

Object interaction using RMI



INF 5040 høst 2006

Lecturer: Frank Eliassen

Frank Eliassen, SRL & Ifi/UiO

1

Plan

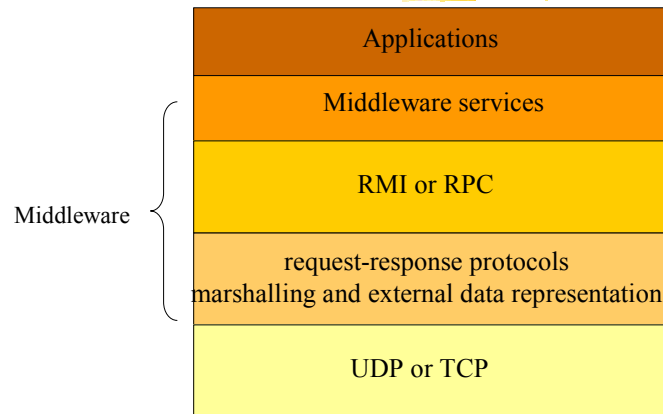


- Principles for realising remote methods invocations (RMI)
- Objekt-servers
- Multi-threaded object servers
- CORBA RMI
- Java RMI

Frank Eliassen, SRL & Ifi/UiO

2

Layering of middleware



Frank Eliassen, SRL & Ifi/UiO

3

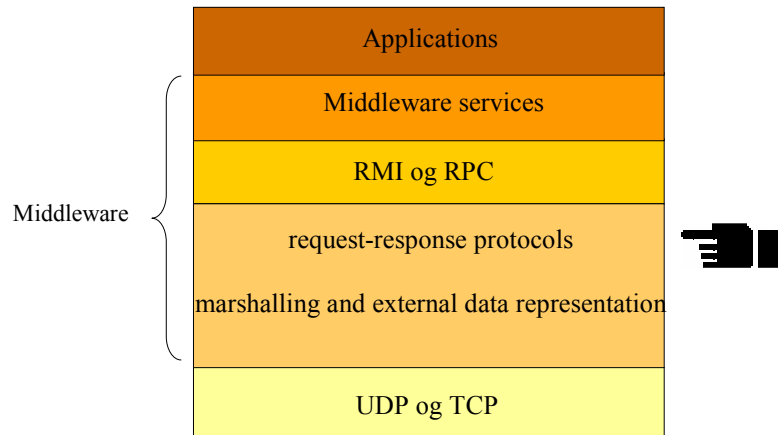
Plan

- Principles for realising remote methods invocations (RMI)
 - Objekt-servers
 - Multi-threaded object servers
 - CORBA RMI
 - Java RMI

Frank Eliassen, SRL & Ifi/UiO

4

Layering of middleware

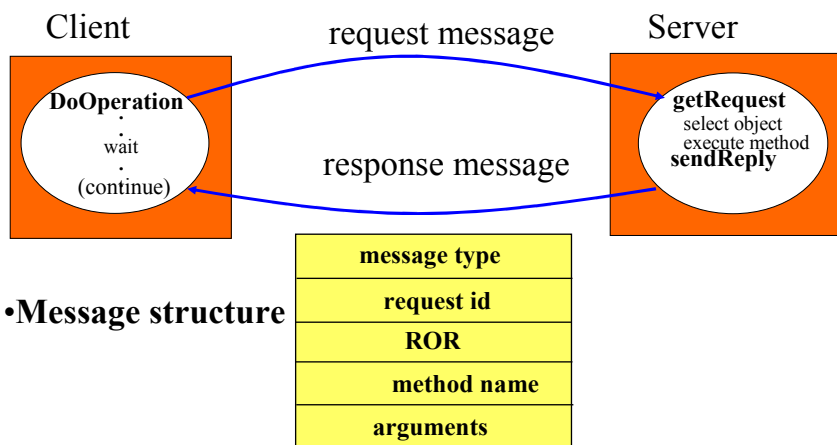


Frank Eliassen, SRL & Ifi/UiO

5

Client-server communication RPC/RMI exchange protocols

- Based on UDP or TCP



Frank Eliassen, SRL & Ifi/UiO

6

Failure model for RPC/RMI exchange protocols over UDP (I)

- Protocol can be exposed to
 - omission failure
 - process crash failure
 - message order not guaranteed
- Failure is detected as *timeout* in the primitive DoOperation:
 - recovery actions depend on the offered delivery guarantee

Frank Eliassen, SRL & Ifi/UiO

7

Failure model for RPC/RMI exchange protocols over UDP (II)

- Timeout DoOperation
 - Send request message repeatedly until
 - response is available, or
 - assume server has failed (max no of retrans.)
- Duplicate request messages
 - occur when request message is sent more than once
 - can lead to operations being executed more than once for the same request
 - => must be able to filter duplicate requests (role of request id)
- Lost response messages
 - server has already sent response message when it receives a duplicate request message
 - => may have to execute the operation again to get the right response
 - OK for operations that are “idempotent”

Frank Eliassen, SRL & Ifi/UiO

8

Failure model for RPC/RMI exchange protocols over UDP(III)

- Logs (histories):
 - used by servers offering operations that are not "idempotent"
 - contains response messages already sent
- Disadvantage of logs:
 - storage requirement
- if a client is allowed to do only one request at a time to the same server, the log can be limited in size (bounded by the number of concurrent clients)
- at reception of the next request message from the same client the server may delete the last response message for that client from the log

