easily obtained.

In XIOS, we used the "probe" technique to search for arrived messages and then performing the receive action. The principle is that sender processes execute the send operations as usual. However, to minimise the time spent on waiting incoming messages, the receiver processe performs in the first place the `MPI_Probe` function to check if a message destinated to it has been published. If yes, the process execute in the second place the `MPI_Recv` to receive the message. In this situation, if we introduce the threads, problems occur. The reason why the "probe" method is not suitable is that messages destinated to one certain

[2]Nucleus for European Modelling of the Ocean alias NEMO is a state-of-the-art modelling framework of ocean related engines. `https://www.nemo-ocean.eu`

[3]the land surface model of the IPSL (Institut Pierre Simon Laplace) Earth System Model. `https://orchidee.ipsl.fr`

[4]The DYNAMICO project develops a new dynamical core for LMD-Z, the atmospheric general circulation model (GCM) part of IPSL-CM Earth System Model. `http://www.lmd.polytechnique.fr/~dubos/DYNAMICO/`

process can be probed by any of its threads. Thus the message can be received by the wrong thread which gives errors.

To solve this problem, we introduce the "matching-probe" technique. The idea of the method is that each process is equiped with a local incoming message queue. All incoming message will be probed, sorted, and then stored in this queue according to their destination rank. Every time we call an MPI function, we firstly call the `MPI_Mprobe` function to get the handle to the incoming message. Then, we identify the destination thread rank and store the message handle inside the local queue of the target thread. After this, we perform the usual "probe" technique upon the local incoming message queue. In this way, we can assure the messages to be received by the right thread.

Another issue remains in this technique: how to identify the receiver's rank? The solution is to use the tag argument. In the MPI environment, a tag is an integer ranging from 0 to 2^{31}. We can explore the large range of the tag to store in it information about the source and destination thread ranks. We choose to limite the first 15 bits for the tag used in the classic MPI communication, the next 8 bits to the sender's thread rank, and the last 8 bits to the receiver's thread rank. In such way, with an extra analysis of the EP tag, we can identify the ranks of the sender and the receiver in any P2P communication. As results, we a thread probes a message, it knows exactly in which local queue should store the probed message.

With the global rank map, tag extension, and the matching-probe techniques, we are able to use any P2P communication in the endpoint environment. For the collective communications, we perform a step-by-step execution and no special technique is required. The most representative functions is the collective communications are `MPI_Gather` and `MPI_Bcast`. A step-by-step execution consists of 3 steps (not necessarily in this order): arrangement of the source data, execution of the MPI function by all master/root threads, distribution or arrangement of the data among threads.

For example, if we want to perform a broadcast operation, 2 steps are needed. Firstly, the root thread, along with the master threads of other processes, perform the classic `MPI_Bcast` operation. Secondly, the root thread, and the master threads send data to threads sharing the same process via local memory transfer. In another example for illustrating the `MPI_Gather` function, we also need 2 steps. First of all, data is gathered from slave threads to the master thread or the root thread. Next, the master thread and the root thread execute the `MPI_Gather` operation of complete the communication. Other collective calls such as `MPI_Scan`, `MPI_Reduce`, `MPI_Scatter` *etc* follow the same principle of step-by-step execution.

3 Performance of LMDZ using EP_XIOS

With the new version of XIOS, we are now capable of taking full advantages of the computing resources allocated by a simulation model when calling XIOS functions. All threads, can participate in XIOS as if they are MPI processes. We have tested the EP_XIOS in LMDZ and the performance results are very encouraging.

In our tests, we used 12 client processor with 8 threads each (96 XIOS clients in total), and one single-thread server processor. We have 2 output

densities. The light output gives mainly 2 dimensional fields while the heavy output records more 3D fields. We also have differente simulation duration settings: 1 day, 5 days, 15 days, and 31 days.

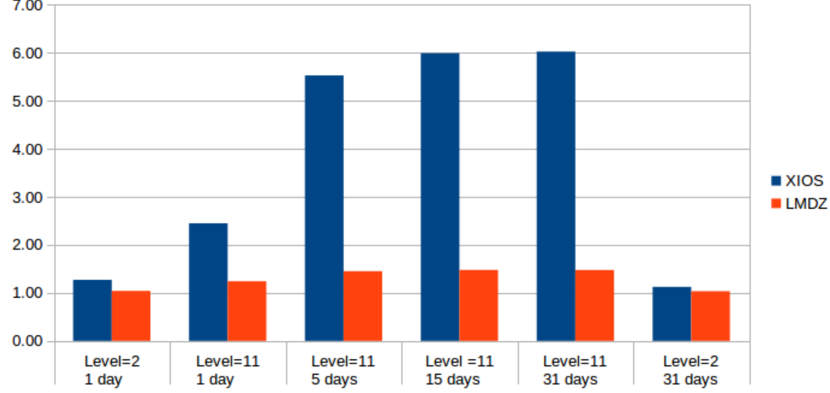

Figure 2: Speedup obtained by using EP in LMDZ simulations.

In this figure, we show the speedup which is computed by $\frac{time_{XIOS}}{time_{EP_XIOS}}$. The blue bars represent speedup of the XIOS file output and the red bars the speedup of LMDZ: calculates + XIOS file output. In all experimens, we can observe a speedup which represents a gain in performance. One important conclusion we can get from this result is that, more dense the output is, more efficient is the EP_XIOS. With 8 threads per process, we can reach a speedup in XIOS upto 6, and a speedup of 1.5 in LMDZ which represents a decrease of the total execution time to 68% ($\approx 1/1.5$). This observation confirmes steadily the importance of using EP in XIOS.

The reason why LMDZ does not show much speedup, is because the model is calcutation dominant: time spent on calculation is much longer than that on the file output. For example, if 30% of the execution time is spent on the output, then with a speepup of 6, we can obtain a decrease in time of 25%. Even the 25% may seems to be small, it is still a gain in performance with existing computing resources.

4 Perspectives of EP_XIOS