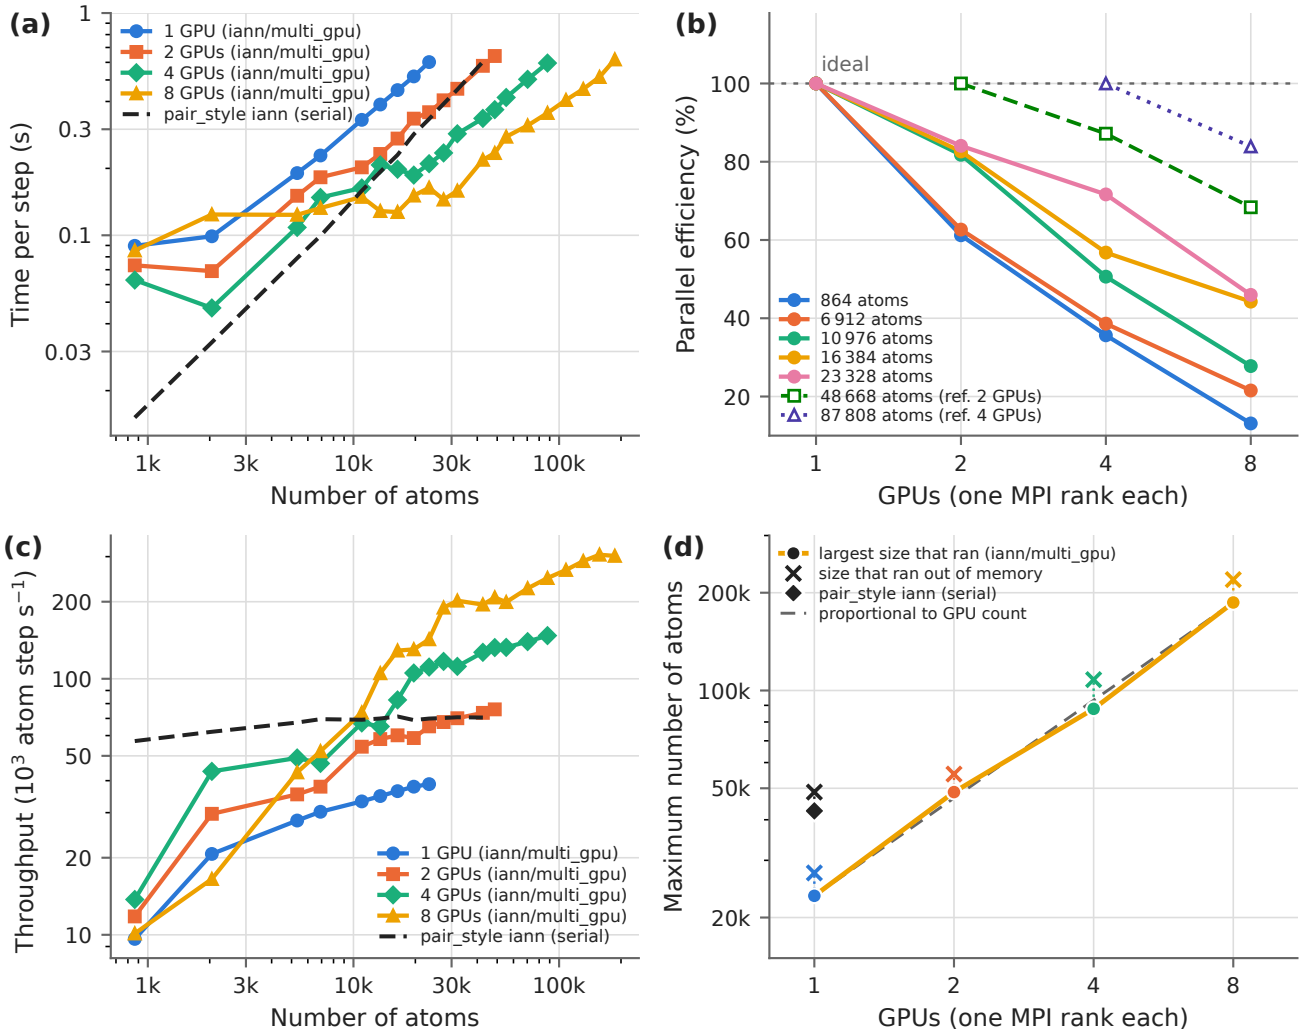

(a)
1 GPU (iann/multi_gpu)
2 GPUs (iann/multi_gpu)
4 GPUs (iann/multi_gpu)
8 GPUs (iann/multi_gpu)
pair_style iann (serial)
Time per step (s)
1
0.3
0.1
0.03
1k
3k
10k
30k
100k
Number of atoms
(b)
ideal
Parallel efficiency (%)
100
80
60
40
20
864 atoms
6 912 atoms
10 976 atoms
16 384 atoms
23 328 atoms
48 668 atoms (ref. 2 GPUs)
87 808 atoms (ref. 4 GPUs)
1
2
4
8
GPUs (one MPI rank each)
(c)
Throughput (10^3 atom step s^{-1})
200
100
50
20
10
1 GPU (iann/multi_gpu)
2 GPUs (iann/multi_gpu)
4 GPUs (iann/multi_gpu)
8 GPUs (iann/multi_gpu)
pair_style iann (serial)
1k
3k
10k
30k
100k
Number of atoms
(d)
largest size that ran (iann/multi_gpu)
size that ran out of memory
pair_style iann (serial)
proportional to GPU count
Maximum number of atoms
200k
100k
50k
20k
1
2
4
8
GPUs (one MPI rank each)