THE BCS PROFESSIONAL EXAMINATION
Professional Graduate Diploma

April 2006

EXAMINERS' REPORT

Distributed & Parallel Systems

There has been some discussion recently about the popularity of this paper and, with only a small number of candidates this year, there is no sign of any increase in interest. There is a noticeable tendency within the paper for candidates to opt for questions directed toward distributed systems (questions one and two) rather than parallel systems (questions three and four). Too few candidates answered questions 3 and 4 for a report to be published on these questions. The style of Question 5 can now be anticipated and remains a popular choice.

Question 1

1. *a)* An enterprise application uses a mix of fixed and mobile clients to deliver services to its customers. List and discuss the four key challenges faced by the distributed system, and suggest how these challenges can be met**.** **(12 marks)**

b) Explain highlighting the notion of physical time, why it is necessary to synchronise computer clocks in distributed systems. Briefly discuss the technique of synchronising the clock of a client process with an external time server. Comment on the issues involved and the ways of addressing them. **(13 marks)**

Answer Pointers

a) Meeting challenges

Heterogeneity: hardware (fixed and mobile clients) and software (application and system software). The best way to meet these challenges is to use middleware which can handle disparate hardware and software by providing a uniform computational model. 3 marks

Security: Firewalls, filters and effective access and authorisation management to provide security against attacks such as packing sniffing, IP spoofing, denial of service etc. 3 marks

Transparency: Access and location transparency challenges met by effective middleware support. 2 marks

Failure handling: Detecting failures, masking failures, recovery from failures and redundancy. 2 marks

b) The problem of maintaining common physical time and limitation to time stamp events at different nodes in a distributed system and order events based on a common time and maintaining global states etc. necessitates synchronising physical clocks. The issues are correctness of clocks and drift. Two modes of synchronisation – internal and external. 4 marks

The time server should maintain the Cordinated Universal Time (UTC). The techniques work as follows:

- the process P request the time in a message m1
- it receives a time value t in message m2
- the process P record the total round-trip time T taken to send m1 and receive m2.
- P could set its clock to t+T/2
- The issues are accuracy of time server and variability of message delays

- Solution is to synchronise time servers and a number of m1 requests and taking the minimum T.

Issues – 2 marks
Algorithm – 7 marks

(Colouris pp 386-389, pp 391-392)

Examiner's Comments

a) The four key challenges are heterogeneity, security, transparency and failure handling. One of the best ways to handle i) heterogeneity challenge is to use middleware which provides a uniform computational model ii) security challenge is through effective authentication and access management assisted by hardware mechanisms like filters and firewalls iii) transparency challenge is again through middleware iv) failure challenge is detecting failures early, masking them if necessary and adapting effective recovery and redundancy techniques. A few students got a couple of the above right and none got all right. The key to the answer lies in the question itself as a mix of fixed and mobile systems indicate heterogeneity and heterogeneous systems have to invariably meet the above challenges.

b) The answer to this part is well discussed in the text 'Distributed Systems- Concepts and Design', George Coulouris, Jean Dollimore and Tim Kindberg (Third Edition). The problem of maintaining common physical time and limitation to time stamp events at different nodes in a distributed system and order events based on a common time and maintaining global states etc.. necessitates synchronising physical clocks. The issues are correctness of clocks and drift. The technique of synchronising the clock of a client process with an external server is to allow the time server to maintain Cordinated Universal Time (UTC) through algorithm discussed in the text. Again only a few students were able to answer the question.

Question 2

2. *a)* Discuss the two approaches used for file sharing in Peer-Peer (P2P) systems highlighting advantages and disadvantages of each approach**.** **(12 marks)**

b) Explain the purpose of replicating data objects in distributed systems. Discuss clearly highlighting issues, an architectural approach for carrying out transactions on replicated objects in distributed systems. **(13 marks)**

Answer Pointers

a) The two approaches:
- Centralised directory, hybrid client-server and P2P architecture (Napster approach)
- search and file transfer (Gnutella approach)

Centralised directory: use of a P2P node launches P2P file sharing application which contacts the centralised directory informing its IP address and names of objects in its local disk for sharing purposes. Similar information is collected from all nodes by the directory server. Mapping of objects and IP addresses are updated in the directory. The directory server periodically sends message to nodes to determine whether they are active.

Advantages - directory query processing
- central directory makes locating objects easy

Disadvantages - single point of failure – the central directory
- performance degradation if node numbers increase

Approaches – 2 marks
Centralised directory 5 marks

Search and transfer: fully distributed P2 approach. Peers form an abstract logical network called an overlay network-a node and edge graph approach. Any time, a peer is in contact with a subset of the number of nodes in the network. The search for an object is through query flooding. A peer sends messages to neighbouring peers and wait for a "query hit". A TCP connection is established between the nodes seeking query and the node with the query hit.

Advantages - pure P2P decentralised approach
- simple architecture

Disadvantages - time consuming query flooding
- traffic and scalability problems

Search and transfer – 5 marks

b) Replication of data objects provides high availability, enhanced performance and faulty tolerance. Replication thus enhances services provided by a distributed system. For example the DNS naming service maintains copies of name-to-attribute mappings for computer across the internet.

A common approach is to use caching of data. For example, browsers and proxy servers cache copies of web resources to avoid latency of fetching. Performance is enhanced.

The issues in availability are service failures and communication disconnection. Replicated copies of data objects kept in back-up server provides continuity of service and automatic access to this server addresses communication disconnection issue.

