Appendix for Parallel Beam Search Algorithms for Domain-Independent Dynamic Programming*

Ryo Kuroiwa, J. Christopher Beck

Department of Mechanical and Industrial Engineering, University of Toronto, Toronto, Canada, ON M5S 3G8
ryo.kuroiwa@mail.utoronto.ca, jcb@mie.utoronto.ca

Experimental Results

We show scatter plots comparing the time to optimally solve a model by 1 and 32 thread solvers in Figures 1-6. The speedup of the parallel DIDP solvers becomes larger as sequential CABS requires more time to solve an instance. The parallel solvers take at least 0.01 seconds to solve an instance even if sequential CABS solves it faster. This overhead possibly comes from the initialization of multiple threads. In SALBP-1, the speedup achieved by the parallel solvers is varied and sometimes super-linear, which is possibly caused by the search behavior change. The number of expansions, shown in Figures 7-12, supports this explanation. While the parallel solvers have almost the same number of expansions as sequential CABS in all instances of TSPTW, CVRP, MOSP, and graph-clear, the expansion ratio ranges from less than 0.1 to more than 10 in SALBP-1. The tendency is more evident in bin packing, and some instances solved by sequential CABS are unsolved by the parallel solvers.

We also show scatter plots comparing time for MIP and CP. In TSPTW and MOSP, parallel CP is slower than sequential CP while it has the better primal gap, as discussed in the paper.

References

Kuroiwa, R.; and Beck, J. C. 2024. Parallel Beam Search Algorithms for Domain-Independent Dynamic Programming. In *Proceedings of the 38th Annual AAAI Conference on Artificial Intelligence (AAAI)*.

*This document is an appendix for a paper published at the 38th Annual AAAI conference on Artificial Intelligence (AAAI-24) (Kuroiwa and Beck 2024).

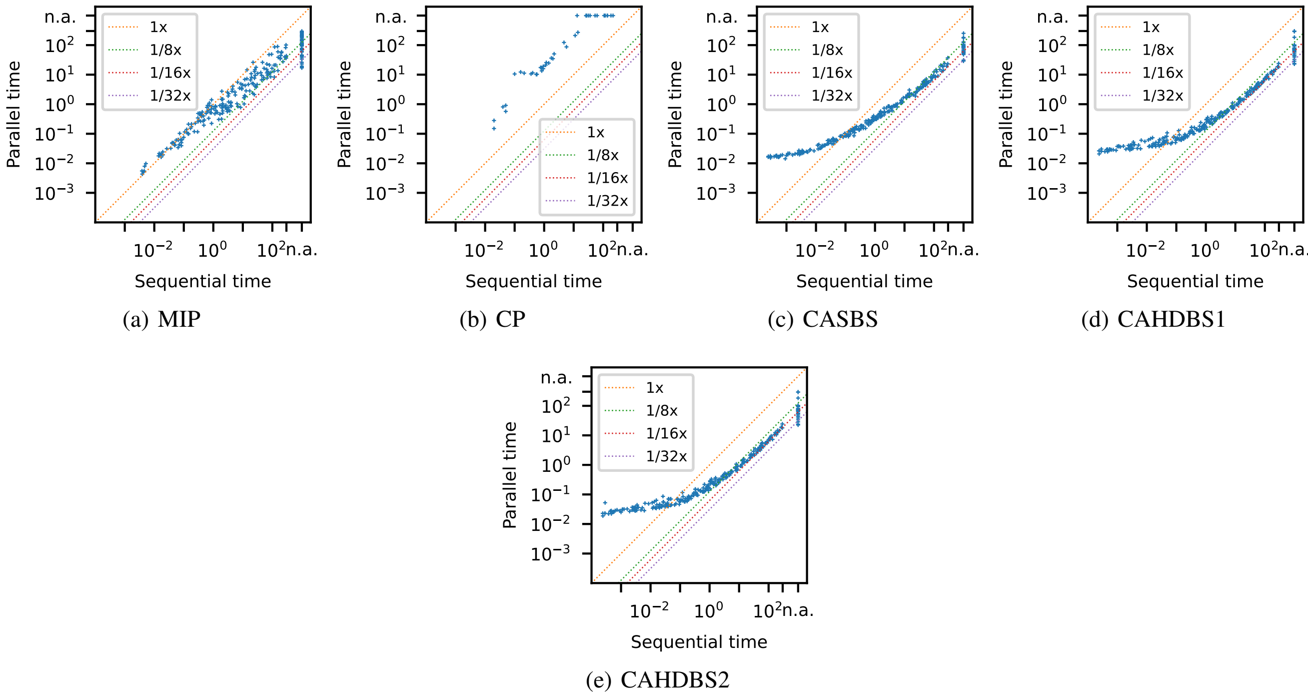

(a) MIP (b) CP (c) CASBS (d) CAHDBS1

(e) CAHDBS2

Figure 1: Comparison of the time to optimally solve a model by 1 and 32 thread solvers in TSPTW. Unsolved instances are shown at ‘n.a.’.

