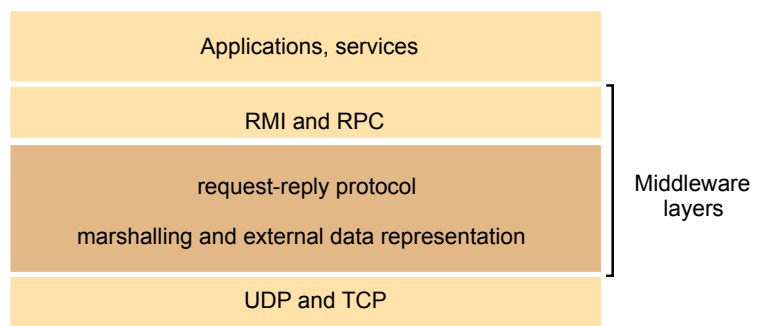


Remote Method Invocation - Design & Implementation

CS 475

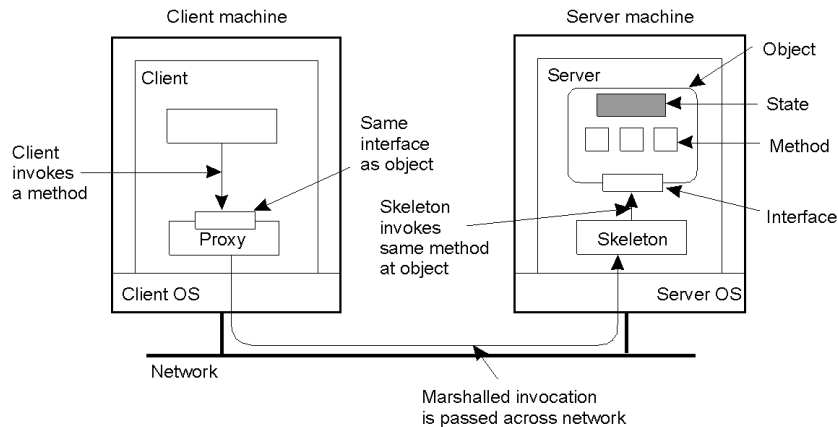
RMI 1

Middleware layers



RMI 2

Distributed Objects



RMI 3

Compile-time vs run-time objects

- Objects can be implemented in many different ways
 - compile-time objects, i.e. instances of classes written in object-oriented languages like Java, C++
 - database objects
 - procedural languages like C, with a appropriate "wrapper code" that gives it the appearance of an object
- Systems like Java RMI support compile-time objects
- Not possible or difficult in language-independent RMI middleware such as CORBA
 - these systems use object adapters
 - implementations of object interfaces are registered at an object adapter, which acts as an intermediary between the client and object implementation

RMI 4

Persistent vs transient objects

- ❑ Persistent objects continue to exist even if they are not contained in the address space of a server process
- ❑ The "state" of a persistent object has to be stored on a persistent store, i.e., secondary storage
- ❑ Invocation requests result in an instance of the object being created in the address space of a running process
 - many policies possible for object instantiation and (de) instantiation
- ❑ Transient objects only exist as long as their container server processes are running

RMI 5

Static vs dynamic remote method invocations

- ❑ Typical way for writing code that uses RMI is similar to the process for writing RPC
 - declare the interface in IDL, compile the IDL file to generate client and server stubs, link them with client and server side code to generate the client and the server executables
 - referred to as static invocation
 - requires the object interface to be known when the client is being developed
- ❑ Dynamic invocation
 - the method invocation is composed at run-time
invoke(object, method, input_parameters, output_parameters)
 - useful for applications where object interfaces are discovered at run-time, e.g. object browser, batch processing systems for object invocations, "agents"

RMI 6

Design Issues for RMI

- ❑ **RMI Invocation Semantics**
 - Invocation semantics depend upon implementation of **Request-Reply protocol** used by RMI
 - **Maybe, At-least-once, At-most-once**
- ❑ **Transparency**
 - Should remote invocations be transparent to the programmer?
 - **Partial failure, higher latency**
 - **Different semantics for remote objects, e.g. difficult to implement "cloning" in the same way for local and remote objects or to support synchronization operations, e.g. wait/notify**
 - **Current consensus: remote invocations should be made transparent in the sense that syntax of a remote invocation is the same as the syntax of local invocation (access transparency) but programmers should be able to distinguish between remote and local objects by looking at their interfaces, e.g. in Java RMI, remote objects implement the *Remote* interface**