Frank Eliassen, SRL & Ifi/UiO

9

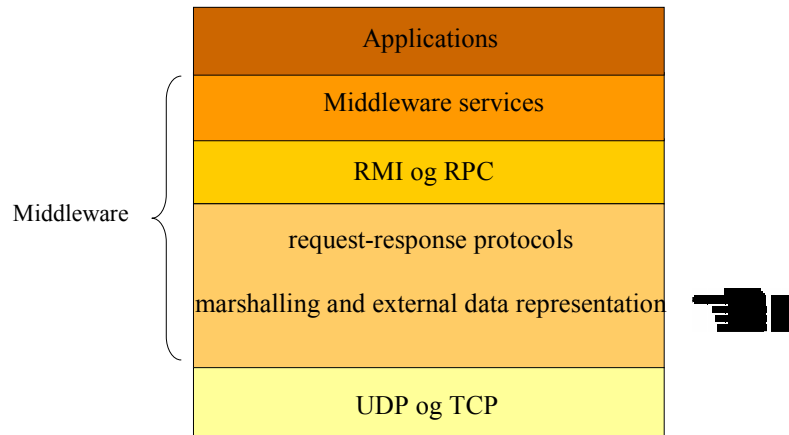
RPC/RMI exchange protocols over TCP

- UDP has limited packet size
 - => need for fragmentation/defragmentation protocols
- RPC/RMI exchange protocols over TCP avoids this problem
 - TCP ensures reliable delivery of byte streams
- Problem:
 - much overhead if the connection has to be created at each request
 - => need for optimization (leave connection open for later reuse)
 - upper bound on number of concurrent TCP-connections could cause problems

Frank Eliassen, SRL & Ifi/UiO

10

Layering of middleware



Frank Eliassen, SRL & Ifi/UiO

11

Marshalling External data representation

- "marshalling"
 - serialize data structures to messages (sequence of data values)
 - translate sequence of data values to an external representation
- "unmarshalling"
 - inverse of "marshalling"
- External data representation (representation "on the wire")
 - a representation of data during transfer of the message
 - Sun XDR (representation of most used data types)
 - ASN.1/BER (ISO standard, based on "type-tags", open)
 - NDR (used in DCE RPC)
 - CDR (used in CORBA RMI, binary layout of IDL types)
 - Java Object Serialization (JOS)
 - XML (used i SOAP: "RMI" protocol for Web Services)

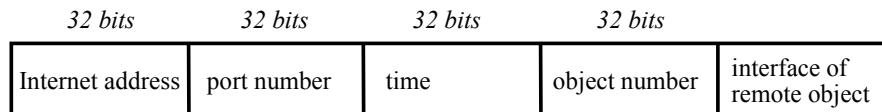
Frank Eliassen, SRL & Ifi/UiO

12

Marshalling Object-references

➤ Remote-object-reference (ROR)

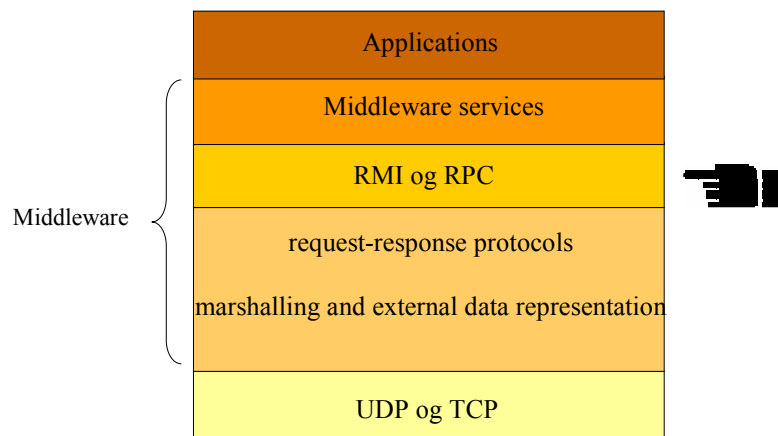
- Identifier for remote objects that is valid in a distributed system
- Must be generated in a way that ensures uniqueness over time and space (=> a ROR can not be reused)
- Example:



Frank Eliassen, SRL & Ifi/UiO

13

Layering of middleware



Frank Eliassen, SRL & Ifi/UiO

14

Classification of RPC/RMI exchange protocols

- Classification after (Spector, 1982):
 - basis for implementing different types of RMI and RPC
- Request (R) protocol
 - Only Request-message. No response message from server
 - No confirmation that operation has been performed
- Request-Reply (RR) protocol
 - Reply-message confirms that the Request-message has been performed
 - A new request from the client confirms reception of Reply-message
- Request-Reply-Acknowledge (RRA) protocol
 - separate message from client to confirm reception of Reply-message
 - tolerates loss of Ack-message
 - Ack with a given *request id* confirms all lower requests ids

Frank Eliassen, SRL & Ifi/UiO

15

RPC/RMI invocation semantics

- Reliability semantics of RPC/RMI under partial failures

	Fault tolerance measures		Invocation semantics
Retransmit Request message	Duplication filtering	Re-execute method <i>or</i> retransmit Reply	
No (R)	—	—	Maybe
Yes (RR)	No	Re-execute method	At-least-once
yes (RR)	Yes	Retransmit Reply	At-most-once

