

TECHNISCHE UNIVERSITEIT EINDHOVEN
Faculteit Wiskunde en Informatica

*Examination Architecture of Distributed Systems (2IMN10 / 2II45),
on Monday November 2, 2015, from 13.30 to 16.30 hours.*

Indicate on your work whether you take this exam in fulfillment of course 2IMN10 or of course 2II45.

Before you start, read the entire exam carefully. Answers to all questions must be motivated and stated clearly. For each question the maximum obtainable score is indicated between parentheses. The total score sums up to 20 points. This is a closed book exam, i.e., you are not allowed to use books or other lecture material when answering the questions.

1. (2 points) Describe the Virtual Machine architectural style using the appropriate vocabulary, name the concepts involved, give a motivation for its usage and mention typical usage.

Answer. See slide 26 of the slide set on architectural styles.

2. (2 points) Describe two approaches to locate mobile entities. For each approach give an example where it is used and indicate the extent to which it provides location transparency.

Answer. Two distinct approaches are:

Pointer forwarding via stubs This is done for remote method invocations to mobile objects. When an object moves to a new server, it leaves behind a client stub at its old location and installs a server stub at its new location. As long as the client stub that is left behind forwards its own location as the reply address, this provides location transparency. If, however, the client stub forwards the address of the original caller as the reply address, location transparency is broken.

Contact via home location This approach is taken by Mobile IP. Here clients contact a mobile host through a home agent that is located at a fixed a priori known home address. Mobile hosts, on their journey, register at so-called foreign agents, which in turn notify the home agent of the current location (the care-off address) of the mobile host. Home agents use tunneling of IP packets to forward messages from clients to hosts. The mobile host may answer via the home agent or directly to the client thus creating triangular routing. In the first case there is complete location transparency. In the latter case, the mobile host uses its home address in its answers, which is noticeable by routers. Alternatively, the home agent may inform the client of the host's care-off address, in which case location transparency is also broken.

3. (2 points) Describe in detail how the URL

`http://www.win.tue.nl/home/wsinmak/Education/2IMN10/ADS.html`

is resolved. In particular, indicate the closure mechanisms, and resolution procedure for the various parts of the URL.

Answer. This URL consists of three components: a scheme **http** that names a protocol, a hostname **www.win.tue.nl** that identifies the host holding the resource, and a pathname **home/wsinmak/Education/2IMN10/ADS.html** that identifies a file (the resource) on that host. It is resolved as follows. The browser extracts the scheme and based on the value found invokes its `Http_client`, which is a plugin in the form of a library. This `Http_client` resolves the hostname by contacting its local DNS-server to obtain an IP-address. The local DNS-server may invoke other DNS-servers to assist in resolving the hostname (see TvS for a description of how this works), but in case the local DNS-server happens to be the TU/e DNS-server it will know the IP-address. Next, the `Http_client` will set up a connection with the host machine, construct an appropriate `http-request` (GET) containing the pathname and invoke the operating system to send the request to the web server at the host machine. Finally, this web server resolves the pathname to obtain the resource.

4. Modifiability is a quality attribute of an architecture.

- (a) (1 point) Which concern is addressed by this quality attribute and what metric is used to measure it?

Answer. Modifiability is concerned with the cost of system change. It is primarily determined by the man-power needed to make the change. Ultimately, however, the cost is expressed in money, which may encompass more than just the cost of man-power.

- (b) (1 point) Name at least two generic tactics to achieve modifiability. For each of them, mention an architectural style that applies the tactic.

Answer.

Localize modification, i.e., keep the number of modules that need to be changed small. The MVC-pattern does so by separating the state, described in the Model, from its representations described by its Views.

Defer binding, to control deployment time and cost of reconfiguration. The SOA-style does so by building applications at runtime from reusable services implemented through components that are independent units of deployment and composition.

5. Consider the Chord scheme for DHTs. Assume a 7-bit identifier space, and assume that the node set N is given by $id(N) = \{8i \mid 0 \leq i < 16\}$.

- (a) (0.5 point) Give the finger table of node 32.

Answer. For a 7-bit identifier space all finger tables have 7 entries. Table FT_{32} is given by:

$FT_{32}[1] = FT_{32}[2] = FT_{32}[3] = FT_{32}[4] = 40$, $FT_{32}[5] = 48$, $FT_{32}[6] = 64$, and $FT_{32}[7] = 96$.

- (b) (1.0 point) What is the maximum number of steps necessary to resolve a key? For your answer, you may assume that each node is aware of the identity of its predecessor and, consequently, resolves all keys for which it is responsible in zero steps.