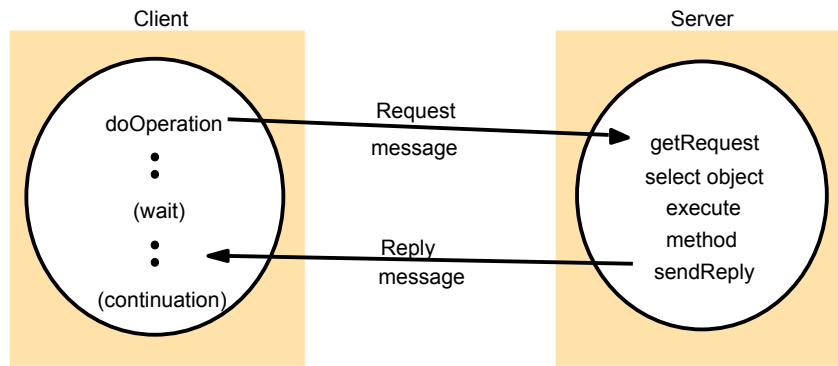
RMI 7

Issues in implementing RMI

- ❑ **Parameter Passing**
- ❑ **Request-Reply Protocol**
 - **Handling failures at client and/or server**
- ❑ **Supporting persistent objects, object adapters, dynamic invocations, etc.**

RMI 8

Request-reply communication



RMI 9

Operations of the request-reply protocol

public byte[] doOperation (RemoteObjectRef o, int methodId, byte[] arguments)
sends a request message to the remote object and returns the reply.

The arguments specify the remote object, the method to be invoked and the arguments of that method.

public byte[] getRequest ();
acquires a client request via the server port.

public void sendReply (byte[] reply, InetAddress clientHost, int clientPort);
sends the reply message reply to the client at its Internet address and port.

RMI 10

Request-reply message structure

| | |
|-----------------|----------------------------------|
| messageType | <i>int (0=Request, 1= Reply)</i> |
| requestId | <i>int</i> |
| objectReference | <i>RemoteObjectRef</i> |
| methodId | <i>int or Method</i> |
| arguments | <i>array of bytes</i> |

RMI 11

Representation of a remote object reference

| | | | | |
|------------------|----------------|----------------|----------------|----------------------------|
| <i>32 bits</i> | <i>32 bits</i> | <i>32 bits</i> | <i>32 bits</i> | |
| Internet address | port number | time | object number | interface of remote object |

RMI 12

CORBA interoperable object references

IOR format

| IDL interface type name | Protocol and address details | | Object key | | |
|---------------------------------|------------------------------|------------------|-------------|--------------|-------------|
| interface repository identifier | IIOP | host domain name | port number | adapter name | object name |

RMI 13

Request-Reply protocol

- ❑ Issues in **marshaling** of parameters and results
 - Input, output, Inout parameters
 - Data representation
 - Handling reference parameters
- ❑ Distributed object references
- ❑ Handling failures in request-reply protocol
 - Partial failure
 - Client, Server, Network

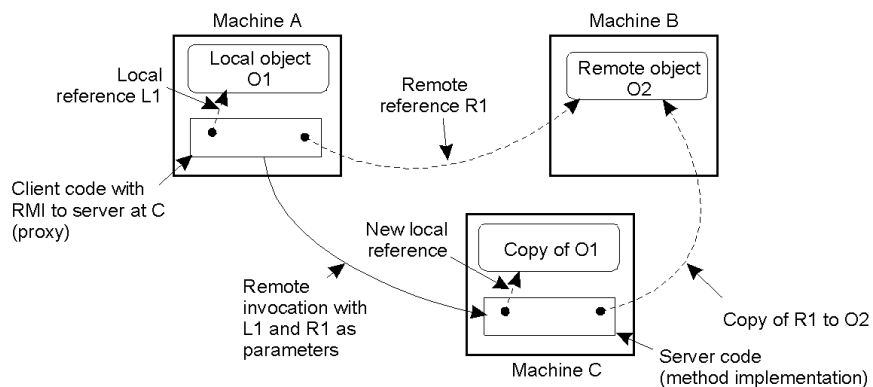
RMI 14

Marshalling

- ❑ Pack method arguments and results into a flat array of bytes
- ❑ Use a canonical representation of data types, e.g. integers, characters, doubles
- ❑ Examples
 - SUN XDR
 - CORBA CDR
 - Java serialization

