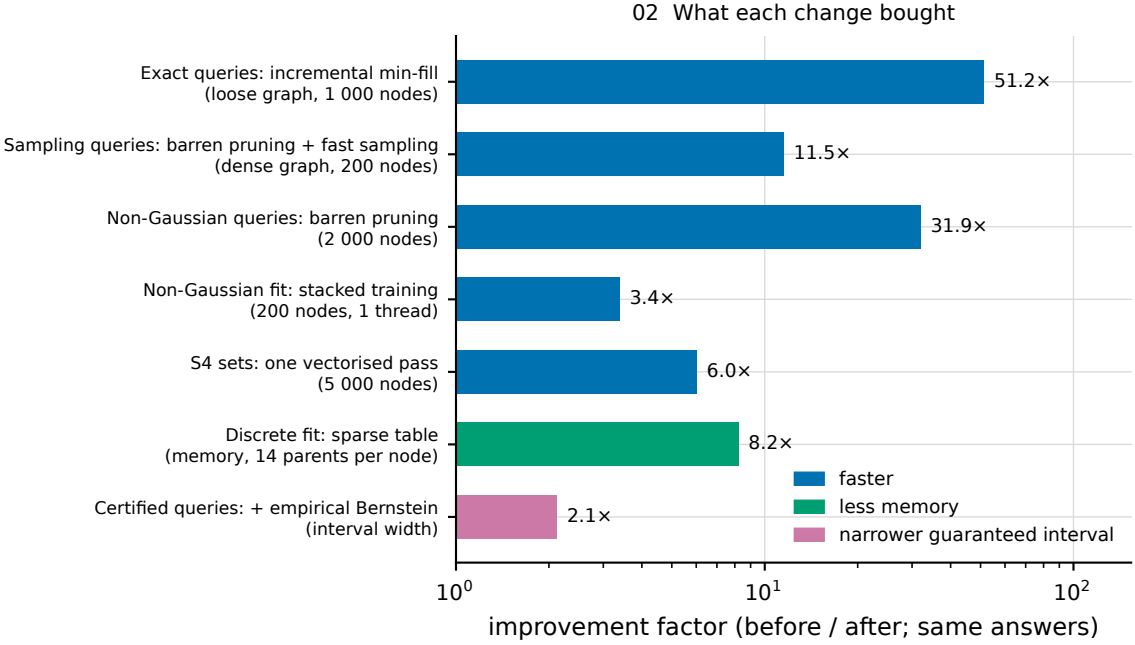

02 What each change bought
Exact queries: incremental min-fill
(loose graph, 1 000 nodes)
51.2×
Sampling queries: barren pruning + fast sampling
(dense graph, 200 nodes)
11.5×
Non-Gaussian queries: barren pruning
(2 000 nodes)
31.9×
Non-Gaussian fit: stacked training
(200 nodes, 1 thread)
3.4×
S4 sets: one vectorised pass
(5 000 nodes)
6.0×
Discrete fit: sparse table
(memory, 14 parents per node)
8.2×
Certified queries: + empirical Bernstein
(interval width)
2.1×
faster
less memory
narrower guaranteed interval
10^0
10^1
10^2
improvement factor (before / after; same answers)