

Slides for Chapter 5: Remove Invocation



From **Coulouris, Dollimore, Kindberg and Blair**
**Distributed Systems:
Concepts and Design**

Edition 5, © Addison-Wesley 2012

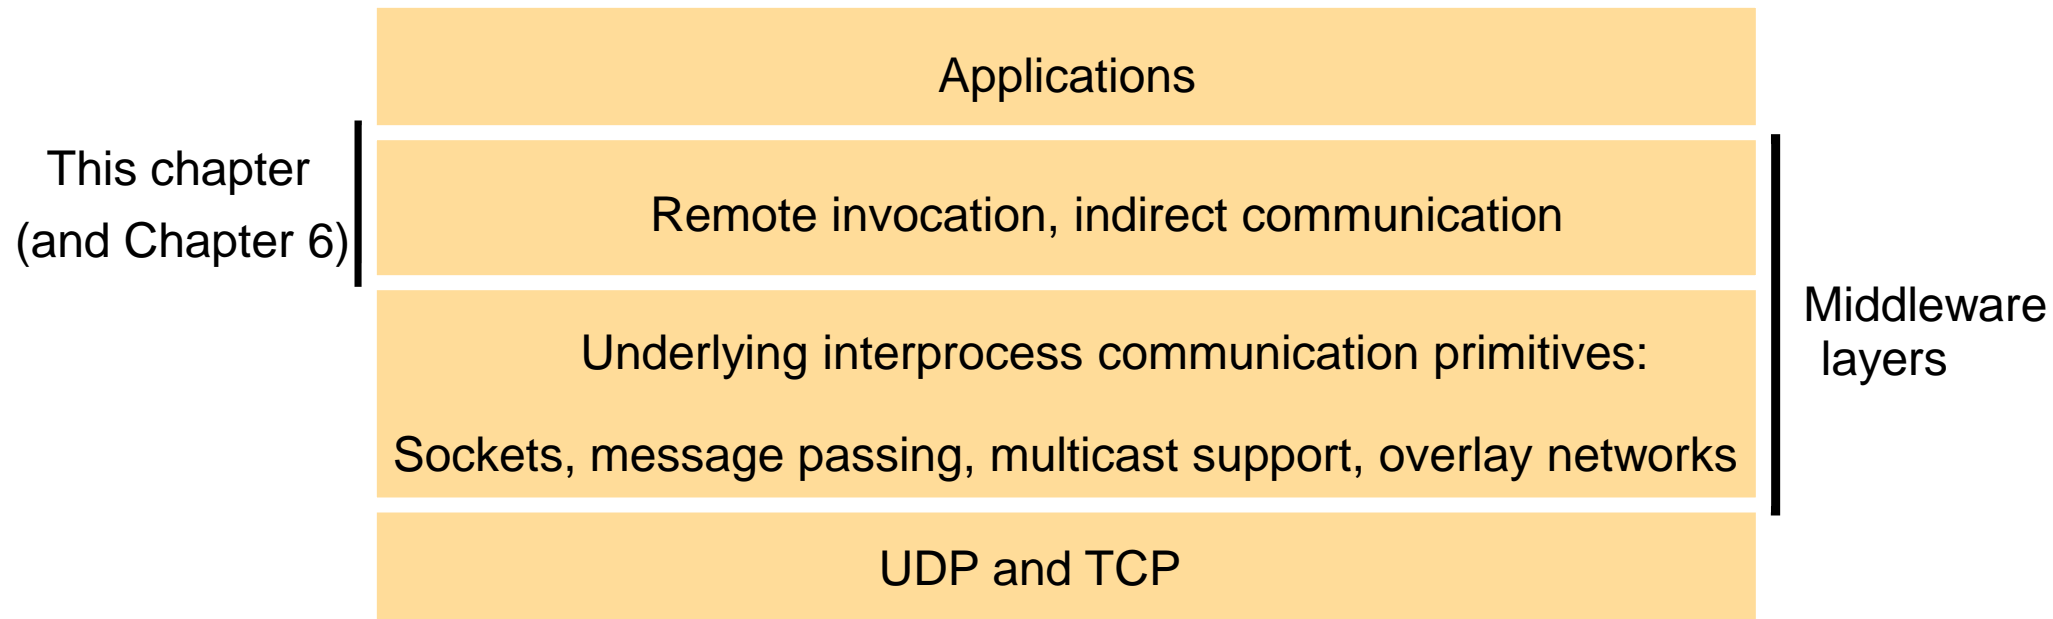
(Not included here: 5.4, 5.5)

Introduction [5.1]

- This chapter: how processes/objects/components/services communication via remote invocation (Chap 2)
- Request-reply
 - Small/thin pattern on top of message passing
 - Can use directly in app (“app protocols”), or build RPC/RMI on
- Remote Procedure Call (RPC)
 - Make a remote procedure look (almost) like a local one to call
- Remote Method Invocation (RMI)
 - Make a remote object look (almost) like a local one to invoke
 - Note: ‘RMI’ is generic category, Java RMI is a specific instance

