

1. What criticism might a computational scientist level against using message passing?
2. Give four reasons why a scientist may want to parallelize a program.
3. Give an example of a *data dependency*.
4. What does MIMD stand for?
5. What is difference between “coarse grain” and “fine grain” parallel computations?
6. Which two variables are related by Amdahl’s law?
7. Why is “latency” of importance for parallel computations?
8. What is meant by a embarrassingly parallel task?