Frank Eliassen, SRL & Ifi/UiO

16

RMI invocation semantics in object and RPC middlewares

- RMI in CORBA and Java have "at-most-once" invocation semantics under partial failures
 - referred to as synchronous requests
- CORBA allows other forms of synchronization that provides other invocation semantics
 - One-way operations: maybe-semantics
- SUN RPC: at-least-once semantics

Frank Eliassen, SRL & Ifi/UiO

17

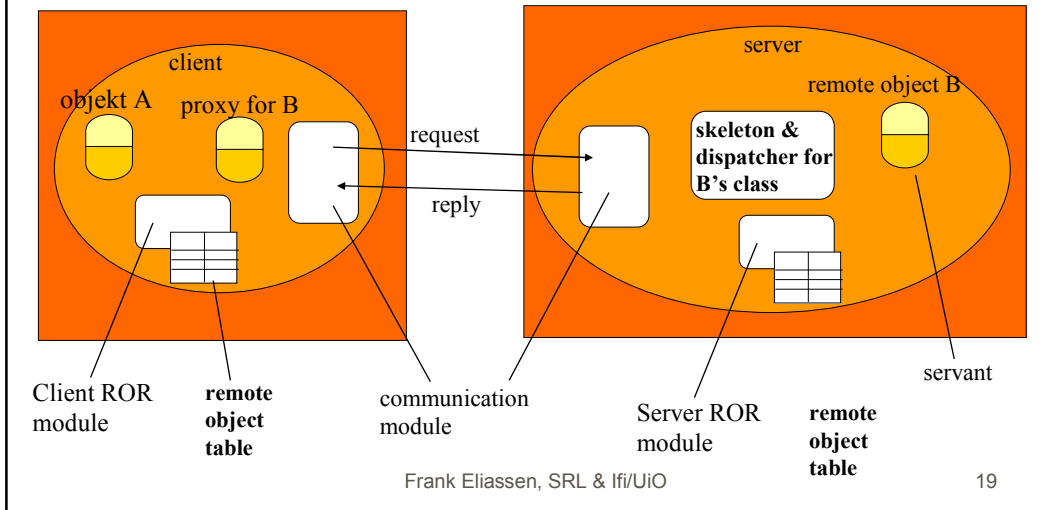
Implementation of RMI

- Three main tasks :
- Interface processing
 - Integration of the RMI mechanism into a programming language.
 - Basis for realizing access transparency
- Communication
 - message exchange (request-reply protocol)
- Object location, binding and activation
 - Locate a suitable server process that hosts the remote object and bind to the server
 - Activate an object-implementation
 - Basis for realizing location transparency

Frank Eliassen, SRL & Ifi/UiO

18

Objects and modules involved in RMI



Elements of the RMI software (I)

➤ RMI interface processing: Client proxy

- Local "proxy" object for each remote object a client holds a ROR ("stand-in" for remote object).
- The class of the proxy-object has the same interface as the class of the remote object. Can perform type checking on arguments
- Performs marshalling of requests and unmarshalling or responses
- Transmits request-messages to the server and receive response messages.

Elements of the RMI software (II)

- RMI interface processing: Dispatcher
 - A server has one dispatcher for each class representing a remote object.
 - Receives requests messages
 - Uses *method id* in the request message to select the appropriate method in the skeleton and passes on the request message

Frank Eliassen, SRL & Ifi/UiO

21

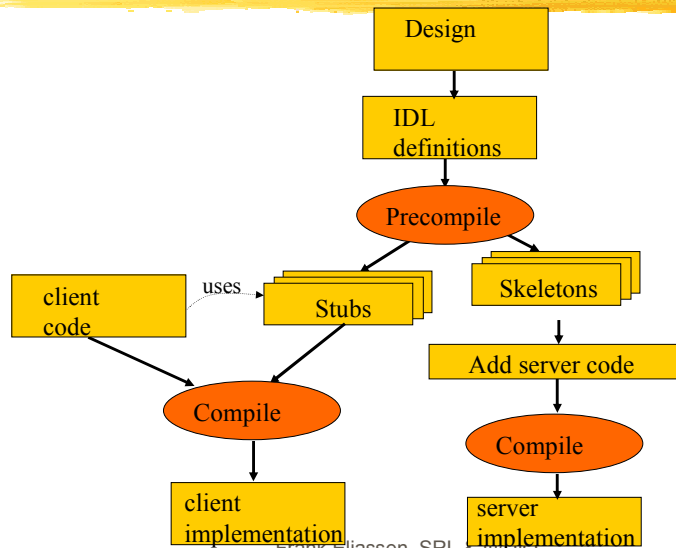
Elements of the RMI software (III)

- RMI interface processing: Skeleton
 - A server has one skeleton for each class representing a remote object
 - Implements the methods in the remote interface
 - A skeleton method unmarshals the arguments in the request message and invokes the corresponding method in the servant.
 - It waits for the invocation to complete and then marshals the result, together with any exceptions, in a reply message to the sending proxy's method.