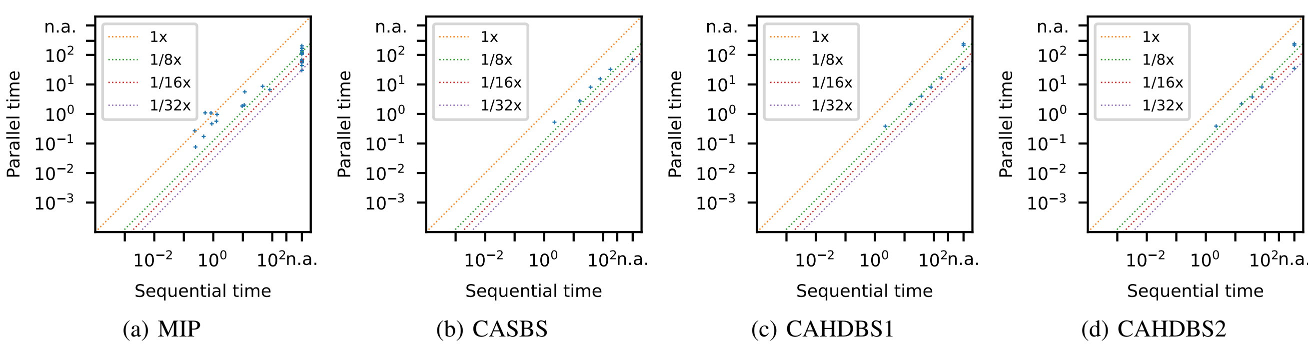

(a) MIP (b) CASBS (c) CAHDBS1 (d) CAHDBS2

Figure 2: Comparison of the time to optimally solve a model by 1 and 32 thread solvers in CVRP. Unsolved instances are shown at ‘n.a.’. CP is not shown as it does not solve any instance optimally.