Figure 5.1

Middleware layers



Request-reply protocols [5.2]

- Support low-level client-server interactions
 - Usually synchronous and reliable
- Built on top of `send` and `receive` operations from Chapter 4
 - Usually use UDP datagrams, could use TCP streams
- Three primitives
 - `doOperation`: client sends request message to server
 - `getRequest`: server receives request msg, selects+invokes oper.
 - `sendReply`: server sends reply message back to (blocked) client

Figure 5.2

Request-reply communication

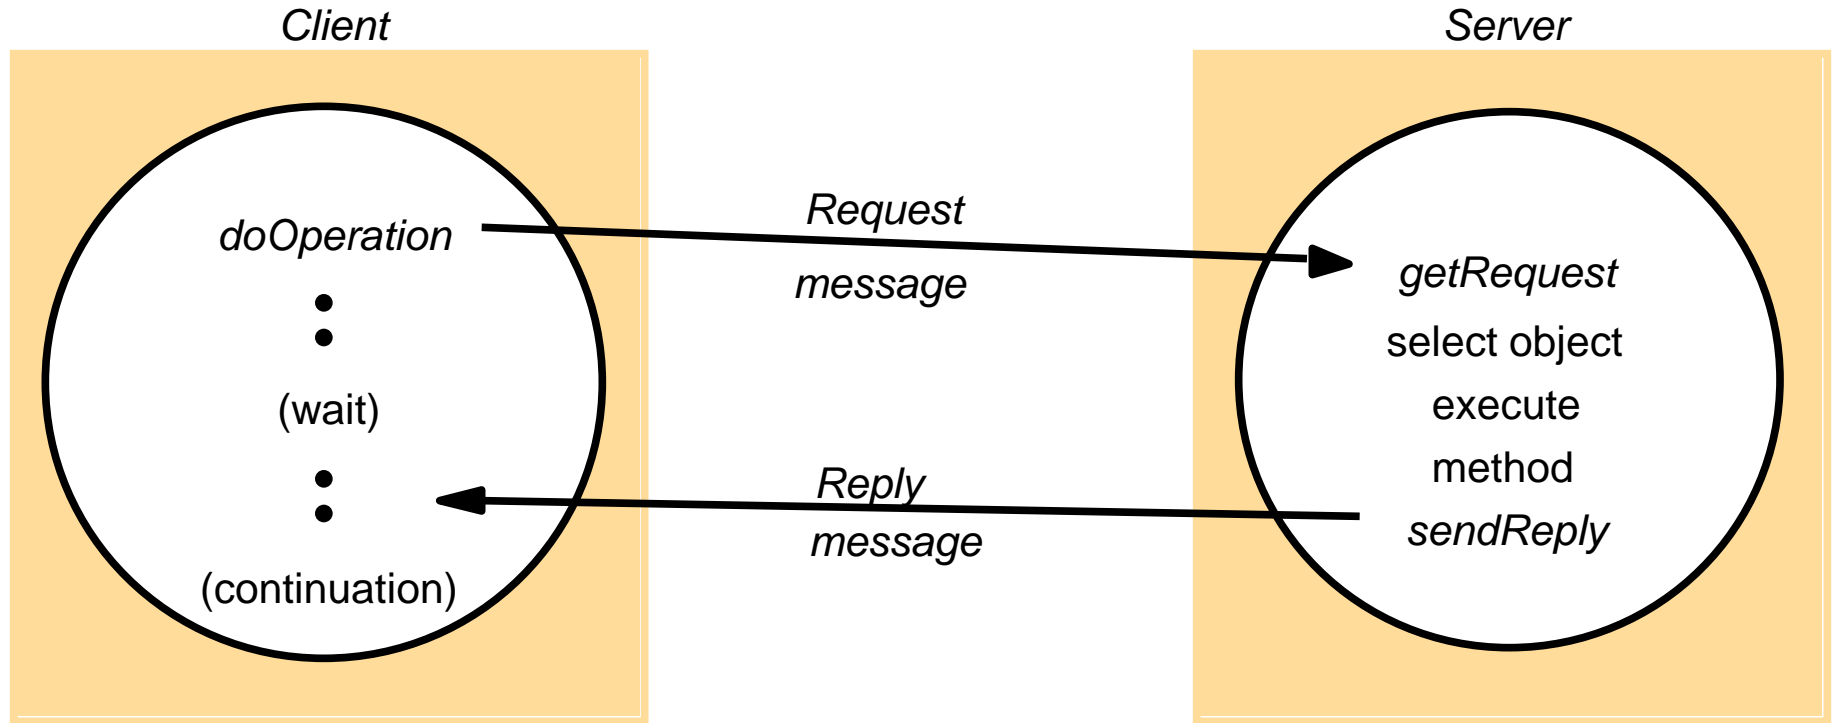


Figure 5.3

Operations of the request-reply protocol

public byte[] doOperation (RemoteRef s, int operationId, byte[] arguments)

sends a request message to the remote server and returns the reply.