Frank Eliassen, SRL & Ifi/UiO

22

Generation of proxies, dispatchers and skeletons



Frank Eliassen, SRL & Ifi/UiO

23

Server and client programs

- Server program contains
 - the classes for the dispatchers and skeletons
 - the implementation classes of all the servants that it supports
 - an initialization section: creates and initializes at least one servant
 - additional servants may be created in response to client requests
 - register zero or more servants with a *binder*
 - potentially one or more *factory methods* that allow clients to request creation of additional servants
- Client program contains
 - the classes and proxies for all the remote objects that it will invoke

Frank Eliassen, SRL & Ifi/UiO

24

RMI binding, object location, and activation

- Binding in RMI
 - corresponds to mapping a symbolic object name to an ROR
 - performed by a "binder" (name service or similar)
- Object location in RMI
 - corresponds to mapping an ROR to a communication identifier
- Activation in RMI
 - corresponds to creating an active object from a corresponding passive object (e.g., on request). Performed by an *activator*
 - register passive objects that are available for activation
 - activate server processes (and activate remote object within them)

Frank Eliassen, SRL & Ifi/UiO

25

Object location

- Object location : corresponds to mapping an ROR to a communication identifier.
 - integrated in ROR
 - Address can be extracted directly from the object reference
 - location service
 - A name server is used by the client proxy at each request
 - cache/broadcast
 - Each client has cache of bindings (ROR, comm. identifier)
 - If ROR not in cache, perform broadcast with ROR
 - Servers that host the object respond with comm.identifier
 - forward pointers or address hint (to e.g., location service)
 - Used at object migration
 - Combinations of the above

Frank Eliassen, SRL & Ifi/UiO

26

Plan

- Principles for realising remote methods invocations (RMI)
- ➤ Object-servers
 - Multi-threaded object servers
 - CORBA RMI
 - Java RMI

Frank Eliassen, SRL & Ifi/UiO

27

Object-server: Server tailored to support distributed objects

- Services realized as objects that the server encapsulates
 - Services can be added or removed by creating and removing remote objects
- Object servers act as places where objects can live
- Object servers activate remote objects on demand
 - Several ways to activate an object

Frank Eliassen, SRL & Ifi/UiO

28

Object servers must assign processing resources to objects when they are activated

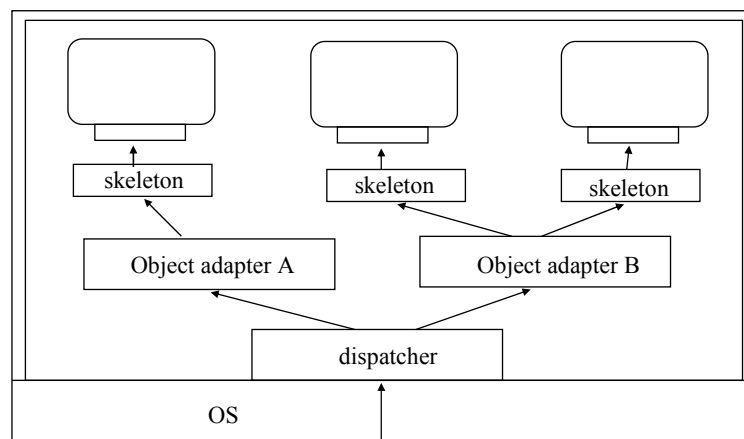
- When an object is activated, which processing resources should be assigned to the implementation?
- Activation policy
 - A particular way of activating an object
 - Different dimensions
 - How to translate between ROR and local implementation?
 - Should the server be single-threaded or multi-threaded?
 - If multi-threaded, how to assign threads to objects and requests?
One thread per object? One per request?
 - Transient vs persistent objects, etc
- No single activation policy that fits all needs
 - Object servers should support several concurrent activation policies
 - Objects can be grouped according to which activation policy they are governed by

Frank Eliassen, SRL & Ifi/UiO

29

Organization of object servers that support different activation policies

- **Object-adapter:** software that implements a specific activation policy (supported by CORBA Portable Object Adapter (POA))



Frank Eliassen, SRL & Ifi/UiO

30

Plan

- Principles for realising remote methods invocations (RMI)
- Objekt-servers
- ➤ Multi-threaded object servers
- CORBA RMI
- Java RMI

Why use several processes or threads in distributed systems?

- Better performance
- Better exploitation of resources

Object-servers must assign processing resources to objects when objects are activated

- When an object is activated, which processing resources should be assigned to its implementation?
 - Create a new process or thread?
 - Are there several ways this can be done?
 - Is there a best way (cf. activation policies)?

Frank Eliassen, SRL & Ifi/UiO

33

Threads are cheaper to manage than processes

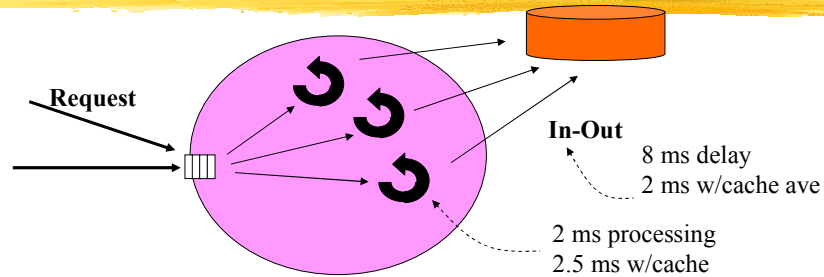
- execution environment (address room) expensive to create
- threads are much cheaper to create and terminate
- threads can share runtime environment (e.g., global variables, open files, timers, signals,)

- BUT: processes define *protection domains*
 - a thread in a process can (normally) not access data and resources in an other process

