

CSIE52400 Distributed Systems, Spring 2013
Final Exam

ID: _____ Dept: _____ Name: _____

1. **(25%) Basic concepts.**

- (a) What does it mean by openness in distributed systems? How to achieve openness? Give an example of an open distributed system and briefly describe why the system can be considered as an open system.

- (b) What is QoE? What are the differences between QoE and QoS? Why do we need to consider QoE instead of just QoS?

(c) How can a distributed system be considered as scalable? What are the possible dimensions of scalability? How to achieve scalability?

(d) What are the roles of proxy and skeleton in RMI mechanism? Are there any similar components in RPC mechanism? Briefly describe the similarities and differences between them.

(e) What is the relationship between blocking/non-blocking and synchronous/asynchronous primitives? Is it possible to have a blocking asynchronous operation? What about nonblocking synchronous operation?

2. (30%) Answer the following TRUE/FALSE questions.

No.	Problem Description	TRUE/FALSE
(a)	In a layered architecture, a component can only call the component immediate above it and below it.	
(b)	In a distributed system, absolute timing of events is usually more important than relative ordering of events.	
(c)	In a synchronous distributed system, since the bounds on execution step, message transmission and clock drift are known, it is possible to use timeouts in system design.	
(d)	Actually distributed systems are mostly asynchronous distributed systems.	
(e)	For the purpose of synchronization in a distributed system, we can simply synchronize the physical clock of each process during initialization and then use the clock time for synchronization.	
(f)	In the Internet protocol layers, the physical layer and data link layer are hardware-dependent while other layers are hardware-independent.	
(g)	An indirect communication mechanism must be both time and space-uncoupled.	
(h)	Concurrent threads within the same process share the same execution environment and therefore do not need to synchronize with each others.	
(i)	In a synchronous communication, the sender must wait until the receiver process has received the message.	
(j)	In general, synchronous communication allows higher degree of parallelism while asynchronous communication leads to easier programming.	
(k)	A distributed system with higher degree of transparency is usually more user friendly than the case with lower degree of transparency.	
(l)	A cluster computing system is a type of distributed system normally with heterogeneous nodes running on high speed networks.	
(m)	The degree of transparency of a network operating system is usually much lower than that of a multiprocessor operating system.	
(n)	A message-passing system can act as a shared memory system using emulation.	
(o)	It is usually easier to program a message-passing system than a shared memory system.	

3. **(10%)** What are the protocol layers of TCP and UDP, respectively, in the OSI model? Compare the relative advantages and disadvantages of TCP and UDP from both protocol characteristics and programming points of view. Give examples to show when it is better to use TCP instead of UDP, and vice versa.

4. **(10%) RPC:** 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 (you can ignore context-switching times):

- (a) if it is single-threaded, and
- (b) if it has two threads that can make requests concurrently on a single processor.

5. **(15%)** Answer the following questions about RMI.
- (a) What are remote interfaces? Why do we need them? What are the differences between the implementation of local vs. remote methods (i.e. methods in the remote interfaces)?
- (b) In remote method invocations, what are static and dynamic invocations? Compare the relative advantages and disadvantages of each invocation mechanism.

- (c) In the RMI assignment, what would happen when a remote method is invoked by two or more clients at the same time? How do you synchronize the updates from different clients? Do we need a multithreaded RMI server?

6. (10%) Answer the following questions about indirect communication models.
- (a) In a publish-subscribe (PS) system, for each type of information below, describe the best subscription model for that type of information. If you think that a particular type of information can be best subscribed by more than one model, briefly explain why it is the case.

Desired information	Subscription model(s)
Scores of 2013 NBA finals	
News articles on global warming	
All stocks with price > \$50 and keep rising for at least two days	
All information about the CSIE department of NDHU	
Sales advertisement of iPad mini around Hualien	

- (b) In a tuple space (TS) model, briefly describe how to retrieve the type of information subscribed by each subscription model in the PS system? Use the examples in the table above to demonstrate your methods.