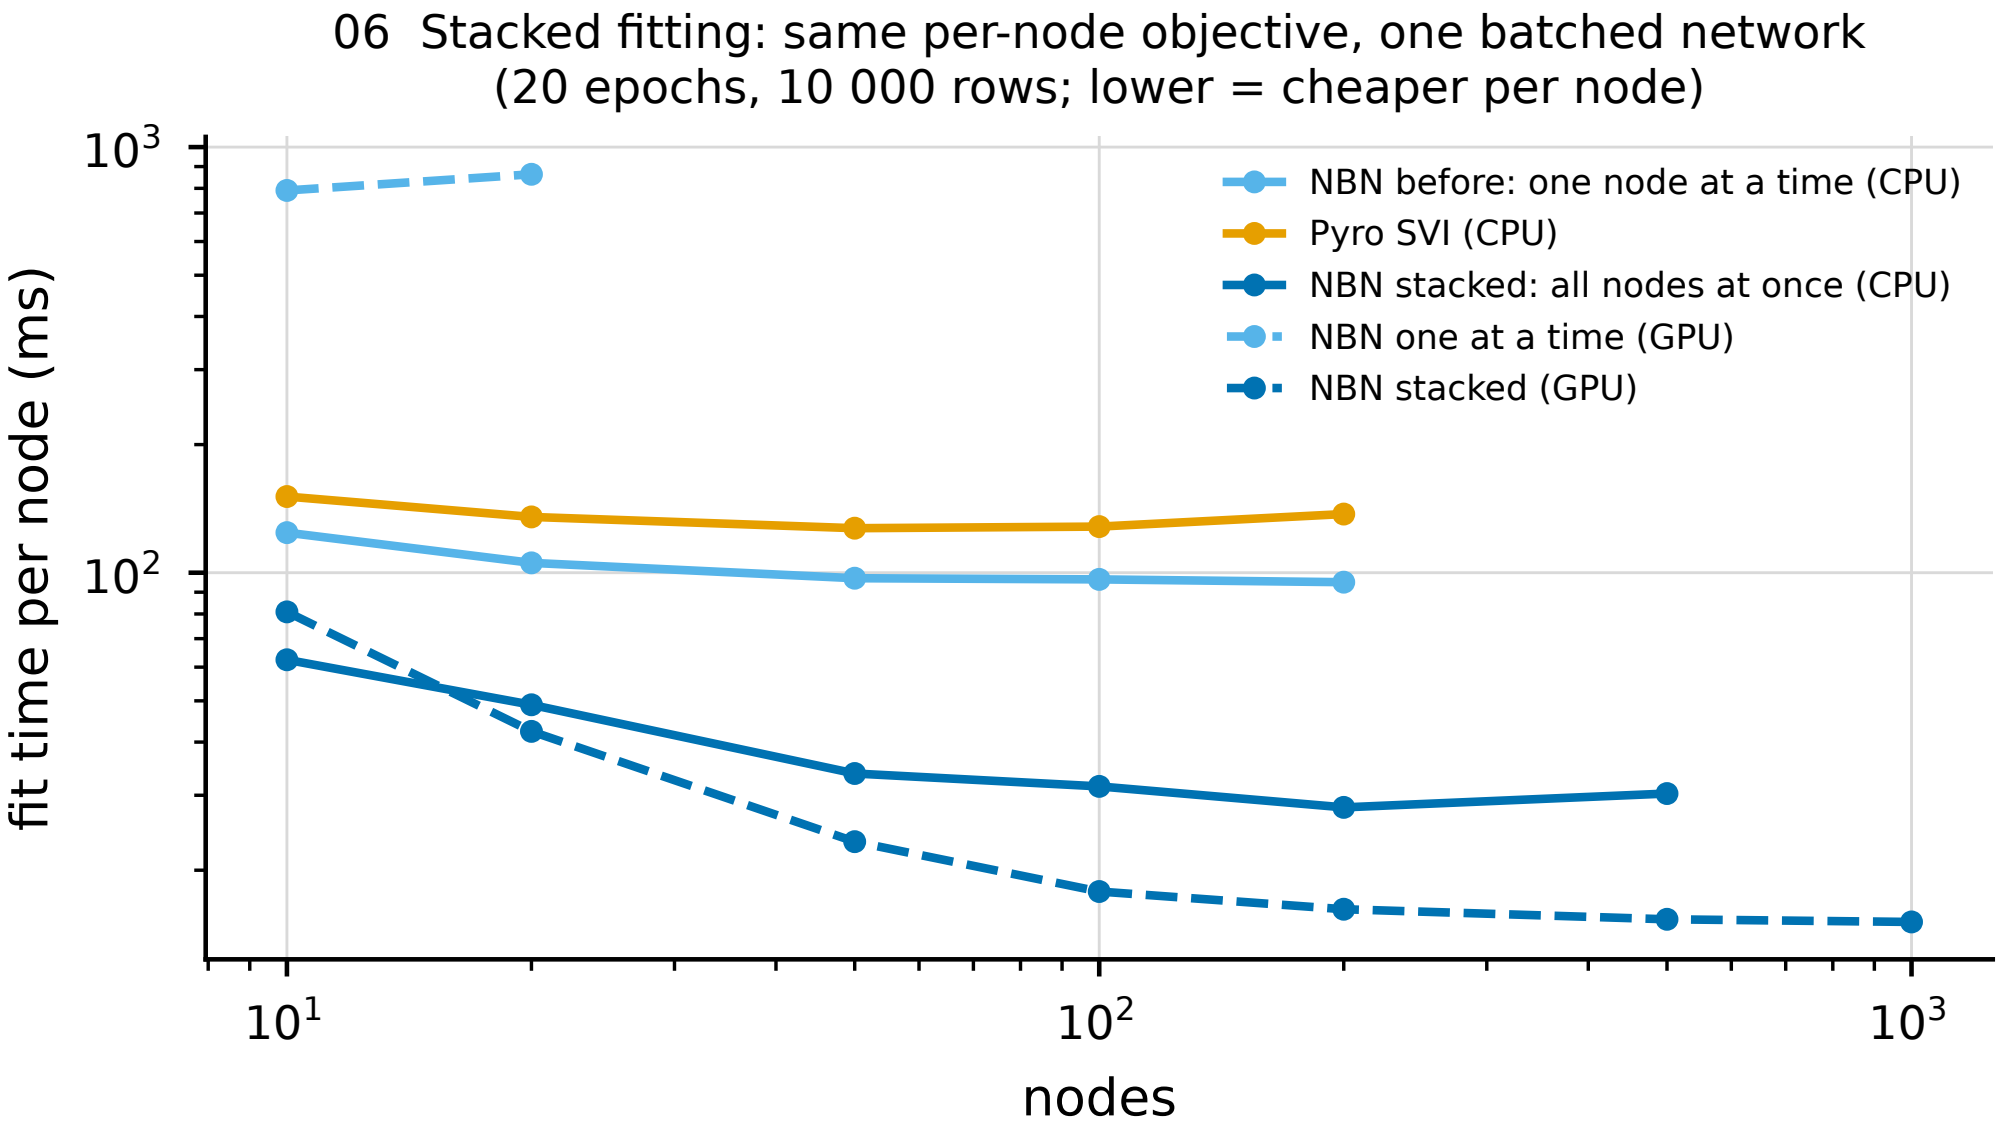

06 Stacked fitting: same per-node objective, one batched network
(20 epochs, 10 000 rows; lower = cheaper per node)
fit time per node (ms)
nodes
10^3
10^2
10^1
10^2
10^3
NBN before: one node at a time (CPU)
Pyro SVI (CPU)
NBN stacked: all nodes at once (CPU)
NBN one at a time (GPU)
NBN stacked (GPU)