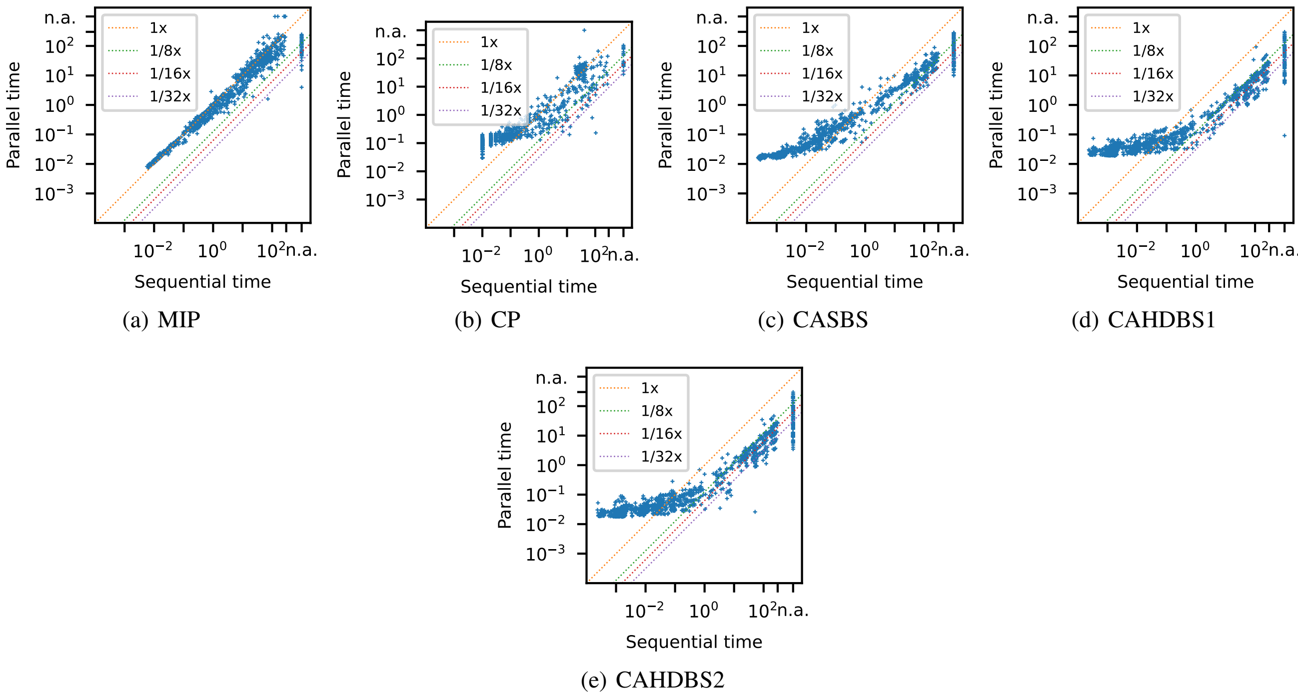

Figure 3: Comparison of the time to optimally solve a model by 1 and 32 thread solvers in SALBP-1. Unsolved instances are shown at 'n.a.'.

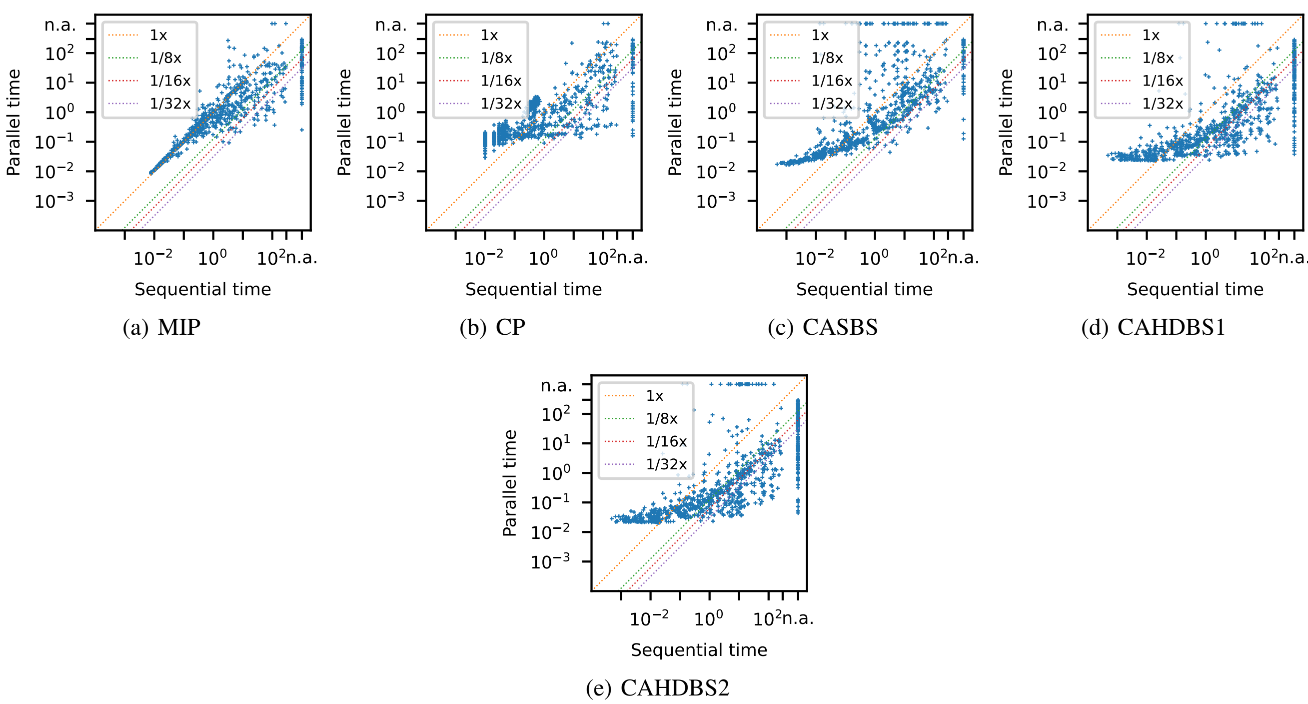

Figure 4: Comparison of the time to optimally solve a model by 1 and 32 thread solvers in bin packing. Unsolved instances are shown at 'n.a.'.

