Faculty EEMCS
Department DIAM

Examination IN4049TU

January 25, 2022, 18:30—21:30

You may use lecture slides/notes, but please write answers in your own words.

Questions 1-7: 12 point; Question 8: 16 points; Grade = (total number of points)/10

1. (a) Modern computers have a hierarchical memory system. What is the main reason of introducing such a memory system?

   (b) Given a parallel program with 4% of the code that cannot be parallelized. What is the maximum speedup if P=20 processors is used for the parallel execution of the program? And what is the maximum speedup when you have infinite many processors available for this program?

2. (a) If caches are considered as parts of the memory system, do you classify today's computers as UMA or as NUMA computers? Please explain.

   (b) Consider a hypercube network and a tree network, compare the advantages and disadvantages of these two networks in terms of performance and cost.

3. (a) We want to send the following data a[2**i*], for *i*=0, …, 100, from processor 1 to processor 2, using MPI_Send and MPI_Recv. Assuming that processor 2 does not know the number of data to be received at the time of writing the program or during compile time. Give an efficient algorithm to implement this communication.

   (b) Describe the main steps/phases of the Multi-Grid method. What are the main reasons for introducing the Multi-Grid method?

4. (a) ORB and tree-partitioning (Costzones) are two load balancing schemes for parallel NBody simulation using the Barnes-Hut algorithm. Describe these two methods. Which one gives a better load balance in general? Please explain.

   (b) Do you agree with the statement "Multi-level graph partitioning is designed to achieve a higher parallel speedup in partitioning a large graph"? Please explain your answer.

5. (a) Do you agree with the statement "SIMD and SIMT have the same meaning, they are used for shared-memory parallel computers and GPUs respectively"? Please explain your answer.

(b) What is the arithmetic intensity of the following piece of code? The variables **a**, **b**, **c** and **d** are of the type 4-byte float. **a** is an n-by-n matrix, and **b**, **c** and **d** are vectors of length n (assume n is large).

```
for (i=0; i<100; i++) {
    c[i]=c[i]+b[i]*2.5;
    for (j=0; j<100; j++)
        c[i]=c[i]+(a[i][j]*b[j]+a[i][j]*d[j])/2;
}
```

Given a processor with a peak performance of 2 GFLOPS, a memory bandwidth of 1.5GB/s, and a L1 cache of 16KB. Is the above code compute-bound or memory-bound?

6. (a) Consider a 2-D Poisson equation discretized with a 5-point central difference stencil into an $m \times n$ grid. For a Gauss-Seidel or SOR iteration, the figure below shows the areas of independent computation, starting from the top left corner with the white rectangular with point 1, followed by the yellow areas with points 2 and 7 can be computed in parallel, and so on. Is the maximum number of independent points equal to $\min\{m, n\}$? Can we have more than $\min\{m, n\}$ independent points at a certain step by combining different parts of multiple yellow and white diagonal areas? Please explain.

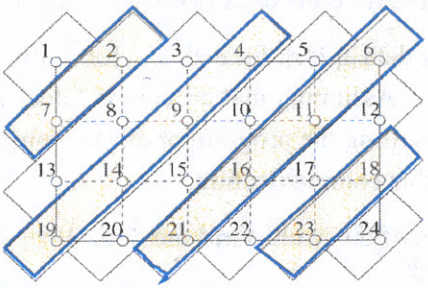

(b) Consider using Gauss-Seidel iterative method to solve a system of linear equations $Ax = b$, where the matrix A is a dense (full) matrix. Can we apply the red-black ordering for the parallelization of the Gauss-Seidel iteration? Please explain.

7. (a) What are the main differences between data parallelism and model parallelism in parallel training of neural networks?

(b) In parallel asynchronous SGD algorithm, asynchrony is introduced in a certain step or operation of the SGD algorithm. Describe this operation and the difference compared to the parallel SGD algorithm.

8. (a) In the Jacobi algorithm, inner products of the residual vector has to be computed to check whether the solution has converged. With "communication avoiding Jacobi" scheme, can we reduce the communication time in the inner product computation? Please explain.

   (b) We consider parallel Jacobi method for solving a problem with 1-dimensional grid of n points. If the 1-dimensional grid is partitioned into P parts each assigned to one of the P processors. What is the data locality of this parallel Jacobi algorithm? And what is the data locality when a k-step communication avoiding Jacobi scheme is implemented? i.e., communicate after every k=5 iterations.

   (c) From the results of (b), can we conclude that communication avoiding Jacobi is more efficient because it has a higher data locality than the conventional Jacobi algorithm? Why? Please explain.

- End -