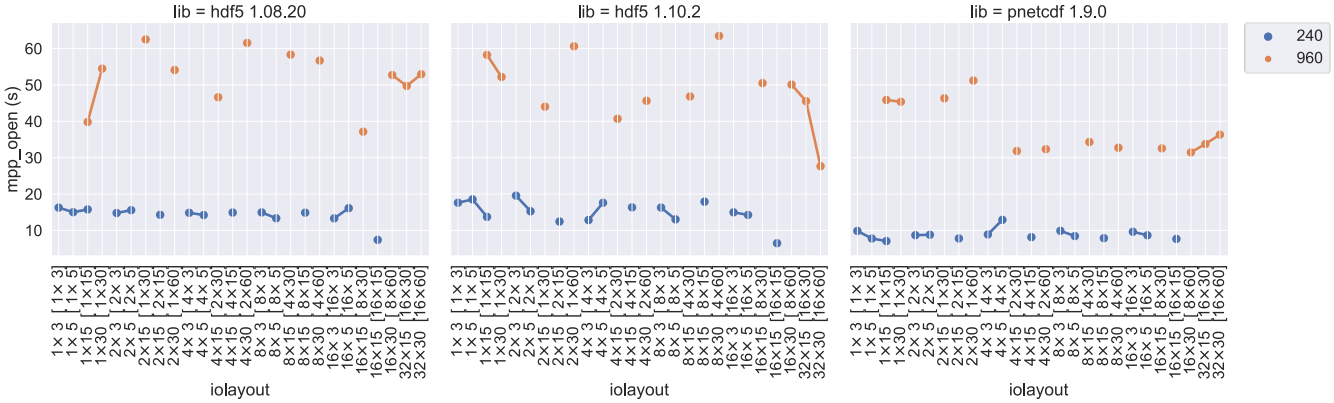

Figure 3. `mpp_open` time (in sec.) versus I/O layout in different libraries and PE numbers. HDF5 times are generally larger than in PnetCDF, and the runtime increases as PEs increase from 240 to 960. The I/O layout together with its PE distribution in [PE per node × nodes] are labelled in X-axis.

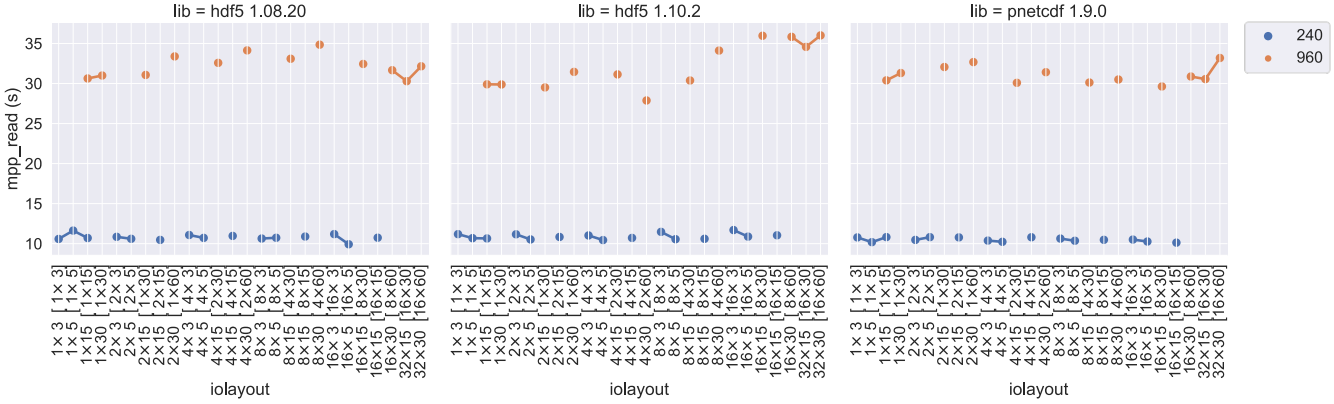

Figure 4. `mpp_read` time (in sec.) versus I/O layout in different libraries and PE numbers. Read operations do not use I/O layout or parallel I/O, and runtimes are largely independent of layout and library. Read times increase significantly as the number of PEs is increased. The I/O layout together with its PE distribution in [PE per node × nodes] are labelled in X-axis.