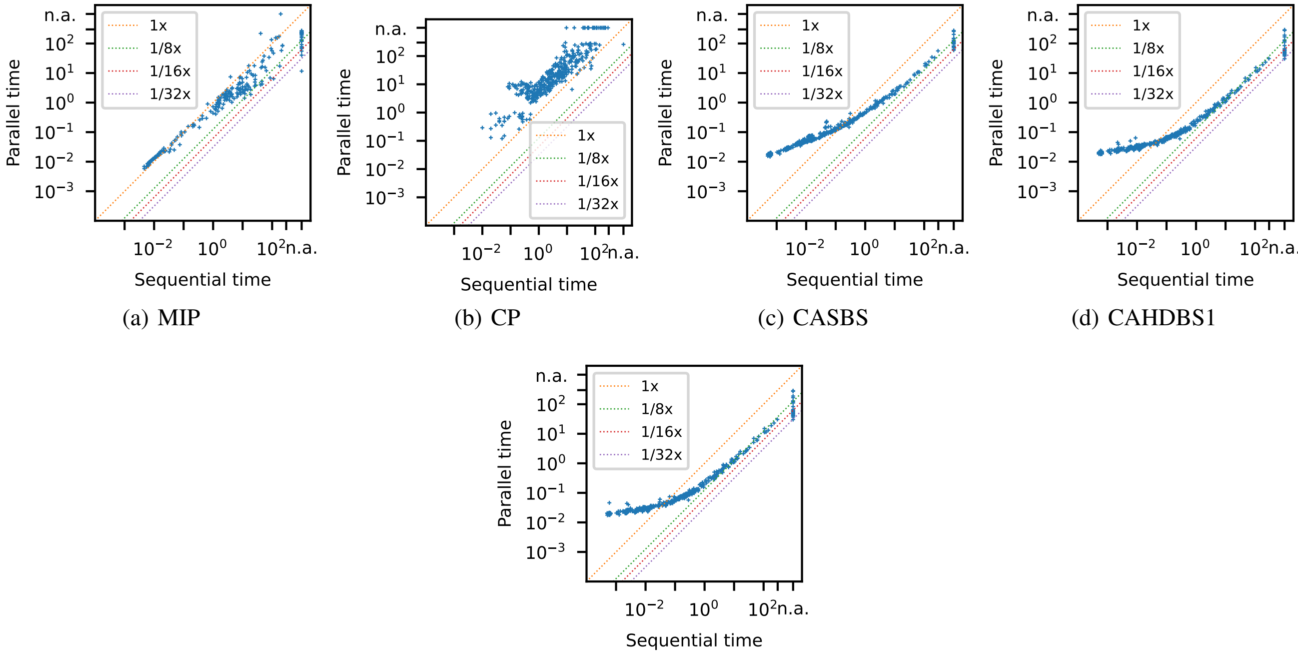

(a) MIP (b) CP (c) CASBS (d) CAHDBS1

(e) CAHDBS2

Figure 5: Comparison of the time to optimally solve a model by 1 and 32 thread solvers in MOSP. Unsolved instances are shown at ‘n.a.’.