RMI 15

Parameter Passing: local vs remote objects



Remote object references are passed by reference where local object references are passed by value

RMI 16

CORBA CDR for constructed types

| <i>Type</i> | <i>Representation</i> |
|-------------------|--|
| <i>sequence</i> | length (unsigned long) followed by elements in order |
| <i>string</i> | length (unsigned long) followed by characters in order (can also have wide characters) |
| <i>array</i> | array elements in order (no length specified because it is fixed) |
| <i>struct</i> | in the order of declaration of the components |
| <i>enumerated</i> | unsigned long (the values are specified by the order declared) |
| <i>union</i> | type tag followed by the selected member |

RMI 17

CORBA CDR message

| <i>index in sequence of bytes</i> | <i>4 bytes</i> | <i>notes on representation</i> |
|---------------------------------------|----------------|------------------------------------|
| 0-3 | 5 | <i>length of string</i> |
| 4-7 | "Smit" | 'Smith' |
| 8-11 | "h__" | |
| 12-15 | 6 | <i>length of string</i> |
| 16-19 | "Lond" | 'London' |
| 20-23 | "on__" | |
| 24-27 | 1934 | <i>unsigned long</i> |

The flattened form represents a *Person* struct with value: {'Smith', 'London', 1934}

RMI 18

Indication of Java serialized form

| <i>Serialized values</i> | | | | <i>Explanation</i> |
|--------------------------|-----------------------|------------------------|-------------------------|--|
| Person | 8-byte version number | | h0 | <i>class name, version number</i> |
| 3 | int year | java.lang.String name: | java.lang.String place: | <i>number, type and name of instance variables</i> |
| 1934 | 5 Smith | 6 London | h1 | <i>values of instance variables</i> |

The true serialized form contains additional type markers; h0 and h1 are handles

RMI 19

RPC exchange protocols

| <i>Name</i> | <i>Messages sent by</i> | | |
|-------------|-------------------------|---------------|--------------------------|
| | <i>Client</i> | <i>Server</i> | <i>Client</i> |
| R | <i>Request</i> | | |
| RR | <i>Request</i> | <i>Reply</i> | |
| RRA | <i>Request</i> | <i>Reply</i> | <i>Acknowledge reply</i> |

RMI 20

Handling failures

- Types of failure
 - Client unable to locate server
 - Request message lost
 - Reply message lost
 - Server crashes after receiving a request
 - Client crashes after sending a request

RMI 21

Handling failures

- Client cannot locate server
 - Reasons
 - Server has crashed
 - Server has moved
 - (RPC systems) client compiled using old version of service interface
 - System must report error (remote exception) to client
 - Loss of transparency

RMI 22

Handling failures

- ❑ Lost request message
 - Retransmit a fixed number of times before throwing an exception
- ❑ Lost reply message
 - Client resubmits request
 - Server choices
 - Re-execute procedure → service should be idempotent so that it can be repeated safely
 - Filter duplicates → server should hold on to results until acknowledged

RMI 23

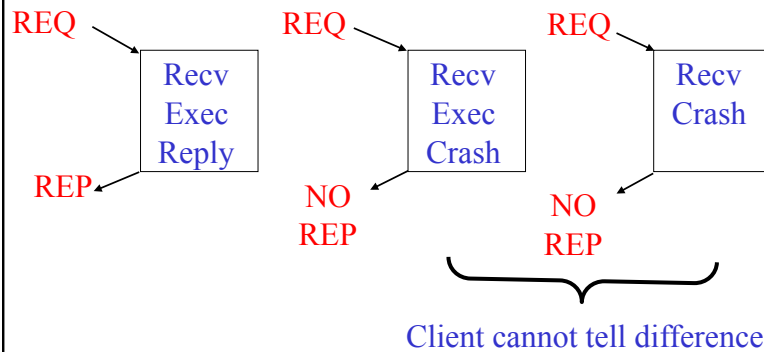
Invocation semantics

| <i>Fault tolerance measures</i> | | | <i>Invocation semantics</i> |
|-----------------------------------|----------------------------|---|-----------------------------|
| <i>Retransmit request message</i> | <i>Duplicate filtering</i> | <i>Re-execute procedure or retransmit reply</i> | |
| No | Not applicable | Not applicable | <i>Maybe</i> |
| Yes | No | Re-execute procedure | <i>At-least-once</i> |
| Yes | Yes | Retransmit reply | <i>At-most-once</i> |