Frank Eliassen, SRL & Ifi/UiO

34

Multi-thread server

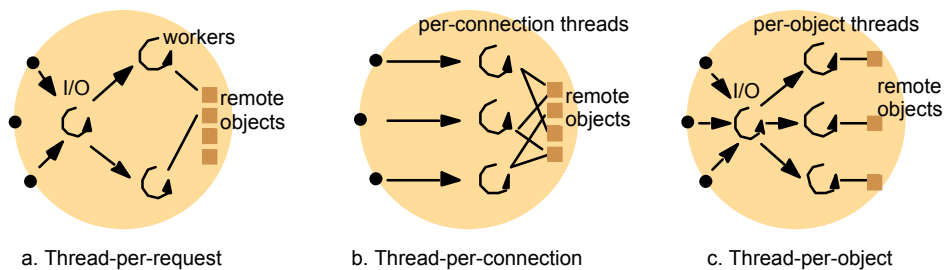


	# processors	# threads	Max. calls/sec
no disk caching	1	1	100
no disk caching	1	2	125
disk caching	1	2	400
disk caching	2	2	444
disk caching	2	3	500

Frank Eliassen, SRL & Ifi/UiO

35

Alternative threading-policies for object activation



a. Thread-per-request

b. Thread-per-connection

c. Thread-per-object

Frank Eliassen, SRL & Ifi/UiO

36

Plan

- Principles for realising remote methods invocations (RMI)
- Objekt-servers
- ➤ CORBA RMI
- Java RMI
- Multi-threaded object servers

Frank Eliassen, SRL & Ifi/UiO

37

CORBA middleware

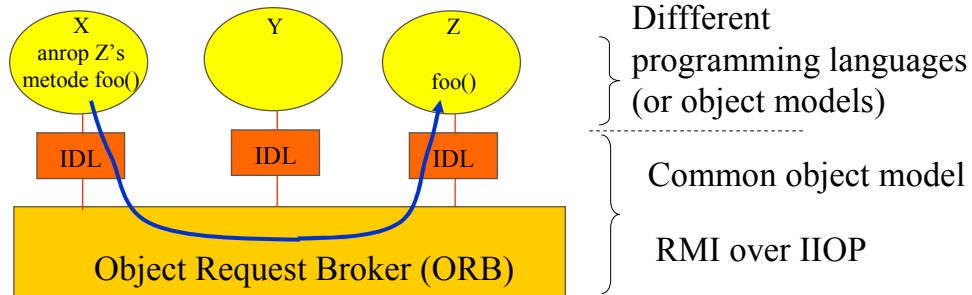
- Offers mechanisms that allow objects to invoke remote methods and receive responses in a transparent way
 - location transparency
 - access transparency
- The core of the architecture is the Object Request Broker (ORB)
- Specification developed by members of the Object Management Group (www.omg.org)

Frank Eliassen, SRL & Ifi/UiO

38

CORBA RMI

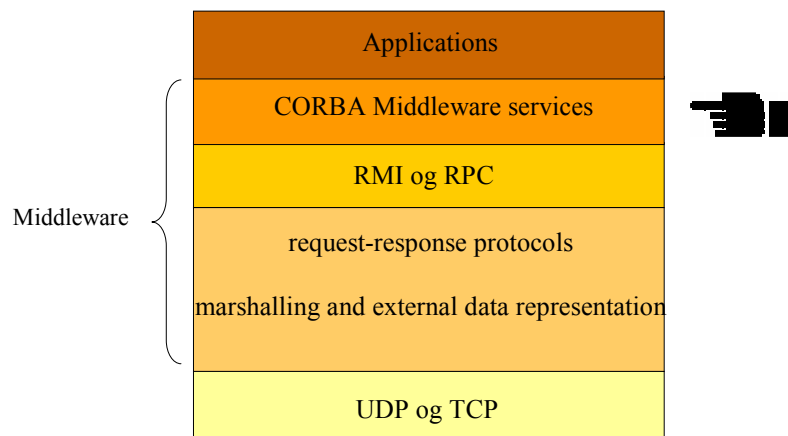
Clients may invoke methods of remote objects without worrying about:
object location, programming language,
operating system platform, communication
protocols or hardware.