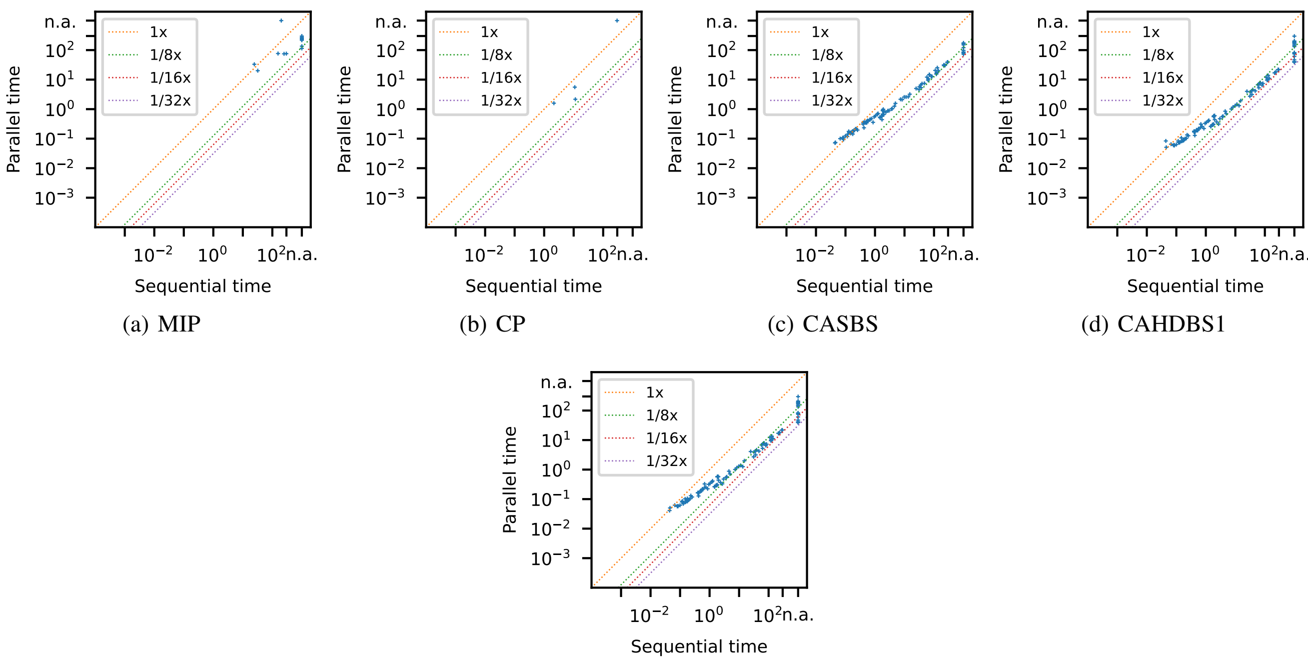

(a) MIP (b) CP (c) CASBS (d) CAHDBS1

(e) CAHDBS2

Figure 6: Comparison of the time to optimally solve a model by 1 and 32 thread solvers in graph-clear. Unsolved instances are shown at ‘n.a.’.

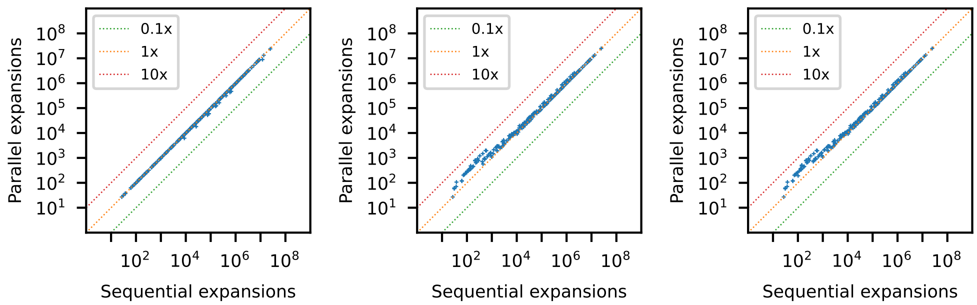

(a) CASBS (b) CAHDBS1 (c) CAHDBS2

Figure 7: Comparison of the numbers of expansions to optimally solve a model by 1 and 32 thread solvers in TSPTW.

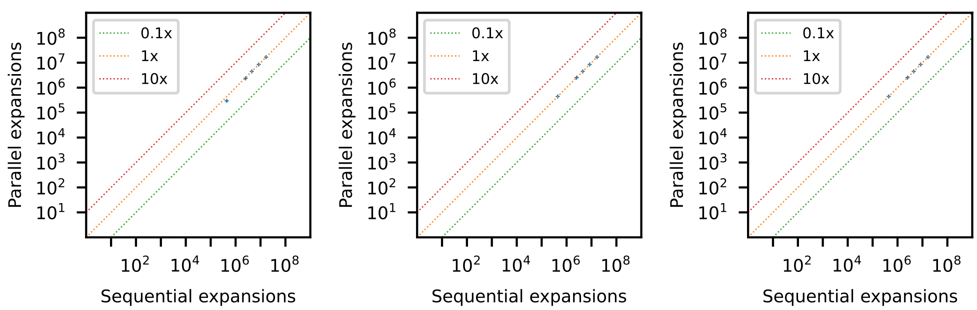

(a) CASBS (b) CAHDBS1 (c) CAHDBS2

Figure 8: Comparison of the numbers of expansions to optimally solve a model by 1 and 32 thread solvers in CVRP.

