M.Berzins

March 12, 2014

These questions are examples of the type of questions used in the mid-term exam for CS6230 the question paper usually consists of three questions. Full marks can be obtained from successful answers to two and a half questions.

SCHOOL OF COMPUTING
UNIVERSITY OF UTAH
CS6230 MID-TERM
Time allowed: 1 hour and Thirty Minutes
Attempt TWO questions.

Question 1

(a) Suppose that an operation with 2^p elements may be recursively decomposed into two operations with 2^{p-1} elements operations. Show with the aid of a diagram how this cab be recursively decomposed into p stages until $p = 1$. [2 marks]

(b) Explain how you would implement such an operation in parallel, assuming that the decomposition at the pth stage involves the communication of 2^{p-1} values between pairs of adjacent processes [2 marks]

(c) Perform a speed-up and efficiency analysis for this case. [4 marks]

(e) Define what is meant by the following terms: speedup, efficiency, weak scalability, strong scalability [2 marks]

(f) Write a sentence or two explaining what is meant by the each of the following terms: isoefficiency isomemory and isotime [3 marks]

(g) A computational problem has has a serial complexity of $O(n^k)$ where $k > 1$. Explain how you would achieve weak scalability for this problem and how you would pick a suitable core count. [4 marks]

(h) Consider the case when the storage required to solve this problem is $O(n^m)$ where there are three cases to consider. In case (i) $m << k$, in case (ii) $m = k$ and in case (iii) $m >> k$. For each case explain the storage implications on weak scalability. [3 marks]

Question 2

(a) Consider applying parallel processors to the task of solving the system of equations defined by

```
a(0,0)x(0)+a(0,1)x(1)+a(0,2)x(2)+a(0,3)x(3)+...+a(0,n)x(n)=b(0)
.                                                          .   .
.                                                          .   .            .
a(n,0)x(0)+a(n,1)x(1)+a(n,2)x(2)+a(n,3)x(3)+...+a(n,n)x(n)=b(n)
```

The algorithm that is used to solve these equations on a serial computer in the form of pseudo code is given by

```
for  {i=0 ;i< n-1; i++}            /* for each row except last
       for {j=i+1 ;j< n; j++}      /* step through subsequent rows
        m = a[j][i]/a[i][j]        /* compute multiplier
           for {k = i; k < n; k++}
              a[j][k] = a[j][k] -a[i][k]*m /* modify last n-i-1
              b[j] = b [j] - b[i]*m        /* elements of row j
           end
       end
end
```

(b) Explain how you would implement this algorithm in parallel assuming that the matrix is distributed in strips or partitioned in a cyclic fashion. [6 marks]

(c) Derive an expression for the communications associated with a value of i. [2 marks]

(d) Explain how these values would be broadcast on P processors using a standard broadcast procedure $log(P)((n-1)t_{startup} + (n+2)(n+1)*0.5-3)t_{data})$
and explain the assumptions behind this cost [4 marks]

(e) Given that the scalar complexity is $(n+3)(n+2)/2-3)$ derive an expression for the speedup and the parallel efficiency. [2 marks]

(f) Consider the effect of using a flat cost faster broadcast procedure on the efficiency. [2 marks]

(g) Consider the effect of $t_{startup} >> t_{data}$ and also of n very large on the efficiency. Determine the conditions under which the algorithm works best and those under which it does not work well. [4 marks]

Question 4

(a) In the context of message passing multiprocessors explain the meaning of:

1. message latency,
2. bandwidth.
3. speedup.
4. efficiency.

[2 marks]

(b) Explain in outline what is meant by the isoefficiency, isomemory and isotime. [3 marks]

(c) 1. Explain what is meant by a hypercube architecture, illustrating this for the case of 16 processors.

2. What does the function MPI_Allreduce do? Explain how you would implement this function efficiently using on a hypercube architecture (you may illustrate your answer by considering an 8 processor hypercube network).

[5 marks]

(d) Consider a bucket sort algorithm in which each processor only has one bucket. Suppose that sorting n numbers with the aid of m buckets takes $nlog(n/m)$ operations. Explain how the algorithn is modified to use p small buckets per processor. [4 marks]

(e) Derive the cost of redistributing p buckets from one processor to all the others and hence calculate the overall cost and then the efficiency. [4 marks]

(f) Which MPI function would you use to implement the redistribution? Where do you think the algorithm may not be scalable on a real calculation? [2 marks]

Question 5

(a) Consider a square mesh of $n \times n$ solution values denoted by u_{ij} and suppose that the calculation

$$u_{ij} = u_{ij} + \sum_{k=-m}^{m} a_{i+k,j} u_{i+k,j} + \sum_{l=-m}^{m} a_{i,j+l} u_{i,j+l} \quad i = 1, ..., n \text{ and } j = 1, ..., n.$$

is to be performed a known number of times. The values outside the range 1 to n may assumed to be known. Assume that this model is mapped onto a parallel computer by partitioning the mesh into (i) r horizontal strips and (ii) s equal size squares.

