Department Computer Science	**Distributed Systems**
VU University	**19.12.2008**

1 Distributing data and processes may help to address size scalability, but may easily introduce geographical scalability. Give a well-known example to illustrate this point, as well as a solution. *5pt*

2a Processes can be decoupled in time and space. For each combination of decoupling (see figure), characterize the type of distributed system through an example. *5pt*

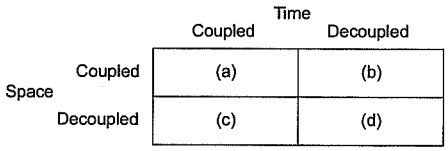

2b Explain the concept of interceptors as used in middleware and why they can be useful. *5pt*

2c Explain how a feedback control loop works by providing an example of a distributed system that fits the following figure. Explain each component, as well as each connection to/from a component. *10pt*

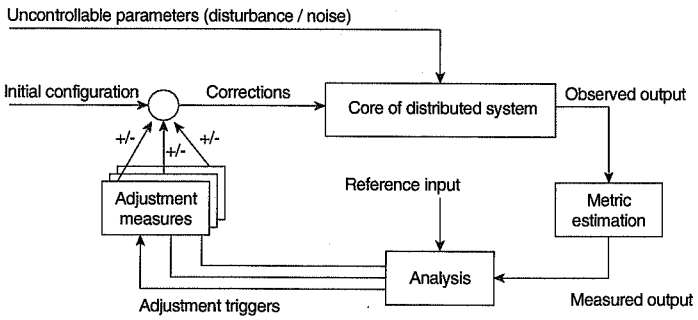

3a Unlike local pointers, having systemwide object references helps to improve access transparency in RPCs. How can such object references be implemented? *Hint: think of how Java realizes remote method invocations.* *5pt*

3b The figure below shows how an RPC system works in practice. Explain what is in the runtime library. *5pt*

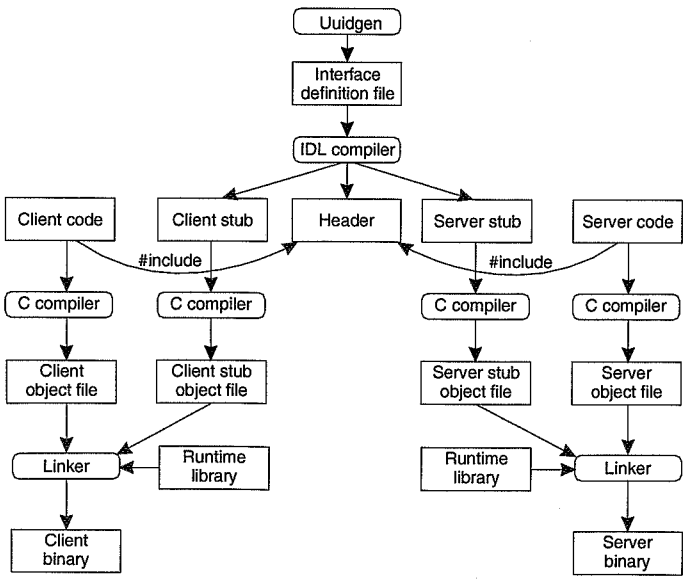

3c In RPC, a client needs to **bind** to a server. What does this mean and how can it be realized? *5pt*

4a Explain how name resolution works in Chord by resolving $k = 30$ starting from node 21 in the following example. Do the same for $k = 19$ from 21. *5pt*

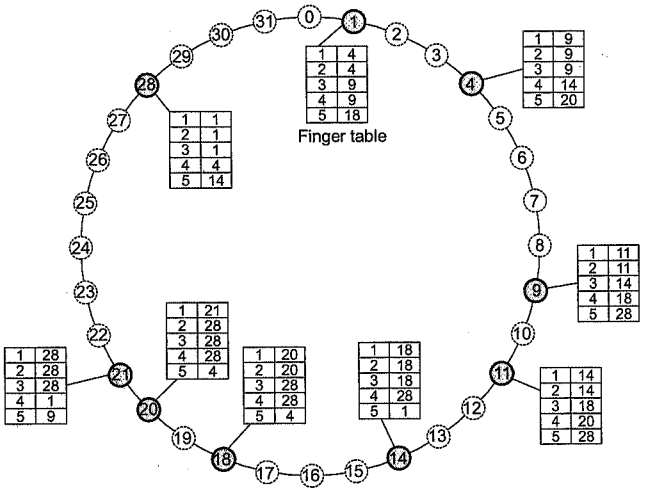

4b In Chord, the finger table entry $FT_p[i]$ of peer p is equal to $succ(p + 2^{i-1})$. Explain how Chord's finger tables can be extended to incorporate proximity routing. *5pt*

4c As in any other naming system, it is possible to look up up a key in Chord recursively or iteratively. Explain the differences, as well as some advantages and disadvantages of either approach. *5pt*

4d Give two approaches to using Chord for implementing a distributed file system. *5pt*

5a Explain how a blocking primary-backup protocol works, as well as its nonblocking variant. Why is a primary-backup protocol in which the primary moves to the location of the writer, never blocking? *5pt*

5b Sketch the design of a simple, centralized consistency protocol for active replication that realizes sequential consistency. *5pt*

6a One can argue that NFS is not really a file system. Explain why. *5pt*

6b Mention two different measures that were designed into NFS version 4 in order to let it operate better in wide-area networks. *5pt*

7a Explain how Akamai uses standard Web-caching techniques to effectively implement server-initiated replication of Web pages. *5pt*

7b Explain how a content-aware cache works in edge-server systems that support replication of Web applications. *5pt*

Grading: *The final grade is calculated by accumulating the scores per question (maximum: 90 points), and adding 10 bonus points. The maximum total is therefore 100 points.*