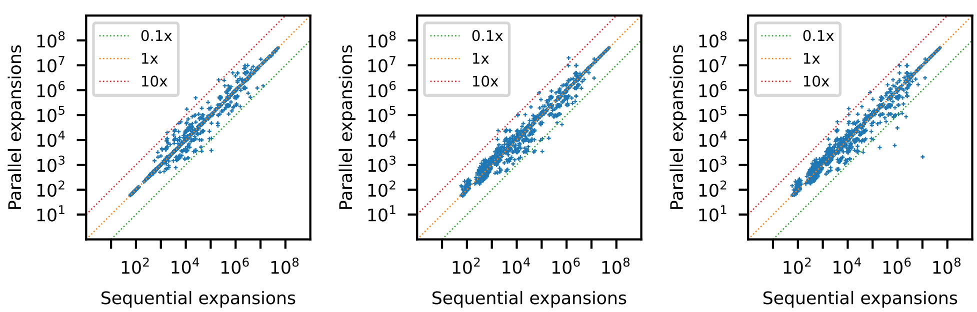

(a) CASBS (b) CAHDBS1 (c) CAHDBS2

Figure 9: Comparison of the numbers of expansions to optimally solve a model by 1 and 32 thread solvers in SALBP-1.

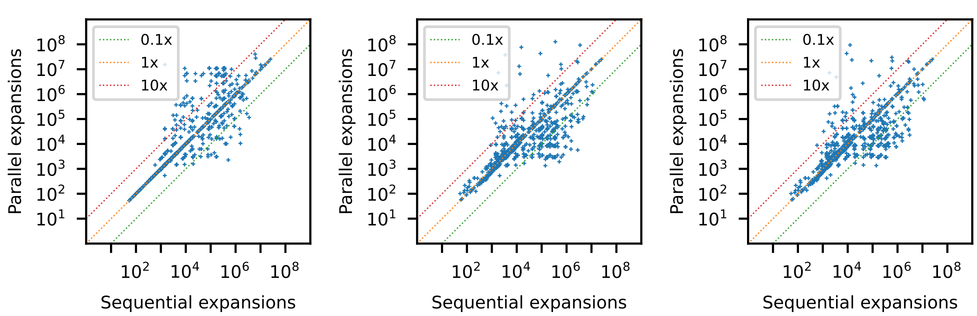

(a) CASBS (b) CAHDBS1 (c) CAHDBS2

Figure 10: Comparison of the numbers of expansions to optimally solve a model by 1 and 32 thread solvers in bin packing.

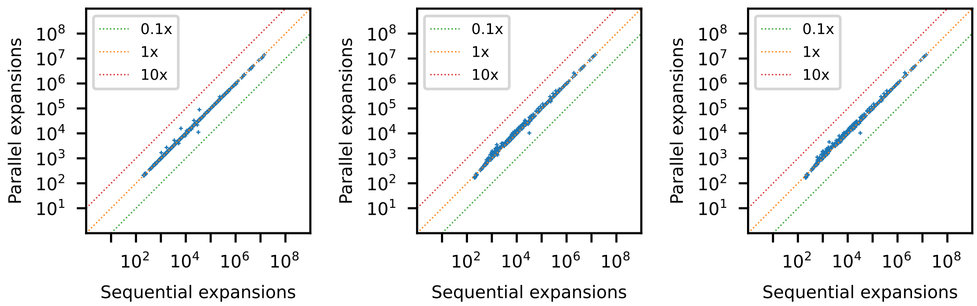

(a) CASBS (b) CAHDBS1 (c) CAHDBS2

Figure 11: Comparison of the numbers of expansions to optimally solve a model by 1 and 32 thread solvers in MOSP.

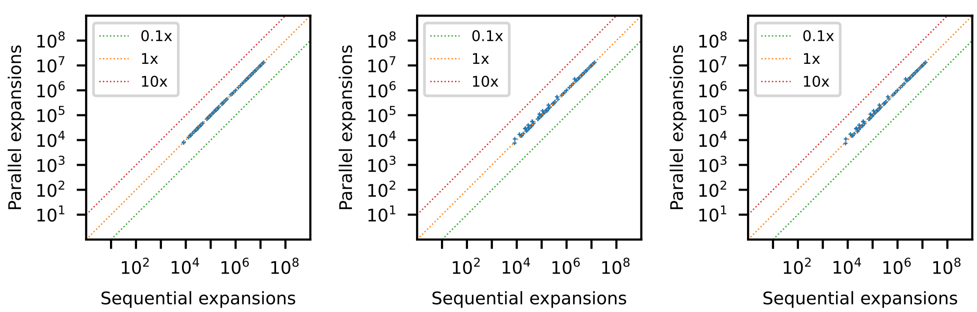

(a) CASBS (b) CAHDBS1 (c) CAHDBS2

Figure 12: Comparison of the numbers of expansions to optimally solve a model by 1 and 32 thread solvers in graph-clear.