

Interactive comment on “A fast input/output library for high resolution climate models” by X. Huang et al.

X. Huang et al.

hxm@tsinghua.edu.cn

Received and published: 9 December 2013

Dear Thomas:

Thanks for your question.

The major reason that we did not compare with PIO is because we are targeting at high-resolution climate models. For the standalone POP and CICE models with NetCDF library that we downloaded from their official website (<http://climate.lanl.gov/Models/>), a resolution of 0.1 degree is directly supported. As far as we know, the POP and CICE models using PIO only exist in the coupled CESM, which provides no direct support for standalone tests with a high resolution of 0.1 degree. We think that porting the entire CESM coupled model from PIO to CFIO involves too much engineering work, and is

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper



beyond the scope of this article.

As J. Dennis has mentioned in his PIO paper (Dennis et al. 2012), the parallel access in PIO calls the functions in the PnetCDF library v1.2.0 directly, while adding the flexibility to reduce the memory usage through data rearrangement. The parallel write bandwidth reported in the PIO paper is around 925 MB/s on 64 cores.

Similarly, our CFIO implementation is also utilizing PnetCDF library to perform the actual parallel I/O operations, while providing automated overlapping of I/O and computing to further reduce the running time. During the last few days, we retested PnetCDF on our supercomputer and the write bandwidth of PnetCDF is around 1130 MB/s on 64 cores. The results are in a similar range and we think that the main difference comes from machine specifications and configurations.

At last, we want to say that our ideas proposed in CFIO shall be complements rather than replacements of existing parallel I/O libraries. One of the main goals of PIO and PnetCDF is to improve the bandwidth as high as possible. The main idea of CFIO is to overlap I/O with computing. We can integrate PIO or PnetCDF into the server side of CFIO to achieve a better performance. Therefore, in our experiments, we are mainly comparing CFIO against the “NO-I/O” case, which means turning off all I/O operations in the experiments. As we mentioned in Sec. 5.4, the bandwidth of CFIO is approximately 10% less than that of PnetCDF because of the overhead that is associated with I/O forwarding operations, but the running time is significantly reduced in real scenarios.

We hope our reply gives you some needed answers.

Best wishes, Xiaomeng

Interactive comment on Geosci. Model Dev. Discuss., 6, 4775, 2013.

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

