THE AUSTRALIAN NATIONAL UNIVERSITY

First Semester Examination – June 2005

COMP3320/COMP6464/Honours

High Performance Scientific Computing

Study Period: 15 minutes

Time Allowed: 3 hours

Permitted Materials: Calculator

Exam questions total 70 marks.

COMP3320 students answer questions 1-6
COMP6464 students answer questions 1-5 and 7
Honours students answer questions 1-5 and 8

Clarity and conciseness in answers will be highly valued
Marks may be lost for supplying irrelevant information

Question 1 [8 marks]

Fundamentals of Numerical Computing

(a) On a computer you are attempting to evaluate e by summing the series

$$e^x = 1 + x + \frac{x^2}{2!} + \frac{x^3}{3!} + \cdots$$

for $x = 1$. Discuss how, as a function of terms in the series, your result will be affected by truncation and rounding errors.

[2 marks]

(b) Assume a decimal (base 10) floating-point number system having machine precision $\epsilon_{\text{mach}} = 10^{-5}$ and an exponent range of ± 20. What is the result of each of the following floating-point arithmetic operations:

i) $1 + 10^{-7}$

ii) $1 + 10^{3}$

iii) $1 + 10^{7}$

iv) $10^{10} + 10^{3}$

v) $10^{10}/10^{-15}$

vi) $10^{-10} + 10^{-15}$

[3 marks]

(c) Each of the following statements are either true or false. Which are they? Give a few sentences to explain your answers.

(i) If two real numbers are exactly representable as floating-point numbers, then the result of any real arithmetic operation on them will also be exactly representable as a floating-point number.

(ii) Between the underflow limit and overflow limit computer implementations of floating point numbers are distributed uniformly.

(iii) In a floating point number system the underflow level is the smallest positive number that perturbs the number 1 when added to it.

[3 marks]

Question 2 [8 marks]

Performance Modeling and Measurement

(a) Explain the difference between elapsed time, user time, and system time.

[2 marks]

(b) You have built a cluster computer using dual CPU nodes linked together via a gigabit Ethernet interconnect. You are trying to develop a performance model for a code running on this system. Parts of the code are sequential running on only one CPU, other parts are OpenMP parallel exploiting just two CPUs on one node of the cluster, while other parts are MPI parallel capable of exploiting the entire cluster. Associated with these three levels of parallelism, you define three computation rates, R_1, R_2 and R_3 (not necessarily in order of parallelism given above). You denote the fraction of work that can be executed at rate R_i as f_i, and $f_1 + f_2 + f_3 = 1$.

(i) Write an expression for the average computation rate (R_a) achieved by the code in terms of the three different R and f values.

(ii) You estimate $R_1 = 190$MIPS, $R_2 = 2880$MIPS and $R_3 = 100$MIPS. Which rate do you think corresponds to which level of parallelism, and how many CPUs do you think are on the cluster? Explain your answer.

(iii) For a particular code you estimate that $f_1 = 0.3$, $f_2 = 0.2$ and $f_3 = 0.5$. Using the values for R given above how much faster is this application when run over the entire cluster, compared to when it is run on a **single** CPU of the cluster?

[6 marks]

Question 3 [16 marks]

High Performance Computer Architecture

(a) Why is branch prediction so important on modern day microprocessors?

[1 mark]

(b) What is the difference between a write-through and copy-back cache? In a modern computer system where are you likely to find each type of cache used?

[2 marks]

(c) Cache misses can be classified as i) conflict misses, ii) compulsory/cold misses, or iii) capacity misses. Define each of these terms. Discuss what hardware/software modifications can be made (if any) in order to minimize each type of cache miss.

[5 marks]

(d) Discuss the performance of the following C loops on a typical modern day processor:

(i) `for (i=0; i<n; i++) a = a + x[i]*x[i];`

(ii) `for (i=0; i<n; i++) x[i] = (a*y[i]+b*y[i])*c;`

(iii) `for (i=0; i<n; i++) x[i] = (y[i]+a*z[i])*(y[i]-a*z[i])/2;`

(iv) `for (i=0; i<n; i++) x[i+1] = x[i] + z[i];`

You should assume that `x, y, z` are double precision floating point arrays, while `a, b, c` are double precision floating point variables. Your discussion should include estimates of the likely cycles per iteration, and how this might change as a function of `n`. Any architectural details that you assume in order to derive your estimates must be clearly identified, as should any assumptions you make concerning data layout. Also, if there are modifications to the code as written that are likely to arise from optimizations undertaken by the compiler you must clearly identify these.

[8 marks]

Question 4 [14 marks]

Shared Memory Parallelism

(a) OpenMP is often referred to as a fork/join parallel model. What is meant by this?

[2 marks]

(b) In the context of parallel processing what is an anti-dependency? Give pseudo-code that illustrates an anti-dependency.

[2 marks]

(c) In the following C code arrays `a` and `b` correspond to distinct memory regions.

```
void mystery(double a[], double b[], int n){
  double factor;
  int i,j,k;
  for (k=0; k<n; k++){
    for (i=k+1; i<n; i++){
      factor=a[i*n+k]/a[k*n+k];
      for (j=k+1; j<n; j++){
        a[i*n+j]=a[i*n+j]-factor*a[k*n+j];
      }
      b[i]=b[i]-factor*b[k];
    }
  }
  return;
}
```

(i) Outline how you would parallelise this code on a shared memory parallel computer using OpenMP. You are not required to reproduce the exact OpenMP syntax, but are required to make the intent of your directives clear.