For the two partitions write down the amount of computation per processor. [4 marks]

(b) Assuming that the communication cost for moving m elements for one processor to another is $t_s + m\, t_m$ where t_s and t_m are known constants, for each partition write down the communications costs per processor.

[4 marks]

(c) For the two partitions write down the speed-ups and efficiencies [4 marks]

(d) Determine the isoefficiency requirements in each case as a function of n, m and p. [4 marks]

(e) Explain how you would implement a barrier synchronization on 2^n processors using

1. A linear approach
2. A tree-based approach
3. A hypercube approach
4. A Butterfly construction

[4 marks]

Question 6

(a) In the context of message passing multiprocessors explain the meaning of:

1. message latency and bandwidth,
2. speedup.
3. efficiency.

[3 marks]

(b) Explain in outline what is meant by the isoefficiency, isomemory and isotime. [3 marks]

(c) Explain with reference to the use of a large parallel computer why isoefficiency, isomemory and isotime are important aspects of different types of parallel performance. [3 marks]

(c)
1. Explain what is meant by a hypercube architecture, illustrating this for the case of 16 processors.
2. Describe two MPI global collective functions and how you would implement both these functions efficiently on a hypercube architecture (you may illustrate your answer by considering an 8 processor hypercube network).

[4 marks]

(c) Perform an isoefficiency analysis of an algorithm at for which n numbers are stored at the nodes of a hypercube and are added together. [3 marks]

(c) Explain how you would implement a barrier synchronization on 2^n processors using

1. A linear approach
2. A tree-based approach
3. A hypercube approach
4. A Butterfly construction

[4 marks]

Question 8

(a) Consider a square mesh of $n \times n$ solution values denoted by u_{ij} and suppose that the calculation

$$u_{ij} = u_{ij} + \sum_{k=-m}^{0} a_{i+k,j} u_{i+k,j} + \sum_{l=-m}^{0} a_{i,j+l} u_{i,j+l} \quad i = 1, \ldots, n \text{ and } j = 1, \ldots, n.$$

is to be performed a known number of times. (Please note that the stencil is slightly different from Laplaces equation.) The values outside the range 1 to n may assumed to be known. Assume that this model is mapped onto a parallel computer by partitioning the mesh into (i) r horizontal strips and (ii) s equal size squares.

For the two partitions write down the amount of computation per processor. [4 marks]

(b) Assuming that the communication cost for moving m elements for one processor to another is $t_s + m\, t_m$ where t_s and t_m are known constants, for each partition write down the communications costs per processor.

[4 marks]

(c) For the two partitions write down the speed-ups and efficiencies [2 marks]

(d) Determine the isoefficiency requirements in each case as a function of n, m and p. [2 marks]

(e) Suppose that the accuracy of the method $E(m)$ is defined in terms of n and m as:

$$E(m) = c \left(\frac{1}{n} \right)^m$$

where c is a constant. The method accuracy parallel efficiency $Meth_{ape}(m1, m2)$ takes into account the accuracy achieved by each of the two methods using different values of m given by $m1$ and $m2$, and is defined by

$$Meth_{ape}(m1, m2) = \frac{E(m1) P_{m1} T_{m1}}{E(m2) P_{m2} T_{m2}}$$

where $m1$ and $m2$ are values of m as in the definition of the method, P is the number of processors and T is the time. Write down an expression for $Meth_{ape}(m1, m2)$.

