THE BCS PROFESSIONAL EXAMINATION
Professional Graduate Diploma

April 2003

EXAMINERS' REPORT

Distributed & Parallel Systems

General

As in previous years all candidates demonstrated a sound knowledge of the topics associated with this subject. There was a 100% pass rate and the average mark was 63%, which is similar to that of previous year, although there was a large increase in the number of candidates.

Many more candidates sat the examination this year compared with last and on the whole did very much better than previously. On reflection it would appear that the questions were more clearly defined than hitherto and possibly therefore more easy for the candidates to prepare for. In particular, question four replicated a familiar format and, as may be expected, proved to be the most popular choice in the set of questions offered.

Some answers continue to be written in English that is not only grammatically awkward but is also hard to understand as a result of poor handwriting.

Question 1

The developer of a large scale distributed system must take into account a number of important design issues.

a) Explain briefly what the terms *performance, availability* and *scalability* mean when applied to a distributed system. Why are these issues so important during the design process? **(13 marks)**

b) Discuss different approaches to load balancing in distributed systems with replicated servers.
(Where possible, you should illustrate your answers with examples from systems you have used or studied).
(12 marks)

Answer Pointers

a) Speed of response, throughput; reliability concerned with 24*7*365 availability and correctness of results; scalability – ability to add user, increase data volumes. These issues need to be addressed early on in the design process, difficult to build in later but equally difficult to assess early on.

b) A good answer would connect the issues a) with the use of replicated servers to enhance performance, reliability and scalability. Static or dynamic load balancing can be used, different trade-offs. With dynamic load balancing, state information management is an important issue in stateful servers.

Comments

This question was in general answered well with candidates showing that they understood the basic issues. More use of metrics in a) would have improved the answers. In b) there was good detailed discussion of different load balancing

algorithms/strategies, but not much attempt to explore the tradeoffs between them. At this level it would be good to see clear evidence of critical evaluation.

Question 2

Internetworking allows computers on independent and incompatible networks to communicate reliably and efficiently.

a) **Explain in general terms how internetworks may be implemented.** **(7 marks)**

b) **Explain in particular how the Internet emerged.** **(7 marks)**

c) **Explain how internetworking protocols enable reliable communication.** **(7 marks)**

d) **Why might it be unwise to use bandwidth as the only measure of network performance?** **(4 marks)**

This question was reasonably popular and those who tackled it, with only one exception, gained at least half marks. There was some weakness in knowledge of the history of the Internet.

Answer Pointers

a) Internetworks are implemented by linking component networks with dedicated packet routing computers (routers) or by general purpose computers (gateways) and adding protocols that support the addressing and transmission of data through what may be thought of as a "virtual network". Some further description of the role of router and gateway is apt, as is an indication that internetworks can be wide-area or single-site implementation each with their distinct advantages and disadvantages. For instance wide-area performance is limited due to cost and physical constraints, added security and management issues whilst single-site performance is not dissimilar to that of the individual component networks though some addressing and management issues need to be resolved.
(CDK p64)

b) The Internet emerged in the early 1970s via the ARPANET, an important feature of which was the development of the TCI/IP protocol suite to support remote login, file transfer and e-mail and more recently via TCP/IP based information services, the World Wide Web.
(CDK p65)

c) Internetwork protocol suites include an application, transport and internetworking (virtual) layer that is responsible for internetwork packet transmission. These layers overlaid on the underlying network of transport, network, data link and physical layers of the real networks that constitute the internetwork.
(CDK p 72)

d) For some applications, such as video, whether there is little interaction between sender and receiver, bandwidth may be the one figure of merit. However many applications are of a request-response nature and so for every large message

there must be a number of small messages. Hence other measures such as time of flight, transmission time, transport latency, sender/receiver overhead may be appropriate.

(H&P p798)

Question 3

a) **Three desirable attributes of parallel algorithms and software are: *concurrency, locality* and *modularity.* What is meant by these terms?**

(8 marks)

b) **Summarise three forms of basic model used to described parallel algorithms and indicate current trends.**

(9 marks)

c) **Describe any parallel algorithm (such as finite differences, pairwise interactions, 2D grid or any other of your choice). Use both text and diagram to aid in your description.**

(8 marks)

The least popular question, with only seven answers attempted. Most candidates gained reasonably good marks with one exceptionally good answer. Answers to part (a) revealed a lack of preparation when compared to part (c), (part (c) had been set in a previous year). Answers to part (b) omitted any mention of the abstract notions of state machines and process algebras.

Answer Pointers

a) Concurrency refer to the ability to perform many actions simultaneously, scalability indicates resilience to increasing processor counts, locality means a high ratio of local: remote memory accesses and modularity mean the decomposition of complex entities into simpler components.

(Foster 1.5)

b) Parallel algorithms may be modelled through the use of state machines and process algebras (as in CSP or LOTOS), though the trend is away from a static task and channel structure and towards a notation that facilitates a (Java) thread-like object-oriented implementation. Other alternatives include message passing (widely used), data parallelism and shared memory models.

(Magee & Kramer p32, Foster, 1.3 and Chapter notes)

c) If the candidate chooses one of the three problems quoted, it is likely that that the explanation will be as found in Foster (for Finite Difference of Pairwise Interactions or in Pfister for 2D grid. However, Foster covers seven other Case Studies (Atmosphere, Floor Plan, Chemistry, Shortest-Path, Convolution, Tuple-Space and Matrix)

Question 4

You have agreed to talk for 30 minutes at the next meeting of your local BCS branch. The title of your talk is "System Security".

Sketch out approximately eight presentation slides, with associated notes, that you would use for your talk.

(25 marks)

This was largely very well done with some candidates being copious in detail. Candidates should remember however, that this question is an exercise in sifting through their knowledge in order to extract and present the key points. For these students, being brief was problematic.

Answer Pointers
The answer should demonstrate understanding of the following:

a) Logical structure of the major s/w parts, layered in to tiers. Client, band d/b and middle tiers offering complex business processing functionality. Relate to ecommerce web systems: browser, web server, application server and several RDBMS. Diagrams and examples to clarify the answer.

b) Provides high live communication support between component applications, supports platform and application inter-operability, location transparency. Examples could include CORBA, Java, RMI or DCOM. Provision of additional services to handle transactions, security etc.

Question 5

With reference to the architectures and support technologies used to develop modern distributed software systems:

a) **What is meant by an n-tier client-server architecture? Discuss how the design of a typical e-commerce web application fits in to this architecture.**

(12 marks)

b) **Evaluate the role played in an advanced client-server system by object-based middleware and illustrate your answer with suitable examples.**

(13 marks)

Answer Pointers
The question defines the audience so the candidate should bear this in mind when answering. The number of slides indicates to the candidate that they should spend approximately 5 minutes on each slide and the candidate should remember only notes are required.

The approaches to be discussed are diverse. A good answer may favour a particular approach but should reflect the diversity.

Comments:
In a) many of the responses showed confusion between the logical concept of a client-server system, and a physical server machine, and therefore candidates found it difficult to correctly identify the separate logical building blocks of such systems. However most answers indicated some awareness of the components involved in an ecommerce system. Overall, the responses to this part of the question were disappointing. However most candidates answered b) knowledgeably and accurately. Again the answers were largely descriptive and more attempt should have been made to evaluate the advantages and disadvantages of modern object based middleware.