Frank Eliassen, SRL & Ifi/UiO

39

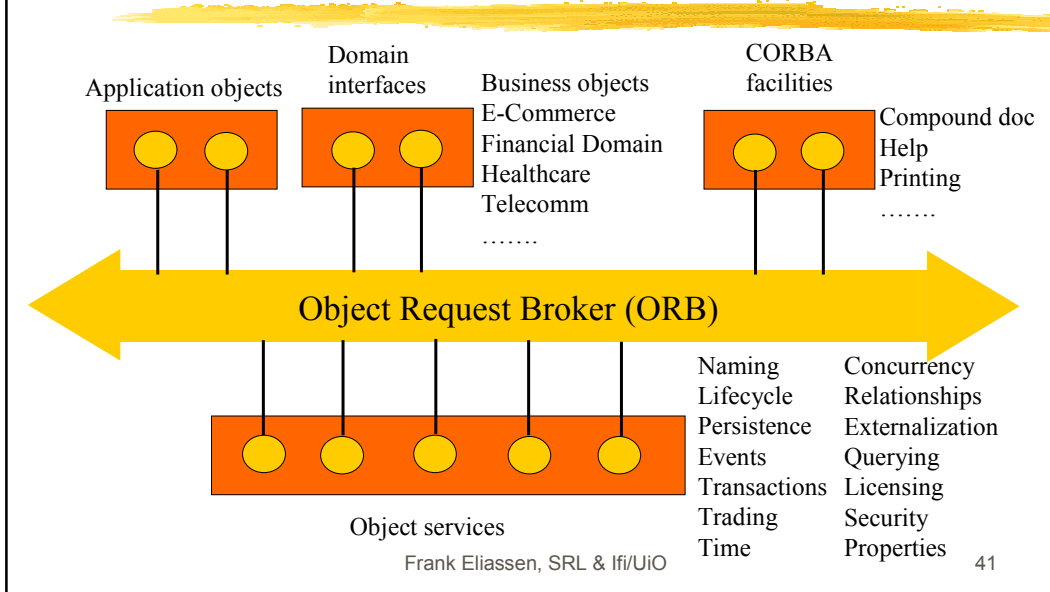
Layering of middleware



Frank Eliassen, SRL & Ifi/UiO

40

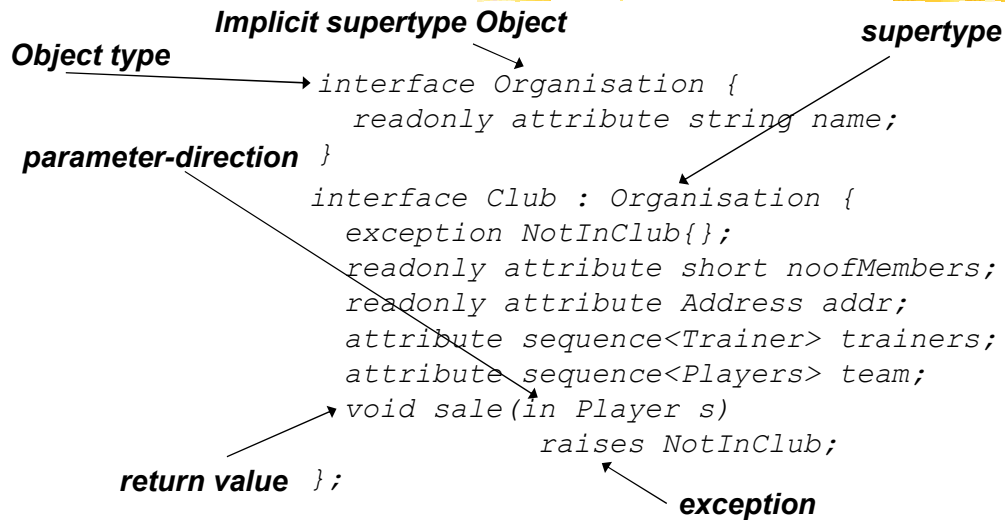
CORBA Services



CORBA IDL

- Language for specifying CORBA object types
- Can express all concepts in the CORBA object model
- OMG/IDL is
 - not dependent on a specific programming language
 - syntactically oriented towards C++
 - not computationally complete
- Different bindings to programming languages available

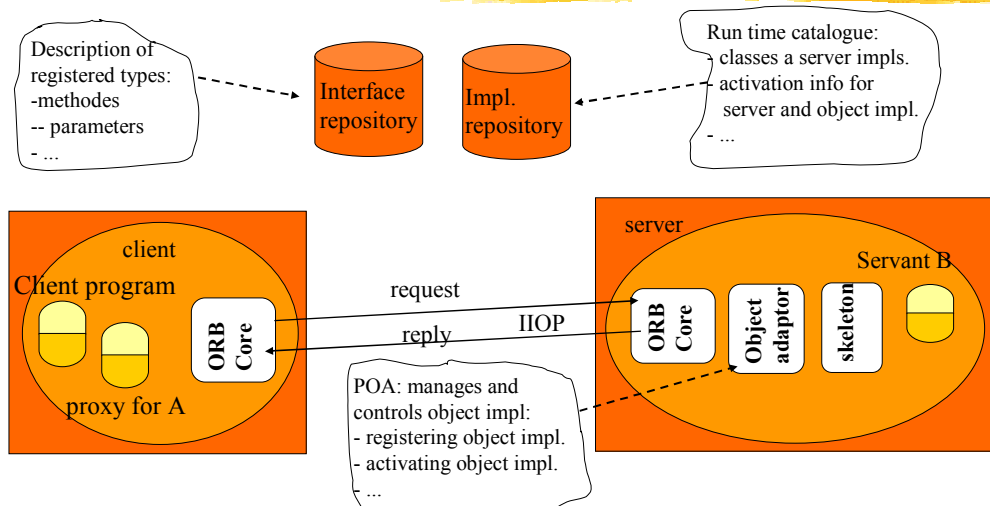
CORBA object model and IDL



Frank Eliassen, SRL & Ifi/UiO

43

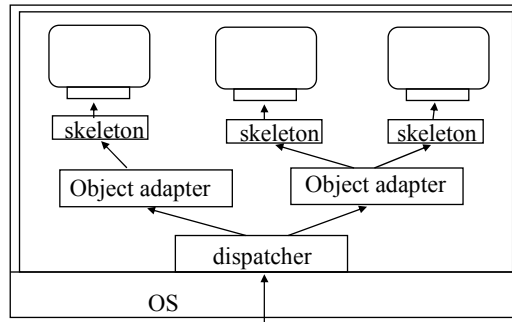
CORBA architecture



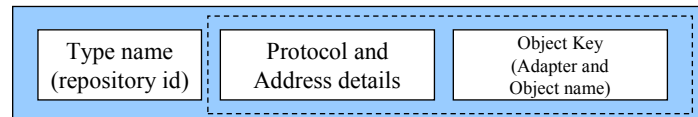
Frank Eliassen, SRL & Ifi/UiO

44

CORBA ROR format (IOR) reflects the organization of object servers



CORBA Interoperable Object Reference (IOR)



