

TENTAMEN / EXAM

TDDC78

Programmering av paralleldatorer /
Programming of parallel computers
2013-10-22, 08:00–12:00, TER1

Christoph Kessler Dept. of Computer Science (IDA)
Linköping University

Hjälpmedel / Allowed aids: Engelsk ordbok /
dictionary from English to your native language

Examinator: Christoph Kessler

Jourhavande lärare:

Christoph Kessler (IDA) 013-28-2406; 0703-666687; visiting at ca. 10:00

Maximalt antal poäng / Max. #points: 40

Betyg / Grading (prel.): The preliminary threshold for passing (grade 3) is at 20p, for grade 4 at 28p, for grade 5 at 34p.

Because of new regulations by Linköping University, we can no longer give ECTS grades.

General instructions

- Please use a new sheet of paper for each assignment. Order your sheets by assignments, number them, and mark them on top with your exam ID number and the course code.
- You may answer in either English or Swedish. **English is preferred** because not all correcting assistants understand Swedish.
- Write clearly. Unreadable text will be ignored.
- Be precise in your statements. Unprecise formulations may lead to a reduction of points.
- Motivate clearly all statements and reasoning.
- Explain calculations and solution procedures.
- The assignments are *not* ordered according to difficulty.

1. (6.5 p.) **Performance tools and analysis**

- (a) (1 p.) Which performance data collection method is required in order to be able to draw a communication statistics diagram (displaying the amount of bytes communicated between every pair of processes)? Justify your answer (technical reasons).
- (b) (1p) Modern tool-suites for performance analysis of parallel programs consist of a collection of several kinds of tools. Give four different kinds of such tools. (*I.e., no concrete tool names, but a short term saying what each kind of tool does.*)
- (c) (1p) When is a parallel algorithm for some problem (asymptotically) *cost-optimal*? (*give a formal definition*)
- (d) (2 p.) (*i*) Derive Amdahl's law and (*ii*) give its interpretation. (*Note: a picture is nice but is not a proof; a calculation is expected for the derivation of Amdahl's Law.*)
- (e) (1p) How is the so-called *peak performance* of a parallel computer system such as Triolith calculated? (give a formula)
- (f) (0.5 p.) How are the supercomputers on the TOP500 list ranked?

2. (2 p.) **Parallel programming models**

Explain the parallel program execution styles *fork-join* and *SPMD*, and explain the main difference. (1p)

Which one is more comfortable for the programmer, and why? (0.5p)

Which one is used in MPI? (0.5p)

3. (4 p.) **Parallel program design methodology**

Foster's design methodology consists of four stages. Name and explain them. Give details about their goals. What are the important properties of the result of each stage? Be thorough!

4. (6 p.) **Parallel computer architecture**

- (a) How does the MSI *write-invalidate protocol* in a bus-based shared memory system work? (2p)
- (b) Define and explain *sequential memory consistency*. (1p)
- (c) (3p) Explain the *2D-Torus* interconnection network topology. (0.5p)
Determine its maximum node degree, the maximum distance of any two nodes, and the maximum aggregated bandwidth (given by the maximum possible number of communications proceeding simultaneously in the network). (1.5p)
Shortly describe one kind of message-passing computation (e.g., from your labs) for which a 2D-Torus network would be more efficient than its 1D equivalent, and why. (1p)

5. (2 p.) **OpenMP**

- (a) In principle, every OpenMP program could likewise be expressed by explicit thread programming, e.g. by using libraries such as Pthreads. What is the main advantage of using OpenMP over such thread programming libraries? Explain your answer. (1p)
- (b) Why is the design of OpenMP helpful for *incremental* parallelization of sequential codes? (1p)

6. (5 p.) **Parallel Basic Linear Algebra**

- (a) (3p) The BLAS (Basic Linear Algebra Subroutines) library API comes in three different levels (BLAS Level 1, 2, 3).
 - (i) Name one representative function from each level and give a short description. (1.5p)
 - (ii) Why is it very important for a supercomputing center to have efficient implementations of BLAS installed? (0.5p)
 - (iii) Which of the levels is most likely to have efficient parallel implementations even where interprocessor communication is expensive, and why? (1p)
- (b) (2p) Consider the following straightforward sequential code for the multiplication of two $N \times N$ square matrices A and B :

```
for (i=0; i<N; i++)
  for (j=0; j<N; j++)
    C[i][j] = 0.0;
for (i=0; i<N; i++)
  for (j=0; j<N; j++)
    for (k=0; k<N; k++)
      C[i][j] = C[i][j] + A[i][k] * B[k][j];
```

Apply *tiling* with tile size $S \times S$ (with $1 < S \ll N$) to the second loop nest. Show the resulting pseudocode.

Why is tiling (with a suitably chosen S) beneficial for the performance of this code on modern processor architectures?

7. (9 p.) **Parallel Solving of Linear Equation Systems**

- (a) (4p) Consider the problem of solving a triangular linear equation system, e.g. a system $Ux = b$ where U is a $n \times n$ upper triangular matrix (i.e., all elements below the main diagonal are zero) by *backward substitution*, with the following sequential pseudocode:

```
 $x_{n-1} \leftarrow b_{n-1}/u_{n-1,n-1};$ 
for  $i = n - 2$  downto 0 do
   $x_i \leftarrow \left( b_i - \sum_{j=i+1}^{n-1} u_{i,j}x_j \right) / u_{i,i};$ 
```

- i. Derive the asymptotic execution time of the above sequential algorithm as an expression in n . (0.5p)
 - ii. Suggest a possible parallelization for a shared memory system. What kind(s) of parallelism can you exploit here? Show the resulting pseudocode (using e.g. OpenMP, pthreads, or skeletons), and derive the parallel execution time as an expression in n and p . You can ignore cache issues for simplicity. Calculate the parallel efficiency, and discuss possible performance issues where the number p of parallel threads is large. (3.5p)
- (b) (5p) *Column pivoting* searches in each step k of Gaussian Elimination within column k for an element $a_{i,k}$ with maximum absolute value and then exchanges entire row i with row k (unless just $i = k$) before the updating of submatrix $A_{k..n-1,k..n-1}$ of step k is being performed.

For a *row-block-wise* distribution of the system matrix A , describe the required additional computation and communication operations for column pivoting, and give its parallel time cost as a formula in k and n for the delay model.

For the above analysis, use the delay model, assuming for simplicity that the underlying interconnection network is fully connected (i.e., each processor has a direct network link to every other processor) with message startup time α (clock cycles) and average (float) word transfer time β (clock cycles), and that a processor can only send or receive one message at a time. If you need to make further assumptions, state them carefully.

8. (1.5 p.) Transformation and Parallelization of Sequential Loops

Given the following C loop:

```
for (i=0; i<N; i++)
    a[i] = 2.0 * a[i+1] / b[i] + c[i];
```

- (a) Explain why the loop iterations cannot be executed in parallel in this form (dependence-based argument). (0.5p)
- (b) Suggest a (semantics-preserving) transformation of the loop (show the new code) such that the loop iterations now could be executed in parallel. Explain why (dependence-based argument). (1p)

9. (3.5 p.) MPI

- (a) How does the *nonblocking* (also called *incomplete*) send routine `MPI_Isend` differ from the ordinary blocking `MPI_Send`? Under what condition is it safe to replace a `MPI_Send` call by `MPI_Isend`? (1.5p)
- (b) (2 p.) Explain the Communicator concept in MPI. How does it support the construction of parallel software components?

10. (0.5 p.) Grid Computing

Name one task performed by *grid middleware*.