

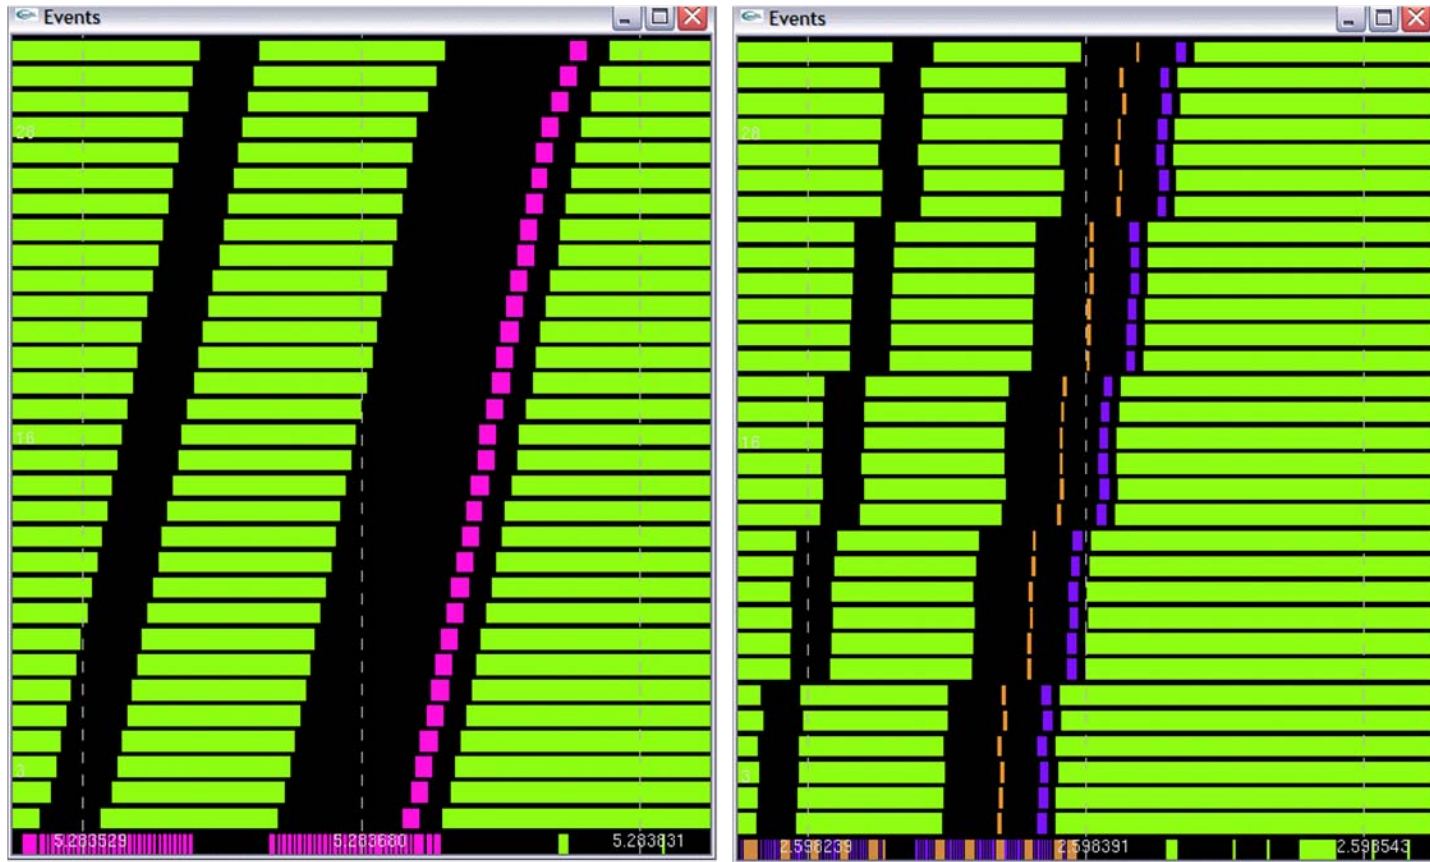


Real example of parallel scaling to large numbers of CPUs




Brief notes of real examples
Nick Allsopp, IBM UK.

