Performance results of the hybrid MPI/OpenMP MrBayes version for a DNA data set with 125 taxa and 19,436 patterns that was run with nruns=2 and nchains=4 for 10,000 generations conducted by Wayne Pfeiffer (SDSC). Note that, with 2 runs executing 4 chains each, the maximum achievable degree of parallelism for the pure MPI versions is 8, i.e., not more than 8 cores can be used in parallel.
Those are evidently not nearly enough generations for convergence, but they are sufficient for timing purposes and for measuring parallel efficiency.
The system used is a cluster with two Intel Clovertown processors per node, i.e., 8 cores per node.

Using 8 MPI processes and 8 OpenMP threads (the maximum possible) leads to a speedup of 5.67 on 8 nodes compared to using MPI alone on a single node, i.e., the response time is much higher :-)

Moreover, on a single node a mix of MPI processes and OpenMP threads is preferable to 8 MPI processes alone or 8 OpenMP threads alone.
In particular, a mix of 2 MPI processes and 4 OpenMP threads per node is 1.38x faster than running 8 MPI processes alone on a single node. This should be typical for multi-gene data sets that have a large number of patterns. The optimum speedups achieved by configurations using 1 node (8 cores), 2 nodes (16 cores), and 4 nodes (32 cores) are marked by yellow color.

Nodes used	Cores used	#MPI processes	#OpenMP threads	MCMC time seconds	speedup
1	8	8	1	1,940	1.00
1	8	4	2	1,675	1.14
1	8	2	4	1,409	1.38
1	8	1	8	1,595	1.22
2	16	8	2	1,017	1.91
2	16	4	4	889	2.18
2	16	2	8	808	2.40
4	32	8	4	546	3.55
4	32	4	8	545	3.56
8	64	8	8	342	5.67