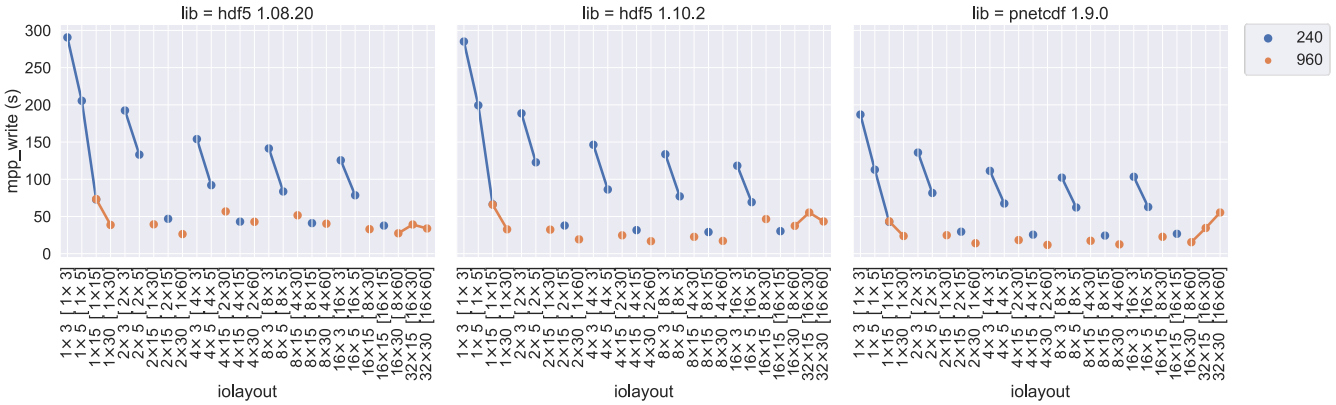

Figure 5. `mpp_write` time versus I/O layout for different library and PE numbers. Write time improves greatly as I/O nodes are increased (grouped curves), and modestly as the I/O PEs per node are increased (across grouped curves). Runtimes are scalarly reduced as PEs are increased. PnetCDF shows modest improvement over HDF5 performance. The I/O layout together with its PE distribution in [PE per node × nodes] are labelled in X-axis.

Figure 6. `mpp_close` time versus I/O layout with different libraries and PE numbers. Contentions within the HDF5 library lead to performance problems, which increase with layout and number of PEs. PnetCDF does not exhibit these issues and close times are negligible. The I/O layout together with its PE distribution in [PE per node × nodes] are labelled in X-axis.

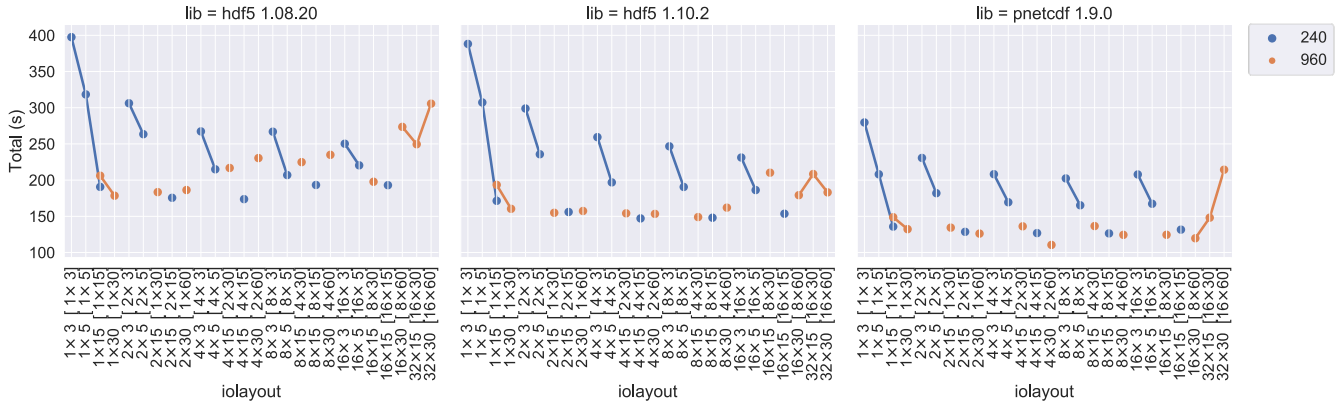

Figure 7. Total elapsed time versus I/O layout for different libraries and PE numbers. Higher contention at 960 PEs can overwhelm the overall performance trends observed at 240 PEs. The I/O layout together with its PE distribution in [PE per node × nodes] are labelled in X-axis.

Table 6. The time metrics of 0.1° model in 720-PE and 1440-PE runs with less I/O frequencies, i.e. write per 1 day and 4 days in 8-day simulations. SIO represents the original single threaded write; PIO represents parallel shared write. The I/O time composes of contributions from `mpp_open`, `mpp_read`, `mpp_write` and `mpp_close`. The I/O time ratio is given between the I/O time and total runtime.

I/O pattern&Format		SIO in netCDF4_classic		PIO in netCDF-4		PIO in netCDF-3	
I/O frequency		1-day	4-day	1-day	4-day	1-day	4-day
720 PEs	Total runtime (sec.)	8114	7817	7685	7569	7666	7469
	I/O time / `mpp_write` (sec.)	494/453	302/265	75/40	62/27	57/17	49/11
	I/O ratio	6.09%	3.87%	0.98%	0.82%	0.74%	0.66%
1440 PEs	Total runtime (sec.)	4118	3743	3547	3578	3518	3549
	I/O time / `mpp_write` (sec.)	452/421	269/238	59/24	48/14	51/14	40/7
	I/O time ratio	10.98%	7.18%	1.67%	1.35%	1.45%	1.14%