Single node serialisation



Using master to groups of workers. Time is along the x-axis and each row is a different MPI task. The master node is the bottom row


Example of domain decomposition with single master I/O




Swapping halo cells

Single node file output – history file

ROMS code example used at UCT



April



August

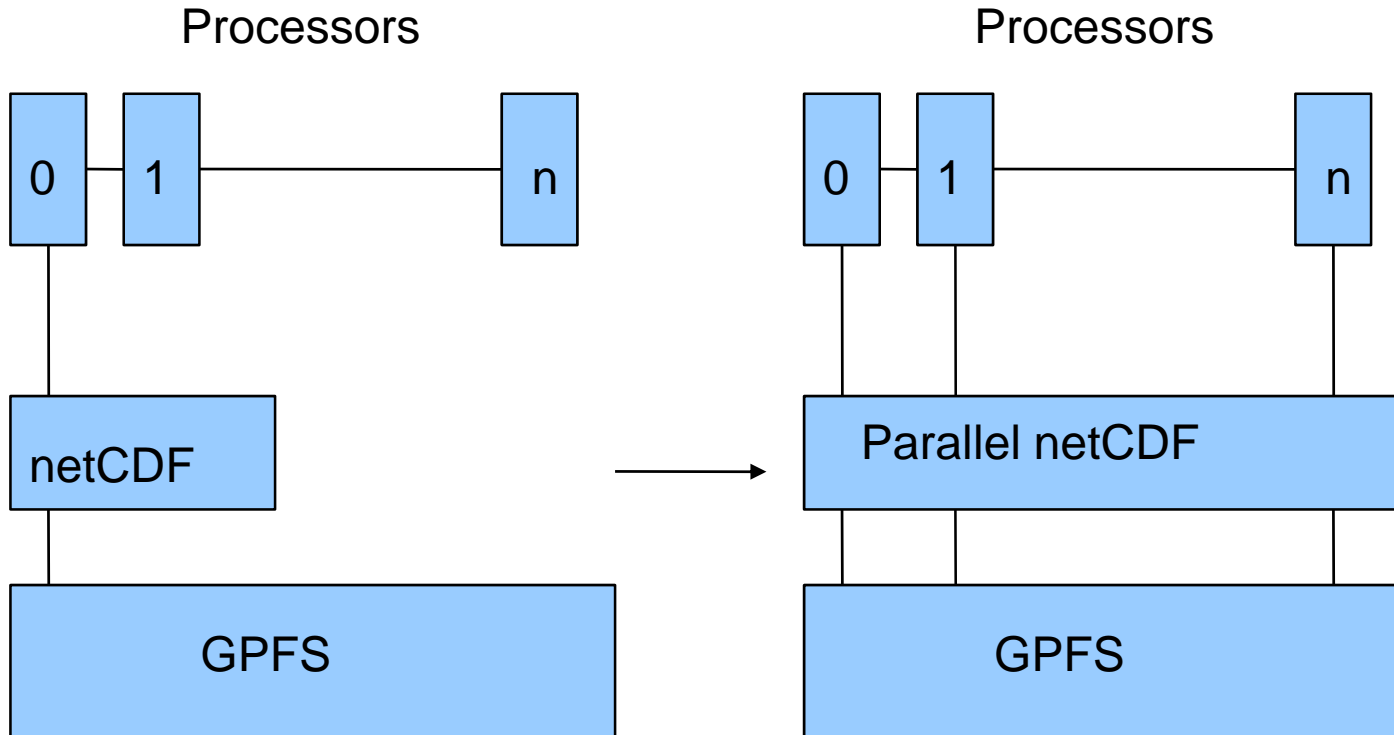
ROMS scaling on Blue Gene / P

BENCHMARK3 2048x265x30

TimeSteps: 00000001 - 00001080) which is 1080 steps and 5 history files.

#cpus	time	I/O%	2D Kernel%	Halo MPI%	total MPI%
64	4154	1.8	40	11	15
128	2272	3.2	41	16	22
256	1191	6.2	41	17	25
512	726	10.0	43	25	37
1024	828	37.9	42	25	45
2048	1246	55.6	42	36	63

ROMS file I/O for History files



Scaling using serial and parallel I/O for History file

Scaling on BG/P

Test Case: benchmark3

