

**Exercise.** This is the solution to exercise 2.17 in the book.

**Solution.** Consider a random search. When using a single processor a certain amount of search space is covered which may delay finding a solution. However, when using more processors, each can explore more of the search space hence a solution is found faster. In this case it is the non-linear nature of the search algorithm that leads to *superlinear* speedup.

Another example could be a program where parallel computations run over Ethernet. If a limited number of nodes are used, the inter-node communication may be limited by Ethernet bandwidth, adding more nodes could reduce the amount of inter-node communication, hence Ethernet bandwidth would not be a bottle-neck.