(ii) Comment on how well you would expect your parallel code to perform as a function of problem size and number of OpenMP threads used.

[10 marks]

Question 5 [17 marks]

Distributed Memory Parallelism

(a) Three processes are running on three processors of a distributed memory multiporocessor. The processes execute the following statements:

Process X	Process Y	Process Z
X1	Y1	Z1
receive i from Y	send p to Z	send a to X
X2	Y2	Z2
receive j from Z	receive q from Z	receive b from Y
X3	Y3	Z3
send k to Y	send r to X	send c to Y
X4	Y4	Z4
send n to Z	receive s from X	receive d from X
X5	Y5	Z5

All data movement between processors is by means of nonblocking *send* and blocking *receive* functions.

(i) What is the difference between blocking and non-blocking communications?

(ii) Show the partial order imposed on statements `X1-X5, Y1-Y5`, and `Z1-Z5` that results from the communications. You should assume that message transmission takes two time units and that statements take one time unit each except for *send* and *receive*, which take no time except for transmission. How many time units are required for the program to complete?

(iii) What would happen if both send and receive were blocking?

[7 marks]

(b) MPI_Allreduce, MPI_Reduce and MPI_Bcast are three MPI functions.

(i) Briefly describe what each of these functions do.

(ii) How would you expect the time taken by each function to vary with the number of processes involved? Explain your answer.

(iii) Do these functions also synchronize the processes involved? Explain your answer.

[10 marks]

COMP3320 STUDENTS ONLY

Question 6 [7 marks]

(a) For several years the Earth Simulator in Japan was the most powerful supercomputer in the world.

(i) Briefly explain why such a powerful machine is required to perform global climate simulations.

(ii) All scientific computations involve approximations. Detail four different approximations used in global climate simulations.

[4 marks]

(b) Computational science is multilingual. Give one positive and one negative relating to the use of i) C, ii) Python, and iii) Fortran95 in computational science applications.

[3 marks]

(c) Optional feedback — but only if you want to and have time:

(i) This exam was

1. Impossibly hard covering material that was not in the course!
2. Very hard and I wish I'd studied more!
3. Challenging but fair.
4. Too simple and won't show how good a student I am!
5. So easy I could have scored 100% without attending any lectures or labs.
6. Other, please elaborate.

(ii) In this exam I estimate that I will score:

1. $\geq 80\%$
2. $\geq 70\%$
3. $\geq 60\%$
4. $\geq 50\%$
5. $< 50\%$ but $\geq 45\%$ so I'll be back for a supplementary (please make it easy!).
6. Too low to estimate!

[0 marks]

COMP6464 STUDENTS ONLY

Question 7 [7 marks]

(a) Molecular dynamics simulations often use periodic boundary conditions.

(i) What are periodic boundary conditions?

(ii) Simulations may use periodic boundary conditions in 1, 2 or 3 dimensions. For each possibility give one example of a physical system for which use of this level of periodic boundary conditions is meaningful.

(iii) Two key parameters for simulations using periodic boundaries are the "box size" and the "cutoff". What do these parameters refer to and how are they related (if at all)?

(iv) From a computational perspective, are calculations that use periodic boundary conditions more or less costly than equivalent calculations that do not use periodic boundary conditions? Explain your answer.

[7 marks]

(b) Optional feedback — but only if you want to and have time:

(i) This exam was

1. Impossibly hard covering material that was not in the course!
2. Very hard and I wish I'd studied more!
3. Challenging but fair.
4. Too simple and won't show how good a student I am!
5. So easy I could have scored 100% without attending any lectures or labs.
6. Other, please elaborate.

(ii) In this exam I estimate that I will score:

1. $\geq 80\%$
2. $\geq 70\%$
3. $\geq 60\%$
4. $\geq 50\%$
5. $< 50\%$ but $\geq 45\%$ so I'll be back for a supplementary (please make it easy!).
6. Too low to estimate!

[0 marks]

HONOURS STUDENTS ONLY

Question 8 [7 marks]

(a) The following C code implements a naïve square in-situ matrix transposition:

```
for (i = 1; i < n; i++)
   for (j = 0; j < i; j++){
      tmp = A[i][j];
      A[i][j] = A[j][i];
      A[j][i] = tmp;
   }
```

(i) For certain matrix dimensions the above code will show poor cache performance. When will this occur?

(ii) Two alternative algorithms for matrix transposition are the "cache blocked" and the "cache oblivious" approaches. Briefly outline how both of these two alternative algorithms work.

(iii) In addition to matrix transposition, another algorithm that is amenable to implementation in a cache oblivious manner is the fast Fourier transform. What are the key ideas behind all cache oblivious approaches.

(iv) In May 2000 the Society for Industrial and Applied Mathematics (SIAM) published an article naming the top 10 algorithms of the 20th century. One of the named algorithms was the fast Fourier transform, name two others.

[7 marks]

(b) Optional feedback — but only if you want to and have time:

(i) This exam was

1. Impossibly hard covering material that was not in the course!
2. Very hard and I wish I'd studied more!
3. Challenging but fair.
4. Too simple and won't show how good a student I am!
5. So easy I could have scored 100% without attending any lectures or labs.
6. Other, please elaborate.

(ii) In this exam I estimate that I will score:

1. $\geq 80\%$
2. $\geq 70\%$
3. $\geq 60\%$
4. $\geq 50\%$
5. $< 50\%$ but $\geq 45\%$ so I'll be back for a supplementary (please make it easy!).
6. Too low to estimate!

[0 marks]