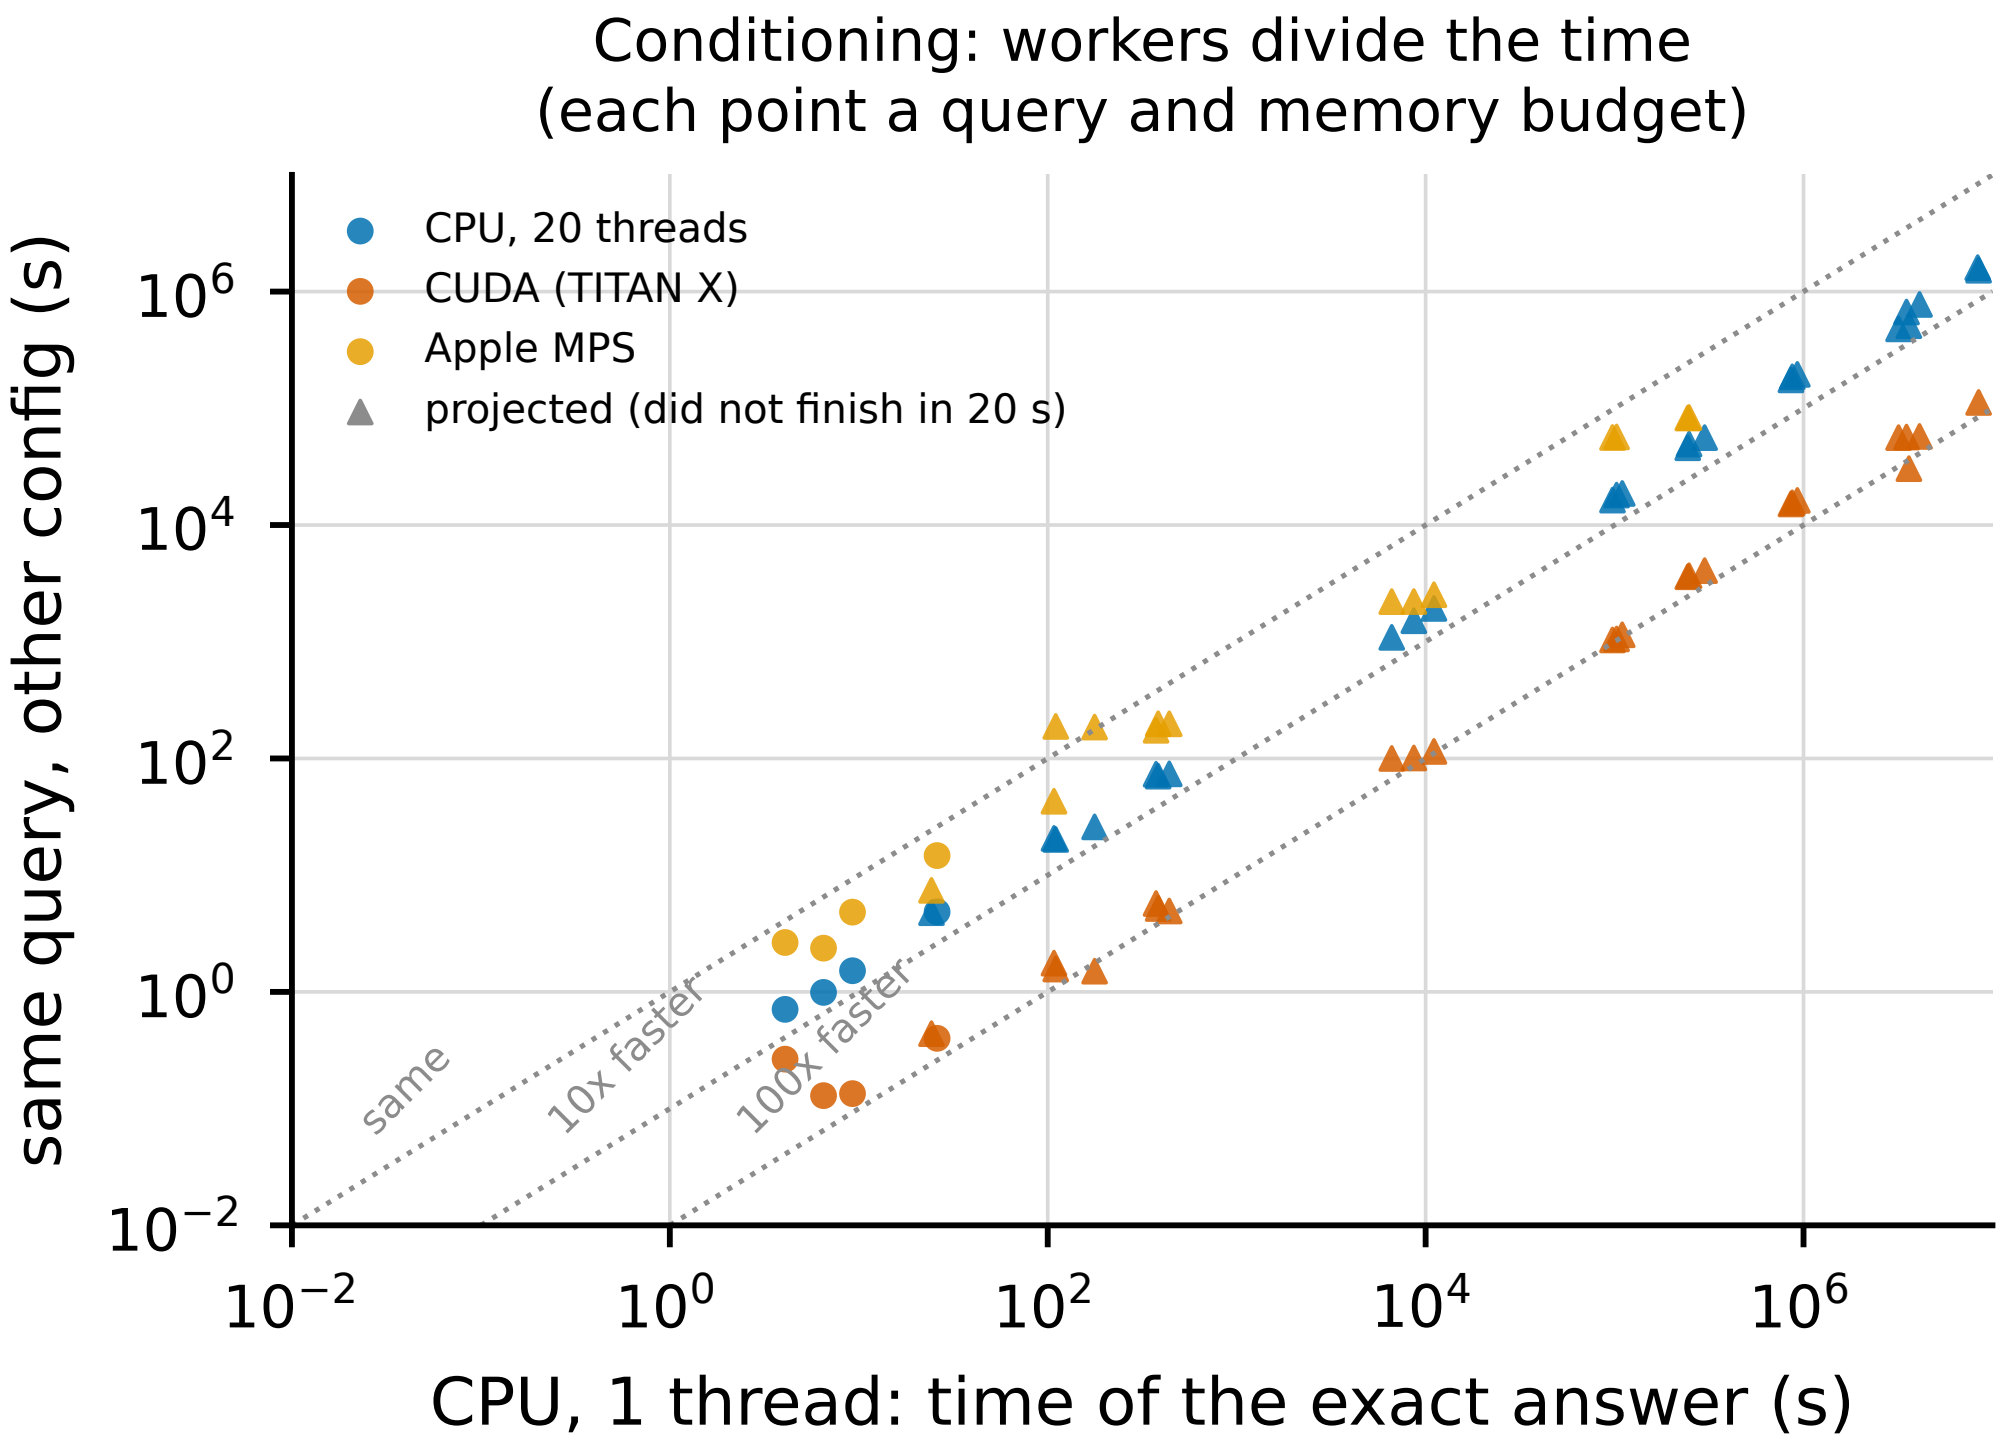

Conditioning: workers divide the time
(each point a query and memory budget)
CPU, 20 threads
CUDA (TITAN X)
Apple MPS
projected (did not finish in 20 s)
same
10x faster
100x faster
same query, other config (s)
CPU, 1 thread: time of the exact answer (s)
10^{-2}
10^{0}
10^{2}
10^{4}
10^{6}