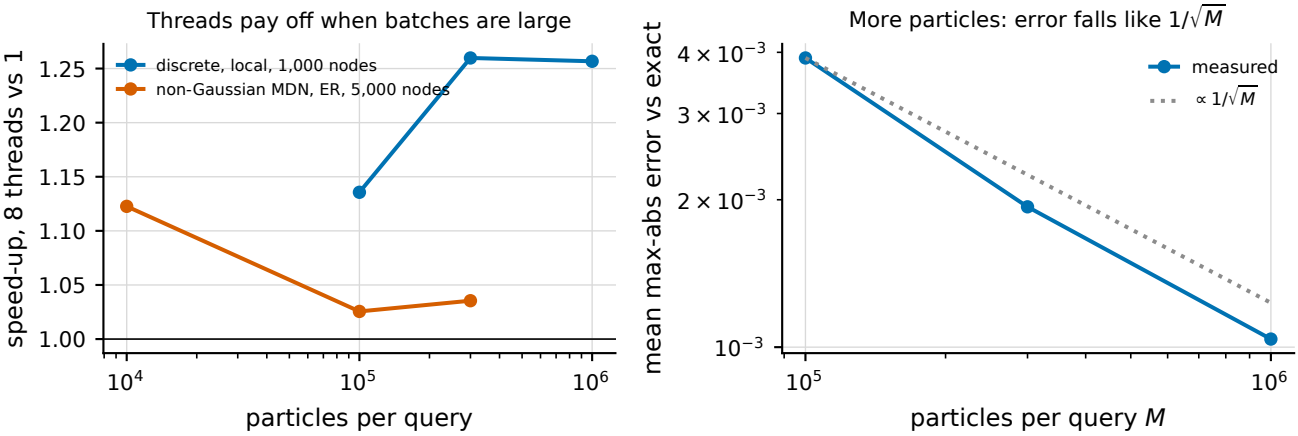

Threads pay off when batches are large
discrete, local, 1,000 nodes
non-Gaussian MDN, ER, 5,000 nodes
speed-up, 8 threads vs 1
1.25
1.20
1.15
1.10
1.05
1.00
10^4
10^5
10^6
particles per query
More particles: error falls like $1/\sqrt{M}$
measured
$\propto 1/\sqrt{M}$
mean max-abs error vs exact
4×10^{-3}
3×10^{-3}
2×10^{-3}
10^{-3}
10^5
10^6
particles per query M