[4 marks]

What is the best possible value of $Meth_{ape}(m1, 1)$, for different choices of $m1$? [4 marks]

Question 9 (a) Explain how the use of Gaussian elimination results in the need to solve upper or lower triangular systems of linear equations. [2 marks]

(b) Consider applying parallel processors to the task of solving the triangular system of equations defined by

```
a(n,0)x(0)+a(n,1)x(1)+a(n,2)x(2)+a(n,3)x(3)+...+a(n,n)x(n)=b(n)
.                                                          .    .
.                                                          .    .
.                                                          .    .
a(5,0)x(0)+a(5,1)x(1)+a(5,2)x(2)+a(5,3)x(3)+...+a(5,5)x(5)=b(5)
a(4,0)x(0)+a(4,1)x(1)+a(4,2)x(2)+a(4,3)x(3)+a(4,4)x(4)    =b(4)
a(3,0)x(0)+a(3,1)x(1)+a(3,2)x(2)+a(3,3)x(3)               =b(3)
a(2,0)x(0)+a(2,1)x(1)+a(2,2)x(2)                          =b(2)
a(1,0)x(0)+a(1,1)x(1)                                     =b(1)
a(0,0)x(0)                                                =b(0)
```

`for x(0),x(1),x(2) in turn`

Write down the algorithm that is used to solve these equations on a serial computer in the form of pseudo code [4 marks]

(c) It is suggested that each of *p* processors be assigned one equation. Write down parallel pseudo code for a processor. [4 marks]

(d) Given that each of the *p* processors is assigned one equation write down how many inputs the *j*th processor receives, how much work it does and what it outputs.(Assume that a broadcast operation takes log(p) units of time). [4 marks]

(e) Hence calculate the total execution time and so the speedup. [4 marks]

(f) How does your algorithm change with respect to efficiency if each processor holds more than one equation? [2 marks]

Question 10

(a) In the context of message passing multiprocessors explain the meaning of:

1. message latency and bandwidth,
2. speedup.
3. efficiency.

[2 marks]

(b) Explain in outline what is meant by the isoefficiency, isomemory and isotime. [3 marks]
Explain with reference to the use of a large parallel computer why isoefficiency, isomemory and isotime are important aspects of different types of parallel performance. [3 marks]

(c) Consider the case of a computational problem with

1. A serial complexity of n^2
2. A parallel complexity of $(n^2)/p$
3. A parallel communications complexity of $c(n + logp)$
4. Parallel memory per processor of $(n^2/p + fn)$.

Derive a bound for 90% efficiency of the form $n \geq$ ..... and show what the associated memory useage is [4 marks]
Derive conditions for isotime and isomemory performance [2 marks]

(d) Describe three MPI forms of local communication and explain how each can help with parallel code safety, efficiency and performance. [3 marks]

(e) Describe what the functions MPI_Alltoall(), MPI_Allreduce, and MPI_Allgather do. Explain what how MPI_Allreduce might be implemented efficiently using either a hypercube approach or a butterfly construct. Illustrate your answers for the case of 16 processors. [5 marks]

Question 11

(a) Consider a square mesh of $n \times n$ solution values denoted by u_{ij} and suppose that the calculation

$$u_{ij} = u_{ij} + \sum_{l=-m}^{0} \sum_{k=-m}^{0} a_{i+k,ji+l} u_{i+k,j+l} \quad i = 1,...,n \text{ and } j = 1,...,n.$$

is to be performed a known number of times. (Please note that the stencil is very different from Laplaces equation.) The values outside the range 1 to n may assumed to be known. Assume that this model is mapped onto a parallel computer by partitioning the mesh into (i) r horizontal strips and (ii) s equal size squares.

For the two partitions write down the amount of computation per processor. [4 marks]

(b) Assuming that the communication cost for moving m elements for one processor to another is $t_s + m\, t_m$ where t_s and t_m are known constants, for each partition write down the communications costs per processor.

[4 marks]

(c) For the two partitions write down the speed-ups and efficiencies [2 marks]

(d) Determine the isoefficiency requirements in each case in terms of n, m and p. [2 marks]

(e) Consider the parallel solution of the triangular system of equations defined by

```
a(n,0)x(0)+a(n,1)x(1)+a(n,2)x(2)+a(n,3)x(3)+...+a(n,n)x(n)=b(n)
.                                                        .   .
.                                                        .   .
.                                                        .   .
a(5,0)x(0)+a(5,1)x(1)+a(5,2)x(2)+a(5,3)x(3)+...+a(5,5)x(5)=b(5)
a(4,0)x(0)+a(4,1)x(1)+a(4,2)x(2)+a(4,3)x(3)+a(4,4)x(4)    =b(4)
a(3,0)x(0)+a(3,1)x(1)+a(3,2)x(2)+a(3,3)x(3)               =b(3)
a(2,0)x(0)+a(2,1)x(1)+a(2,2)x(2)                          =b(2)
a(1,0)x(0)+a(1,1)x(1)                                     =b(1)
a(0,0)x(0)                                                =b(0)
```

