Exam Parallel Programming 23 October 2008
Department of Computer Science, Faculty of Sciences

1. (a) How is the *efficiency* of a parallel program defined?
   (b) Why is it difficult to determine the efficiency if the program is executed on a computational grid?

2. Given below is a sequential algorithm that makes a fixed number of sweeps over an N-by-M array and during each sweep updates elements:

```
float G[1:N, 1:M], Gnew[1:N, 1:M];
for (step = 0; step < NSTEPS; step++) {
  for (i = 3; i < N-1; i++) {
    for (j = 3; j < M-1; j++) {
      Gnew[i,j] = (G[i,j] + G[i-1,j] + G[i-2,j] +  G[i+1,j] + G[i+2,j]
                   + G[i,j-1] + G[i,j-2] + G[i,j+1] + G[i, j+2]) / 9;
    }
  }
  G = Gnew;
}
```

   What is the communication scheme if this algorithm is parallelized by partitioning the array row-wise over P processors (i.e., giving each processor N/P consecutive rows)? Give the pseudo-code for the parallel algorithm, making clear how the processors exchange data.

3. (a) What are the diameter and bisection width of a 3-dimensional mesh (lattice) topology that has 64 nodes in total?
   (b) Some machines contain multiple networks with different topologies. For example, the Blue Gene has a network with a 3D mesh topology and another network with a tree topology. Explain why it is useful to have different networks and give examples to illustrate this.

4. Message passing constructs can use (1) either synchronous or asynchronous sending and (2) either explicit or implicit receiving of messages. Explain which form(s) of sending and receiving are used in the message passing primitives of SR, MPI, and Java (RMI).

5. Consider the following HPF (High Performance Fortran) program fragment:

```
real A(5,5), B(5,5)
!HPF$ PROCESSORS P(5)
!HPF$ ALIGN B(I,J) WITH A(J,I)
!HPF$ DISTRIBUTE A(BLOCK,*) ONTO P
```

   How are the arrays A and B distributed over the 5 processors?

6. Consider an N-body system with 9 bodies, which will be simulated on 3 processors using the Barnes-Hut algorithm. The figure below shows two different distributions of the 9 bodies over the 3 processors; the number in each body indicates the CPU to which the body is assigned.

   Explain why the distribution may affect the performance of the parallel program. Which of the two distributions will give the best speedups? Why?

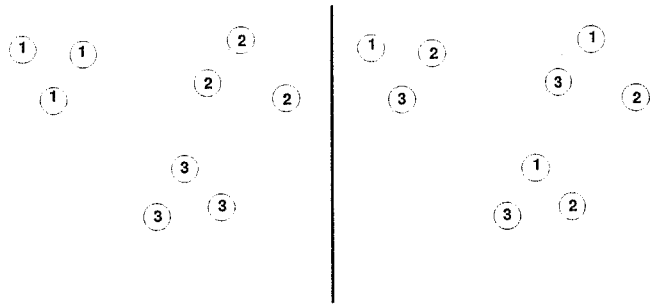

7. Consider the following four different parallel search algorihms:
   1 IDA* (a search algorithm based on work-stealing) without a shared transposition table
   2 IDA* with a shared replicated transposition table
   3 IDA* with a shared partitioned transposition table
   4 Transposition-Driven Search (TDS)

   (a) The four algorithms differ in the number of search-nodes they analyse (expand and evaluate). Rank the four algorithms in order of increasing number of nodes searched and explain your ordering.

   (b) The four algorithms also have different communication overheads for handling transposition table lookups and stores. Discuss for each algorithm from what type of communication overhead it suffers.

8. The direction of research in the field of Grid Computing is driven by a certain 'visionary aim', or 'promise'.

   (a) Explain what is meant by this 'promise of the Grid'.

   (b) Explain in what ways this promise is realized by the Parallel-Horus (or Jorus) framework for multimedia computing and its Ibis-based extensions. Discuss at least three techniques (solutions) that this framework provides to realize the promise of the grid.

Points

1a	1b	2	3a	3b	4	5	6	7a	7b	8a	8b
6	6	12	6	6	12	8	8	6	6	6	8

Total: 90 (+ 10 = 100)