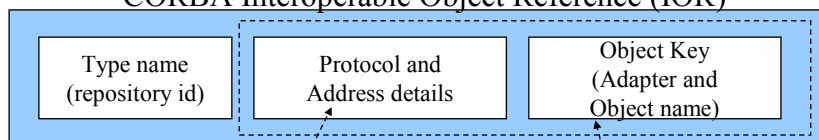
Frank Eliassen, SRL & Ifi/UiO

45

CORBA RMI Object Location (I)

- Object location in RMI corresponds to mapping object references (ROR) to "servants"
 - servant: implementation of one or more CORBA objects
- ROR in CORBA: Interoperable Object Reference (IOR)
- Location process:
 - based on information encoded in the object reference

CORBA Interoperable Object Reference (IOR)



IOP: host name/port no Proprietary format (of the ORB that created the IOR)

Frank Eliassen, SRL & Ifi/UiO

46


CORBA RMI Object Location (II)

- Transient IOR
 - Valid only as long the corresponding server process is available
 - After the server process has terminated, the IOR will never be valid again
 - The location of the server process is encoded into the IOR of the object.
- Persistent IOR
 - continue to function (denote same CORBA object) even when the server process terminates and later starts up again

Frank Eliassen, SRL & Ifi/UiO

47

Plan

- Principles for realising remote methods invocations (RMI)
- Objekt-servers
- Multi-threaded object servers
- CORBA RMI
-  ➤ Java RMI

Frank Eliassen, SRL & Ifi/UiO

48

➤ Java Remote Method Invocation (RMI)



Frank Eliassen, SRL & Ifi/UiO

49

Java RMI

- Remote Method Invocation (RMI) supports communication between different Java Virtual Machines (VM), and possibly over a network
- Provides tight integration with Java
- Minimizes changes in the Java language/VM
- Works for homogeneous environments (Java)
- Clients can be implemented as *applet* or *application*

Frank Eliassen, SRL & Ifi/UiO

50

Java Object Model

- Interfaces and Remote Objects
- Classes
- Attributes
- Operations/methods
- Exceptions
- Inheritance

Frank Eliassen, SRL & Ifi/UiO

51

Java interfaces to remote objects

- Based on the ordinary Java interface concept
- RMI does not have a separate language (IDL) for defining remote interfaces
- Pre-defined interface `Remote`
- All RMI communication is based on interfaces that extends `java.rmi.Remote`
- Remote classes implement `java.rmi.Remote`
- Remote objects are instances of remote class

Frank Eliassen, SRL & Ifi/UiO

52

Java remote interface: Example

interface name *declares the Team interface as "remote"*

```
interface Team extends Remote {
public:
    String nama() throws RemoteException;
    Trainer[] trained_by() throws RemoteException;
    Club club() throws RemoteException;
    Player[] player() throws RemoteException;
    void chooseKeeper(Date d) throws RemoteException;
    void print() throws RemoteException;
};
```

remote operation

Frank Eliassen, SRL & Ifi/UiO

53

Java RMI parameter passing

- Atomic types transferred *by value*
- Remote objects transferred *by reference*
- None-remote objects transferred *by value*

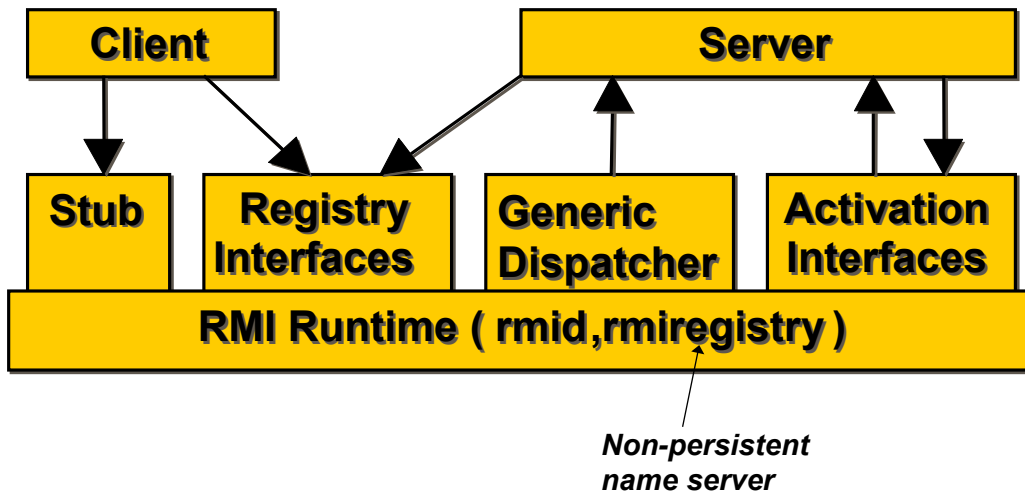
```
class Address {
    public:
        String street;
        String zipcode;
        String town;
};
interface Club extends Organisation, Remote {
    public:
        Address addr() throws RemoteException;
        ...
};
```

