High-Performance Computing

– Exercises: Foundations of Parallel Hardware and Parallel Software –

Christian Plessl

High-Performance IT Systems Group

Paderborn University

UNIVERSITÄT PADERBORN
Die Universität der Informationsgesellschaft

version 1.1.0 2017-10-13

Paderborn Center for Parallel Computing

Pacheco Ex 2.3 + Ex 2.5

2.3. Recall the example involving cache reads of a two-dimensional array (page 22). How does a larger matrix and a larger cache affect the performance of the two pairs of nested loops? What happens if `MAX = 8` and the cache can store four lines? How many misses occur in the reads of `A` in the first pair of nested loops? How many misses occur in the second pair?

2.5. Does the addition of cache and virtual memory to a von Neumann system change its designation as an SISD system? What about the addition of pipelining? Multiple issue? Hardware multithreading?

Pacheco Ex 2.10

2.10. Suppose a program must execute 10^{12} instructions in order to solve a particular problem. Suppose further that a single processor system can solve the problem in 10^6 seconds (about 11.6 days). So, on average, the single processor system executes 10^6 or a million instructions per second. Now suppose that the program has been parallelized for execution on a distributed-memory system. Suppose also that if the parallel program uses p processors, each processor will execute $10^{12}/p$ instructions and each processor must send $10^9(p-1)$ messages. Finally, suppose that there is no additional overhead in executing the parallel program. That is, the program will complete after each processor has executed all of its instructions and sent all of its messages, and there won't be any delays due to things such as waiting for messages.

a. Suppose it takes 10^{-9} seconds to send a message. How long will it take the program to run with 1000 processors, if each processor is as fast as the single processor on which the serial program was run?

b. Suppose it takes 10^{-3} seconds to send a message. How long will it take the program to run with 1000 processors?

Pacheco Ex 2.13 + Ex 2.19

2.13. **a.** Sketch a four-dimensional hypercube.

b. Use the inductive definition of a hypercube to explain why the bisection width of a hypercube is $p/2$.

2.19. Suppose $T_{\text{serial}} = n$ and $T_{\text{parallel}} = n/p + \log_2(p)$, where times are in microseconds. If we increase p by a factor of k, find a formula for how much we'll need to increase n in order to maintain constant efficiency. How much should we increase n by if we double the number of processes from 8 to 16? Is the parallel program scalable?

Change log

- **1.1.0 (2017-10-13)**
  - updated for winter term 2017/18
- **1.0.1 (2016-11-09)**
  - add Ex 2.19
- **1.0.0 (2016-11-04)**
  - initial version of slides