`for x(0),x(1),x(2) in turn`

(f) It is suggested that each of p processors be assigned one equation. Write down parallel pseudo code for a processor. [4 marks]

(g) Given that each of the p processors is assigned one equation write down how many inputs the jth processor receives, how much work it does and what it outputs.(Assume that a broadcast operation takes log(p) units of time). [2 marks]

(h) Hence calculate the total execution time and so the speedup. [2 marks]

Question 12

(a) Explain the meaning of the term overhead function in connection with the execution of a parallel program. [2 marks]

(b) suppose that the fraction of a code whose serial execution is T_s that is inherently serial is f. Derive Amdahl's law for the maximum attainable speedup $S(p)$ using p processors

$$S(p) = \frac{p}{1+(p-1)f}$$

When $f = 0.1$ explain what is the maxiumum possible speedup? Explain why for many practical problems the predictions from Amdahl's law tend not to be realistic. [2 marks]

(c) Explain the reasoning behind the Karp-Flatt metric and used Amdahl's law to derive this metric in terms of the observed speedup and the number of processors. [4 marks]

(d) The following timings are obtained for two almost identical codes running with between 1539 and 98304 cores. Using 1536 cores as the base case the following results were obtained while trying to estimate and improve the scalabilty of the Uintah code.

P	1536	3072	6144	12288	24576	49152	98304
P/1536	1	2	4	8	16	32	64
Code A	240	120	61	31	17	9	4.6
Code B	330	190	120	57	33	20	12

where $P/1536$ is the multiplier of the number of processors over the base case. Calculate the speedup over the base case for both codes for $P/1536 = 4, 16$ and 64. [3 marks]

(e) Calculate the value of the Karp-Flatt metric for both codes, when $P/1535 = 4, 16$ and 64, by using these $P/1536$ values in the metric calculation. [6 marks]

(f) By comparing the calculated values of the Karp-Flatt metric for both cases explain, as best you can, the performance of the code.

[3 marks]

Question 13

(a) Explain how it is possible to model the effect of memory transfer of the performance of an algorithm by assuming just 2 levels in the hierarchy, fast and slow. Suppose that all data is initially in slow memory. If m is the number of memory elements (words) moved between fast and slow memory, Let t_m be the time per slow memory operation, f be the number of arithmetic operations t_f be the time per arithmetic operation where $t_f << t_m$ and le t $q = f/m$ be the average number of flops per slow element access. Derive an expression for the minimum possible time when all data is in fast memory [2 marks]

(b) Derive an expression for the actual time [1 mark]

(c) Explain why it is important for q to be large. [1 mark]

(d) Consider the standard approach to implementing matrix matrix multiplication C = C + A*B

```
for i = 1 to n
 {read row i of A into fast memory}
  for j = 1 to n
      {read C(i,j) into fast memory}
      {read column j of B into fast memory}
      for k = 1 to n
          C(i,j) = C(i,j) + A(i,k) * B(k,j)
      {write C(i,j) back to slow memory}
```

Show that this approach needs = $n^3 + 3n^2$ reads and hence derive an expression for the value of q. [4 marks]

(e) Now consider the block based approach Let A,B,C be n by n matrices split into N by N matrices of b by b sub blocks where block size is b=n / N

```
  for i = 1 to N
       for j = 1 to N
       {read block C(i,j) into fast memory}
       for k = 1 to N
                  {read block A(i,k) into fast memory}
                  {read block B(k,j) into fast memory}
                  C(i,j) = C(i,j) + A(i,k) * B(k,j) {do a matrix multiply on b
       {write block C(i,j) back to slow memory}
```

Show that this approach needs = $n^3 + 3n^2$ reads and hence derive an expression for the value of q. [4 marks]

(f) Strassen's algorithm for multiplying the two 2x2 matrices A and B with entries aij and bij respectively to get a third matrix C is given by:

```
d1= (a11+a22) * (b11+b22)
d2= (a12-a22) * (b21+b22)
d3= (a11-a21) * (b11+b12)
d4= (a11+a12) * (b22)
d5= (a21+a22) * (b11)
d6=(a11) * (b12-b22)
d7=(a22) *(-b11+b21)
C11 = d1 + d2 d4 + d7
C12 = d4 + d6
C21 = d5 + d7
C22 = d1 - d3 d5 + d6
```

Calculate the computational complexity of this algorithm and contrast it with the standard approach. [4 marks]

(g) Describe a recursive algorithm based upon Strassen's algorithm for matrix-matrix multiply. How is the complexity changed if the elements considered are not just scalars but matrices themselves? [4 marks]