Dept. Comp. Science **Midterm Distributed Sys.**
Vrije Universiteit **25.10.2005**

1a Give an example of a 3-tiered client-server architecture. *5pt*

One, nowadays very common example, is the organization of a Web site. It consists of a Web server (first tier), which talks to an application server (second tier), which, in turn, communicates with a database server (third tier). Other examples include a clint machine with the user interface communicating with an application server, in turn communicating with a database server.

1b Despite the fact that multi-tiered architectures do not really solve any problems inherent to distributed systems, they have one practical advantage. What is this advantage? *5pt*

These architectures considerably simplify systems management by physically decoupling the logical application layers across multiple machines.

2a What is the difference between synchronous and isochronous transmission mode? *5pt*

In synchronous transmission mode, only a maximum end-to-end delay for each unit of a data stream is given. This is useful for real-time sensor data to guarantee that observations are accurate. In isochronous transmission mode, there is also a minimum end-to-end delay specified generally to ensure that a receiver can process incoming data units in real time. It is primarily used for multimedia streaming.

2b What is meant by stream synchronization? *5pt*

Stream synchronization refers to the fact that packets of two or more different (continuous or discrete) data streams should be delivered in a specific order at the receiver. So, for example, if PDU[i,j] denotes the jth packet from stream i, synchronization may dictate that packets are to be delivered in the order PDU[1,1], PDU[2,1], PDU[1,2], PDU[2,2], PDU[1,3], PDU[2,3], etc.

3a What is the difference between an iterative and a concurrent server? *5pt*

An iterative server handles a single request at a time, mostly by means of a single thread. A concurrent server can handle multiple requests at the same time, often by means of one thread that picks up incoming requests in order to pass it on to another thread that completes the request processing.

3b Explain how you would organize an object server so that it can *simultaneously* support the iterative and concurrent way of handling requests (for different objects). *5pt*

The key to this answer are object adapters with their activiation policy. By simply having one adapter follow the policy that a request is to be handled by the adapter, and another that it spawns a new thread on every request, both types of request handling can be handled.

4a Explain how Lamport's way of adjusting logical clocks works. Be precise! *5pt*

You need to explain that every process P_i maintains a local counter C_i that is incremented at every event, notably when sending and receiving a message. When process P_i sends a message m, it receives a timestamp $T(m) = C_i$. When m is received by P_j, the latter sets its local clock to $\max\{C_j, T(m)\} + 1$.

4b Explain how Lamport timestamps can be used to realize totally ordered multicasting. *10pt*

Grading: *The final grade is calculated by accumulating the scores per question (maximum: 45 points), and adding 5 bonus points. The maximum total is therefore 50 points.*