

CS 5523 Lecture 14: Review

- *Questions on laboratory 2*
- *Brief introduction to laboratory 3*
- *Comments on the exam*
- *Review questions Chapter 5*
- *Review questions Chapter 6*

Discussion Questions from CDK Chapter 5

[CDK 5.1] The Election interface provides two remote methods:

- *vote: with two parameters through which the client supplies the name of a candidate (a string) and the 'voter's number' (an integer used to ensure each user votes once only). The voter's numbers are allocated sparsely from the range of integers to make them hard to guess.*
- *result: with two parameters through which the server supplies the client with the name of a candidate and the number of votes for that candidate.*

Which of the parameters of these two procedures are input and which are output parameters?

Discussion Questions from CDK Chapter 5

[CDK 5.2]

Discuss the invocation semantics that can be achieved when the request-reply protocol is implemented over a TCP/IP connection, which guarantees that data is delivered in the order sent, without loss or duplication. Take into account all of the conditions causing a connection to be broken.

Discussion Questions from CDK Chapter 5

[CDK 5.3]

Define the interface to the Election service in CORBA IDL and Java RMI. Note that CORBA IDL provides the type long for 32 bit integers. Compare the methods in the two languages for specifying input and output arguments.

Discussion Questions from CDK Chapter 5

[CDK 5.4]

The Election service must ensure that a vote is recorded whenever any user thinks they have cast a vote. Discuss the effect of maybe call semantics on the Election service.

Would at-least-once call semantics be acceptable for the Election service or would you recommend at-most-once call semantics?

Discussion Questions from CDK Chapter 5

[CDK 5.5]

A request-reply protocol is implemented over a communication service with omission failures to provide at-least-once RMI invocation semantics. In the first case the implementor assumes an asynchronous distributed system. In the second case the implementor assumes that the maximum time for the communication and the execution of a remote method is T . In what way does the latter assumption simplify the implementation?

Discussion Questions from CDK Chapter 5

[CDK 5.12] A client makes remote procedure calls to a server. The client takes 5 milliseconds to compute the arguments for each request, and the server takes 10 milliseconds to process each request. The local operating system processing time for each send or receive operation is 0.5 milliseconds, and the network time to transmit each request or reply message is 3 milliseconds. Marshalling or unmarshalling takes 0.5 milliseconds per message. Calculate the time taken by the client to generate and return from two requests: (i) if it is single-threaded, and (ii) if it has two threads that can make requests concurrently on a single processor. You can ignore context-switching times. Is there a need for asynchronous RPC if client and server processes are threaded?

Discussion questions from CDK Chapter 6

CDK [6.4]

Should signal (software interrupt) handlers belong to a process or to a thread?

Discussion questions from CDK Chapter 6

CDK [6.8]

A file server uses caching, and achieves a hit rate of 80%. File operations in the server cost 5 ms of CPU time when the server finds the requested block in the cache, and take an additional 15 ms of disk I/O time otherwise.

Explaining any assumptions you make, estimate the server's throughput capacity (average requests/sec) if it is:

- i) single-threaded;*
- ii) two-threaded, running on a single processor;*
- iii) two-threaded, running on a two-processor computer.*

Discussion questions from CDK Chapter 6

CDK [6.9]

Compare the worker pool multi-threading architecture with the thread-per-request architecture.

Discussion questions from CDK Chapter 6

CDK [6.10]

What thread operations are the most significant in cost?

Discussion questions from CDK Chapter 6

CDK [6.14]

Explain the factors that motivate the hybrid scheduling approach of the 'scheduler activations' design (instead of pure user-level or kernel-level scheduling).

Discussion questions from CDK Chapter 6

CDK [6.23]

Explain the program linkage requirements that must be met if a server is to be dynamically loaded into the kernel's address space, and how these differ from the case of executing a server at user level.

Discussion questions from CDK Chapter 6

CDK [6.24]

How could an interrupt be communicated to a user-level server?

Discussion questions from CDK Chapter 6

CDK [6.25]

On a certain computer we estimate that, regardless of the OS it runs, thread scheduling costs about 50 μ s, a null procedure call 1 μ s, a context switch to the kernel 20 μ s and a domain transition 40 μ s. For each of Mach and SPIN, estimate the cost to a client of calling a dynamically loaded null procedure.