High availability coupled with consistency of replicated data updating of all copies of replicated data ensures fault tolerant service.

2 marks each (high availability, improved performance and fault tolerance) x 3 = 6 marks

Architectural characteristics:
replication managers
Client front end
Primary copy replication
Example: the Gossip architecture

The issues are consistency of data objects and concurrency control.

Solution: Read-one/write all procedure and two-phase commit protocol.

Front end clients communicate with a replication manager to perform an operation on the data object. This manager is responsible for co-operation with other replication managers. A read request can be performed by one replica manager but a write request must be performed by all replica managers. The operations are update response, co-ordination, execution and agreement, the latter two are group operations. Co-ordination is carried out through a two-phase commit protocols.

Read-one/write all operation; write lock is used by all replication managers for write operation and read operation is performed by a single replication manager by setting the read lock.

7 marks

(Colouris pp 554-556, 574, 592-594)

Examiner's Comments

a) Those who answered this question got the two approaches right which are: centralised directory, hybrid client-server and P2P architecture (Napster approach) and search and file transfer (Gnutella approach). In the first, the main advantage is direct query processing and the disadvantage is performance degradation as nodes increase, where as in the second, the main advantage is simple architectural approach and the disadvantage is time consuming query flooding

b) Some of the best answers from those whose attempted this question were produced here. The purpose of replicating data objects is to provide high availability, enhanced performance and fault tolerance. The common approach is to cache data and the issues for discussion are consistency of data objects and concurrency control.

Question 3

3. *a)* There always seems to be a demand for greater computational power from computer systems than is currently possible. From where does this demand arise? **(7 marks)**

b) In terms of speed-up factor and scalablility:

i) what is the potential for increased computational speed within a multiprocessor system? **(8 marks)**

ii) what is meant by *super linear speed up*? **(4 marks)**

c) A multiprocessor consists of 100 processors, each capable of a peak execution rate of 2 Gflops. What is the performance of the system as measured in Gflops, when 10% of the code is sequential and 90% is executed in parallel? **(6 marks)**

Answer Pointers

a) The demand for computational speed often comes from numerical simulation of scientific and engineering problems that must be completed in a reasonable time. Examples lie within the fields of weather forecasting or the motion of astronomical or molecular bodies. Alternatively the emergence of the internet has given rise to an increasing demand from users that require web servers to handle increasingly higher rates of request. (Other answers accepted).

7 marks

b) in determining how much faster a multiprocessor solves a problem, the time (t_p) is compared to the time using a single processor system (t_s) using the best sequential algorithm. This is the Speedup Factor, $S(p) = t_s / t_p$.

The maximum speedup is limited by period due to idle processors due to extra computations that may be required in the parallel version that are not required in the sequential version or due to communication time between processes.

Combined architecture/algorithmic scalability suggests that increased problem size can be accommodated with increased system size making use of constant problem-size scaling (Amdahl's assumption), time-constrained scaling (Gustafson's assumption) or memory-constrained scaling. 8 marks

Super linear speed up is really an invalid concept and usually applies only because the original sequential algorithm is suboptimal. 4 marks

c) Using Amdahl's Law or otherwise, S(p) = ~10 6 marks

(See Wilkinson, p3-12)

Examiner's Comments

No Examiner's Comments have been published for this question.

Question 4

4. *a)* Shifting, scaling, rotation and clipping are graphical operations that can be performed upon a stored image. Explain why these operations may be considered *naturally* parallel? **(4 marks)**

b) Summarise three different programming strategies that may apply to problems that can be divided into independent parts. **(9 marks)**

c) For each strategy summarised above, describe an application that illustrates the use of that strategy. **(12 marks)**

Answer Pointers

a) The operations require geometrical transformations on the 2D pixel coordinates (which are shifted, magnified or reduced, rotated or contained respectively). Since these operations are totally independent from the transformations on other pixels and result in simply an updated bitmap they are considered truly natural, (Wilkinson, p81) 4 marks

b) Strategies that may be summarised are: partitioning and divide and conquer, pipelining, synchronous computations, and load balancing, (but other alternative lists are acceptable). 9 marks

c) For each strategy several example applications are covered in the literature. One typical illustrative application for each of the strategies is given respectively below:
bucket sort, prime number generation, heat distribution, shortest path problem. 12 marks

(Wilkinson, Chapters 4, 5, 6 and 7)

Examiner's Comments

No Examiner's Comments have been published for this question.

Question 5

5. You have agreed to talk for 30 minutes at the next meeting of your local BCS branch. The title of your talk is *Trends in Distributed and Parallel Systems: Looking from the Past and Present toward the Future.*

Sketch out approximately eight presentation slides, with associated notes, that you would use for your talk. **(25 marks)**

Answer Pointers

This type of question makes a regular appearance at the end of the question paper each year. This year, the topic is more demanding since it requires candidates to speculate about the future based upon their experience and knowledge.

The number of slides indicate to the candidate that they should spend approximately 5 minutes on each slide.

(Marks: 5 for correctness, 5 for clarity, 5 for completeness, 5 for relevance and 5 for depth).

Examiner's Comments

For this question, candidates generally gathered marks within the 15 to 20 (out of 25) marks band. Whilst this represented a satisfactory performance it nonetheless was gained as a result of regurgitated bookwork, rather than tailoring the answer to suit particular points in the question. It may be worthwhile indicating to candidates in future that marks will be awarded equally for correctness, clarity, completeness, relevance and depth in their answer to this question.