


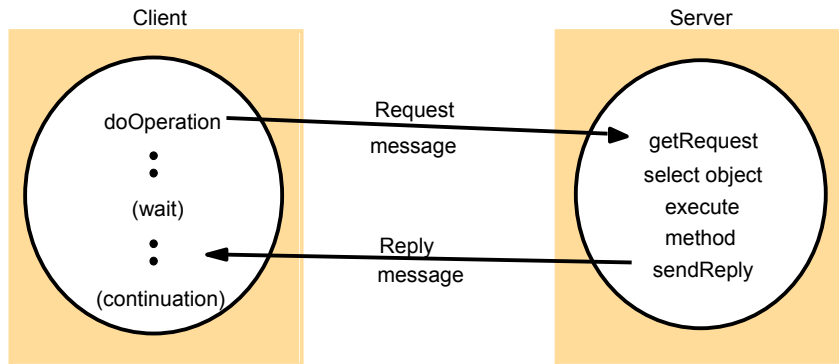
RMI: Design & Implementation

Operating Systems

Middleware layers



Request-reply communication



RMI 3

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
 - 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 4

Request-Reply protocol

- ❑ Issues in **marshaling** of parameters and results
 - Input, output, Inout parameters
 - Data representation
 - Passing pointers? (e.g., call by reference in C)
- ❑ Distributed object references
- ❑ Handling failures in request-reply protocol
 - Partial failure
 - Client, Server, Network

RMI 5

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 6

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 7

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 8

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 9

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	"Smith"	'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 10

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

Representation of a remote object reference

<i>32 bits</i>	<i>32 bits</i>	<i>32 bits</i>	<i>32 bits</i>	
IP address	port number	time	object number	interface of remote object

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

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 14

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

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

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 17


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 18

Handling failures

❑ Server crashes



RMI 19

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 20

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 21

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)

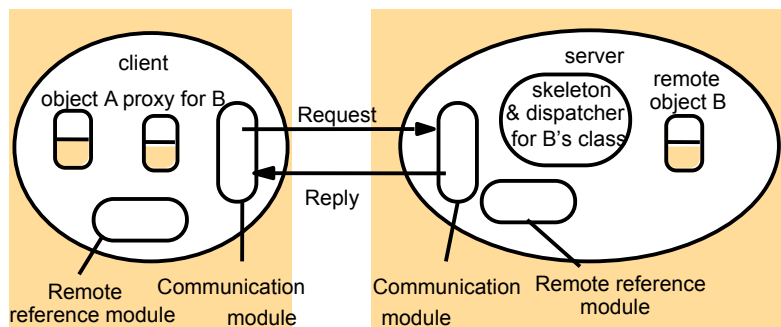
RMI 22

RMI Implementation

- 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 23

The role of proxy and skeleton in remote method invocation



RMI 24

RMI Implementation

- RMI software
 - Generated by IDL compiler
 - Proxy
 - Behaves like remote object to clients (invoker)
 - Marshals arguments, forwards message to remote object, unmarshals results, returns results to client
 - Skeleton
 - Server side stub;
 - Unmarshals arguments, invokes method, marshals results and sends to sending proxy's method
 - Dispatcher
 - Receives the request message from communication module, passes on the message to the appropriate method in the skeleton
- Server and Client programs

RMI 25

RMI Implementation

- Binder
 - Client programs need a means of obtaining a remote object reference
 - Binder is a service that maintains a mapping from textual names to remote object references
 - Servers need to register the services they are exporting with the binder
 - Java RMIregistry, CORBA Naming service
- Server threads
 - Several choices: thread per object, thread per invocation
 - Remote method invocations must allow for concurrent execution

RMI 26

Optional Assignment : Implementing RPC/RMI

□ RPC

- Implement communication module using UDP as transport protocol
 - Main challenge: support at most once semantics
- Hand-construct client & server stubs, dispatcher for arithmetic server program
 - Marshalling arguments & results

□ RMI

- Implement communication module using TCP
- Implement generic dispatcher making use of Java reflection
- Implement remote reference module
- Hand-construct client stub & skeleton
 - Marshalling arguments & results (can use Java serialization)
 - Can construct generic skeleton using Java reflection

RMI 27

RMI - Other topics

□ 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 28

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 29

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 30

Classes supporting Java RMI



RMI 31

Readings

- Coulouris - Chapters 4,5
- WWW (see links on class web page)
 - Java RMI tutorial on web
 - "A Young Persons Guide to SOAP"

RMI 32