Answer. Each time the resolution process forwards the lookup query from node n to node $(n + x) \bmod 2^7$, the increment x satisfies $x \in \{8, 16, 32, 64\}$. Furthermore, in any sequence of forwarding operations, the increments are almost always decreasing, with the exception of the last two increments which may both be 8. For $x = 64$ the next increment is smaller, because otherwise the lookup query travels at least once around the ring (since there are only $2^7 = 2 * 64$ identifiers) which cannot be the case. For $x \in \{16, 32\}$, the next increment is smaller, because otherwise the query would have been forwarded from n to $(n + 2x) \bmod 2^7$. So all decreasing sequences have length at most 4, the only exception being the sequence 64, 32, 16, 8, 8. However, this sequence cannot occur, since $64 + 32 + 16 + 8 + 8 = 128$ which would mean that the lookup has been forwarded once around the ring. Therefore, the maximum number of steps is 4.

- (c) (0.5 point) Give a key that requires the maximum number of steps to be resolved, when starting the resolution at node 32. Indicate the nodes that are visited when resolving the given key.

Answer. Starting at node 32, any key k with $17 \leq k \leq 23$ will need 4 steps to be resolved. For $k = 17$ the path traversed is $32 \rightarrow 96 \rightarrow (128 = 0) \rightarrow 16 \rightarrow 24$. Note that the algorithm described in TvS also takes 4 steps for key 24, but since node 32 knows that node 24 is its predecessor, it could take advantage of that knowledge and forward the query directly to its predecessor.

6. (2 points) Protocols used for updating a cache from a central data store can be classified either as push-based or as pull-based. Explain the difference and indicate for each protocol type a situation in which it is the preferred one.

Answer. See TvS, Section 7.4.3. pp. 303-305.

7. Often, replication is applied in a system architecture to meet quality requirements.

- (1 point) Name four quality drivers for the usage of replication.

Answer. Four important drivers are: *availability, reliability, performance and scalability*.

- (1 point) For each of the named qualities indicate a concrete context in which it is used, and explicitly identify what is replicated.

Answer.

availability and reliability. Redundancy allows switching to a backup server in case the primary server fails. This increases both availability and reliability. In this case both the primary server and its state need to be replicated. Another reason to temporarily switch to a backup server is to do periodic maintenance or update servers. So maintainability, updatability and serviceability can also be seen as quality drivers.

performance. By replication of processing elements and distributing, but not necessarily replicating, the data to be processed, performance, in particular throughput, is increased by concurrency. As another example, sending a query for domain name resolution to the geographically nearest DNS root server (replica), diminishes the round trip communication time and therefore response time.

scalability. In any P2P system, each peer, in its role as client, increases system load, but also, in its role as server, provides resources. Hence the balance between system load and system capacity (i.e., amount of resources) is maintained, and thus scalability is realized.

8. Indicate for the following statements whether they are true or false. Motivate your answer with a short argument.

(a) (0,75 point) Concurrency issues are best addressed in the development view.

False. Concurrency is addressed in the process view, where interaction between processes is a concern. It is also visible in the deployment view, because there it can be seen that multiple physical nodes are deployed to provide a single service.

(b) (0.75 point) All DNS servers implement iterative name resolution.

True. Iterative resolution is the default way of doing resolution. DNS servers at the administrative and managerial level often also provide recursive resolution, but the root servers will never do so.

(c) (0.75 point) An adapter is an architectural element that provides life cycle management.

True. Adapters can be used to create and destroy objects, to activate or deactivate processes, etc. All of these are considered life cycle management activities.

(d) (0.75 point) Using the blackboard architectural style provides referential decoupling.

True. In the blackboard style, multiple knowledge sources (KSs) interact to solve a problem without being aware of each other's existence or identity. A separate controller that is aware of their capabilities schedules their activities.

(e) (0.75 point) Systematic usage of asynchronous RPCs improves scalability of a system through latency hiding.

True. Usage of Asynchronous RPCs allow a client to continue with other work

while waiting for a reply. The time that a client would otherwise have been waiting includes the round-trip time for request and reply, which is therefore hidden. As a system grows in number of processing element, in all likelihood also the communication between those elements will grow, but since asynchronous RPCs are used systematically, all communication will benefit from latency hiding. Hence the effect of communication overhead, which in general hampers scalability, is kept to a minimum.

- (f) (0.75 point) An architectural viewpoint is a collection of models
False. A viewpoint is a collection of patterns, templates, and conventions for constructing models that address a specific (set of) concern(s). In addition, a viewpoint identifies stake-holders that have an interest in those concerns. A collection of models constructed according to the principles and guidelines of a particular viewpoint, on the other hand, is called a view.
- (g) (0.75 point) According to Bondi, a system has structural scalability, if it is possible to continuously extend its resources to cope with ever increasing load.
False. According to Bondi, a system has structural scalability, if its implementation and standards do not impede the growth of the number of objects it encompasses.
- (h) (0.75 point) Anycasting is only used in local area networks.
False. Anycast is, e.g., used to contact a DNS root server. It is implemented in the network layer by the Border Gateway Protocol that routes traffic *between* networks.