Returns a copy of the Address-object

Frank Eliassen, SRL & Ifi/UiO

54

Architecture of Java RMI



Frank Eliassen, SRL & Ifi/UiO

55

Summary

- Implementation of RMI
 - proxies, skeletons, dispatcher
 - interface processing, binding, location, activation
- Invocation semantics (under partial failure)
 - maybe, at-least-once, at-most-once
 - Reliability of RMI is at best "at-most-once"
- Multi-threaded servers
 - can in some cases be used to increase the throughput (method calls/time unit) if, e.g., I/O is the bottleneck
- Principles of CORBA
 - Clients may invoke methods of remote objects without worrying about: object location, programming language, operating system platform, communication protocols or hardware.
- Principles of Java RMI
 - Similar to CORBA but limited to a Java environment

Frank Eliassen, SRL & Ifi/UiO

56

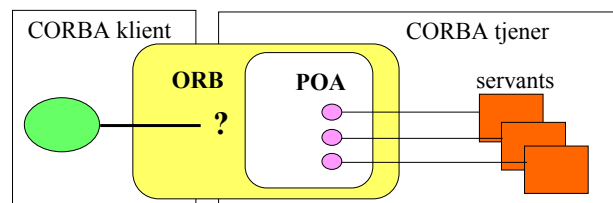
Extra slides

Frank Eliassen, SRL & Ifi/UiO

57

CORBA Portable object adapter (POA)

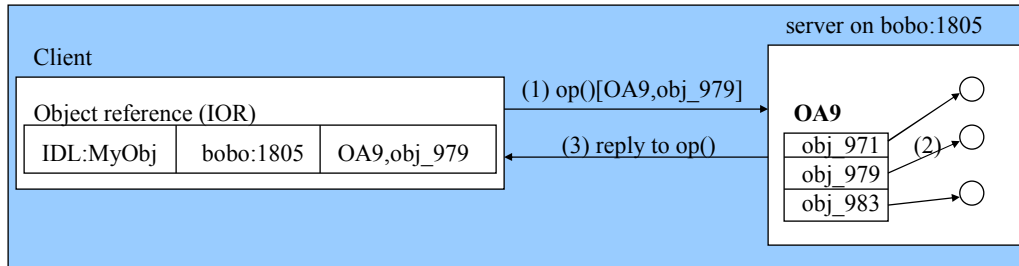
- Enables portability of object implementations across different ORBs
- Supports light weight transient object and persistent object identifiers (e.g., for objects stored in databases)
- Supports transparent object activations
- Extensible mechanism for activation policies
- Several POAs in one single server



Frank Eliassen, SRL & Ifi/UiO

58

Locating transient IORs over IIOP



Frank Eliassen, SRL & Ifi/UiO

59

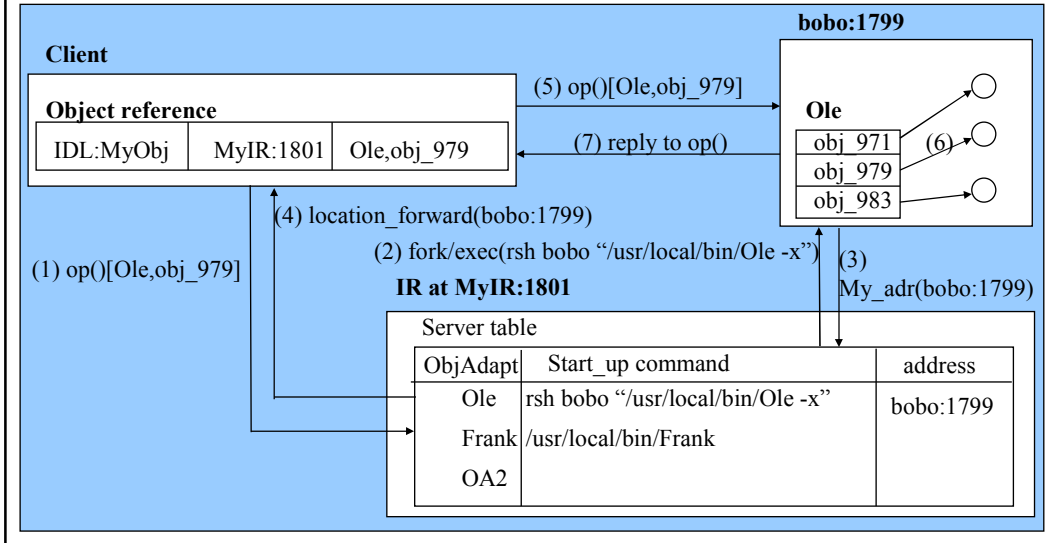
Locating persistent IOR over IIOP

- uses Implementation Repository (IR) as activator.
- IR:
 - handles process/thread-creation and -termination, a.m.
 - is not portable (specific to an ORB implementation)
 - not standardized
 - tailored to specific environments
 - not possible to write specifications that cover all environments
 - communication between an ORB and its IR is not visible to the client
 - Object migration, scaling, performance, and fault tolerance are dependent on IR
 - Implemented usually as a process at a fixed address
 - a set of host machines that is configured under the same IR is denoted a "location domain"

Frank Eliassen, SRL & Ifi/UiO

60

Locating persistent IOR over IIOP



Java RMI development process