RMI 24

Handling failures

❑ Server crashes



RMI 25

Handling failures

❑ Server crashes

- At least once (keep trying till server comes up again)
- At most once (return immediately)
- *Exactly once impossible to achieve*

❑ SUN RPC

- At least once semantics on successful call and maybe semantics if unsuccessful call

❑ CORBA, Java RMI

- at most once semantics

RMI 26

Handling failures

- Implementing the request-reply protocol on top of TCP
 - Does it provide applications with different invocation semantics?
 - NO!
 - TCP does not help with server crashes
 - If a connection is broken, the end pts do not have any guarantees about the delivery of messages that may have been in transit

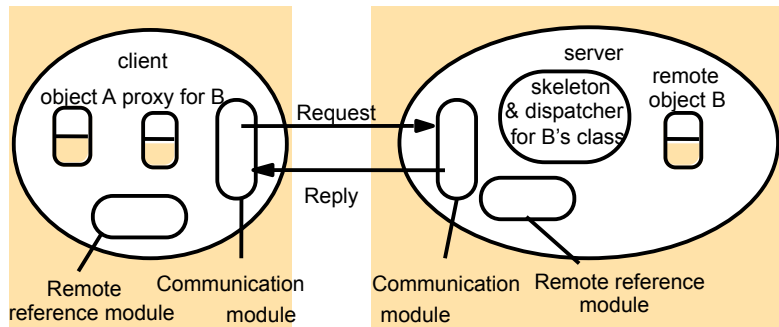
RMI 27

Handling failures

- Client crashes
 - If client crashes before RPC returns, we have an "orphan" computation at server
 - Wastes resources, could also start other computations
 - Orphan detection
 - Reincarnation (client broadcasts new "epoch" when it comes up again)
 - Expiration (RPC has fixed amount of time T to do work)

RMI 28

RMI Software Components



RMI 29

RMI Software Components

- ❑ **Communication module**
 - Implements the request-reply protocol
- ❑ **Remote reference module**
 - Responsible for translating between local and remote object references and for creating remote object references
 - Maintains remote object table that maintains a mapping between local & remote object references
 - E.g. Object Adaptor in CORBA

RMI 30

RMI - Object Activation

- ❑ Activation of remote objects
 - Some applications require that information survive for long periods of times
 - However, objects not in use all the time, so keeping them in running processes is a potential waste of resources
 - Object can be activated on demand
 - E.g. standard TCP services such as FTP on UNIX machines are activated by inetd

RMI 31

Object Activation

- ❑ Active and passive objects
 - Active object = instantiated in a running process
 - Passive object = not currently active but can be made active
 - Implementation of its methods, and marshalled state stored on disk
- ❑ Activator responsible for
 - Registering passive objects that are available for activation
 - Starting named server processes and activating remote objects in them
 - Keeping track of locations of servers for remote objects that it has already activated
- ❑ Examples: CORBA implementation repository, JAVA RMI has one activator on each server computer

RMI 32

RMI - Other topics

- ❑ Persistent object stores
 - An object that is guaranteed to live between activations of processes is called a persistent object
 - Stores the state of an object in a marshalled (serialized) form on disk
- ❑ Location service
 - Objects can migrate from one system to another during their lifetime
 - Maintains mapping between object references and the location of an object

RMI 33

RMI - Other topics

- ❑ Distributed Garbage Collection
 - Needed for reclaiming space on servers
- ❑ Passing "behavior"
 - Java allows objects (data + code) to be passed by value
 - If the class for an object passed by value is not present in a JVM, its code is downloaded automatically
 - See Java RMI tutorial example
- ❑ Use of **Reflection** in Java RMI
 - Allows construction of generic dispatcher and skeleton

RMI 34

Distributed Garbage Collection

- ❑ Java approach based on reference counting
 - Each server process maintains a list of remote processes that hold remote object references for its remote objects
 - When a client first acquires a remote reference to an object, it makes an `addRef()` invocation to server before creating a proxy
 - When a clients local garbage collector notices that a proxy is no longer reachable, it makes a `removeRef()` invocation to the server before deleting the proxy
 - When the local garbage collector on the server notices that the list of client processes that have a remote reference to an object is empty, it will delete the object (unless there are any local objects that have a reference to the object)
- ❑ Other approaches
 - "Evictor" pattern
 - Leases