The arguments specify the remote server, the operation to be invoked and the arguments of that operation.

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.

Figure 5.4

Request-reply message structure

messageType	<i>int (0=Request, 1= Reply)</i>
requestId	<i>int</i>
remoteReference	<i>RemoteRef</i>
operationId	<i>int or Operation</i>
arguments	<i>array of bytes</i>

Request-reply protocols (cont.)

- Message identifiers: must identify request uniquely
 - `requestId`: usually a sequence counter (makes unique at client)
 - Client/sender identifier endpoint (with `requestId`, globally unique)
- Failure model
 - Over UDP: omission, misordering
 - Over UDP or TCP: server crash failure (later, Byzantine...)
- Timeouts: `doOperation` uses when blocked for reply
 - Options to use?
- Duplicate request msgs: server may get >1 times
 - how? problem?
 - Soln: server tracks what got from client (how?)

Request-reply protocols (cont.)

- Lost reply messages
 - Idempotent operation: just redo
 - Else store reply history (how many? How to use?)
- Q: should client and/or server ACK messages?

Figure 5.5

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>

Using TPC streams to implement request-reply protocol

- Advantages

- Never need multi-packet protocols
- “Reliable”

- Disadvantages

- More CPU intensive: scale

HTTP RR protocol

- HTTP protocol specifies
 - Messages in RR exchange
 - Methods
 - Arguments
 - Results
 - Marshalling rules
 - Content negotiation
 - Authentication
- Implemented over TCP streams
 - Early versions: new connection for each request (later persistent)
- Request & reply msgs marshalled into ASCII
- Resource data can be represented as a byte sequence

HTTP methods


- GET (URL, ...) : send data from file or program output
 - Can add args, e.g. formdata
 - Can be conditional on modification date of resource
 - Can grab only parts of the data
 - Only for idempotent requests
- HEAD: like GET, but no return data, just meta-data
- POST (URL, data, ...) : Run resource (program usually) that can handle the data
 - Program may change data on the server
 - Uses?

HTTP methods (cont.)

- `PUT (URL, data, ...)`: store data at URL; idempotent
- `DELETE (URL)`: delete resource at the URL; idempotent
- `OPTIONS`: tells client allowable methods & requirements
- `TRACE`: like ping for IP

Figure 5.6

HTTP request message




<i>method</i>	<i>URL or pathname</i>	<i>HTTP version</i>	<i>headers</i>	<i>message body</i>
GET	//www.dcs.qmw.ac.uk/index.html	HTTP/ 1.1		

- Message body optional
- Headers can contain
 - Request modifiers
 - Client info (last modify, acceptable content types, ...)
 - Authentication info

Figure 5.7

HTTP *Reply* message



<i>HTTP version</i>	<i>status code</i>	<i>reason</i>	<i>headers</i>	<i>message body</i>
HTTP/1.1	200	OK		resource data

- Status and reason: did the server succeed
- Header can pass
 - Extra info about server or access to resource (e.g., challenge)

Remote procedure call [5.3]

- Design issues
 - Style of programming promoted by RPC: using interfaces
 - Call semantics
 - Transparency

Programming with interfaces

- Explicit interface
 - Hide a lot of implementation details
 - Tell exactly how a client can access the server
- Keeping implementation separate from interface
 - Good idea? Why?
- Differences from local procedure interface
 - Can't access shared memory variables between client and server
 - Call by reference does not make sense for RPC
 - Parameters are `in`, `out`, or `inout`
 - Can't pass pointers
- IDL originally developed for RPC

Figure 5.8

CORBA IDL example

```
// In file Person.idl
struct Person {
    string name;
    string place;
    long year;
} ;

interface PersonList {
    readonly attribute string listname;
    void addPerson(in Person p) ;
    void getPerson(in string name, out Person p) ;
    long number() ;
};
```

RPC call semantics

- Choices for implementing `doOperation`
 - Retry request message
 - Duplicate request filtering at server
 - Retransmission of results: keep reply history, or re-execute procedure

Figure 5.9

Call semantics

<i>Fault tolerance measures</i>			<i>Call 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>

- How can each of these happen?
- What would you call local procedure call semantics?

Transparency

- RPC tries to offer at least location and access transparency
- Does client need to know call semantics?
- Implement RPC with stub over an RR protocol (Fig 5.10)

- Note: not covering Sun RPC (5.3.3), not testable

Figure 5.10

Role of client and server stub procedures in RPC

