

Distributed Systems

03. Remote Procedure Calls

Paul Krzyzanowski
Rutgers University
Fall 2016

September 19, 2016

© 2014-2016 Paul Krzyzanowski

1

Socket-based communication

- Socket API: all we get from the OS to access the network
- Socket = distinct end-to-end communication channels
- **Read/write model**
 - Send a bunch of bytes
 - Read a bunch of bytes
 - Send a bunch of bytes
 - Read a bunch of bytes
 - ...
- Application implements its protocol
- Line-oriented, text-based protocols common
 - Not efficient but easy to debug & use

September 19, 2016

© 2014-2016 Paul Krzyzanowski

2

Sample SMTP Interaction

```
$ telnet cs.rutgers.edu 25
Trying 128.8.4.2...
Connected to cs.rutgers.edu.
Escape character is '^]'.
220 aramis.rutgers.edu ESMTP Sendmail 8.11.7p3+Sun/8.8.8; Mon, 19 Sep 2016 12:12:01 -0400 (EDT)
HELO pk.org
250 aramis.rutgers.edu Hello aramis.rutgers.edu [128.6.4.2], pleased to meet you
MAIL FROM:<pxk@cs.rutgers.edu>
250 2.1.0 <pxk@cs.rutgers.edu>... Sender ok
RCPT TO:<p@pk.org>
250 2.1.5 <p@pk.org>... Recipient ok
DATA
354 Enter mail, end with "." on a line by itself
From: Paul Krzyzanowski <pxk@cs.rutgers.edu>
Subject: test message
Date: Mon, 17 Feb 2020 17:00:16 -0500
To: Whoever <testuser@pk.org>

Hi,
This is a test
.
250 2.0.0 s8FGEZS13883 Message accepted for delivery
quit
221 2.0.0 aramis.rutgers.edu closing connection
```

SMTP = Simple Mail Transfer Protocol

This is the message body.
Headers may define the structure of the
message but are ignored for delivery.

September 19, 2016

© 2014-2016 Paul Krzyzanowski

3

Problems with the sockets API

The `sockets` interface forces a read/write mechanism

Programming is often easier with a functional interface

To make distributed computing look more like centralized computing, I/O (read/write) is not the way to go

September 19, 2016

© 2014-2016 Paul Krzyzanowski

4

RPC

1984: Birrell & Nelson

– Mechanism to call procedures on other machines

Remote Procedure Call

September 19, 2016

© 2014-2016 Paul Krzyzanowski

5

Regular procedure calls

You write:

```
x = f(a, "test", 5);
```

The compiler parses this and generates code to:

- Push the value 5 on the stack
- Push the address of the string "test" on the stack
- Push the current value of a on the stack
- Generate a call to the function f

In compiling `f`, the compiler generates code to:

- Push registers that will be clobbered on the stack to save the values
- Adjust the stack to make room for local and temporary variables
- Before a return, unadjust the stack, put the return data in a register, and issue a return instruction

September 19, 2016

© 2014-2016 Paul Krzyzanowski

6

Implementing RPC

No architectural support for remote procedure calls

Simulate it with tools we have (local procedure calls)

Simulation makes RPC a **language-level construct** instead of an **operating system construct**

The compiler creates code to send messages to invoke remote functions

The OS gives us sockets

September 19, 2016 © 2014-2016 Paul Krzyzanowski 7

Implementing RPC

The trick:

Create **stub functions** to make it appear to the user that the call is local

On the client
The stub function has the function's interface Packages parameters and calls the server

On the server
The stub function (skeleton) receives the request and calls the local function

September 19, 2016 © 2014-2016 Paul Krzyzanowski 8

Stub functions

1. Client calls stub (params on stack)

September 19, 2016 © 2014-2016 Paul Krzyzanowski 9

Stub functions

2. Stub **marshals** params to net message

Marshalling = put parameters in a form suitable for transmission over a network (serialized)

September 19, 2016 © 2014-2016 Paul Krzyzanowski 10

Stub functions

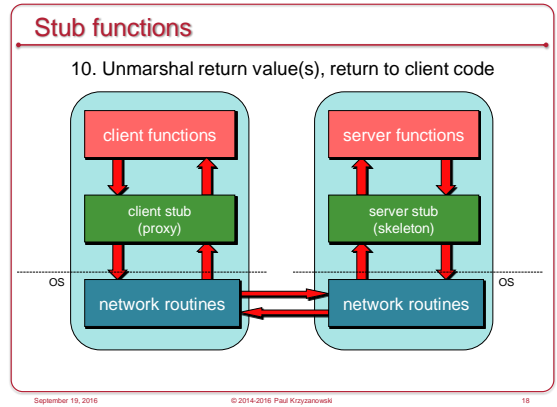
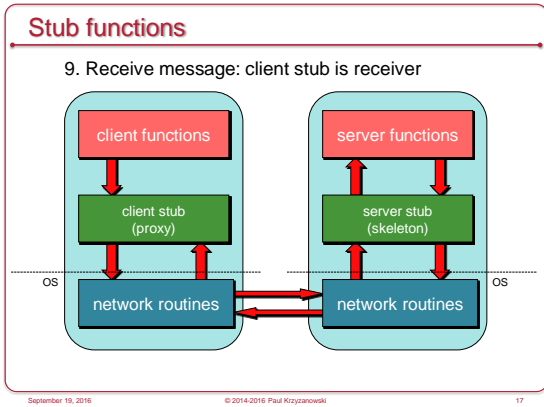
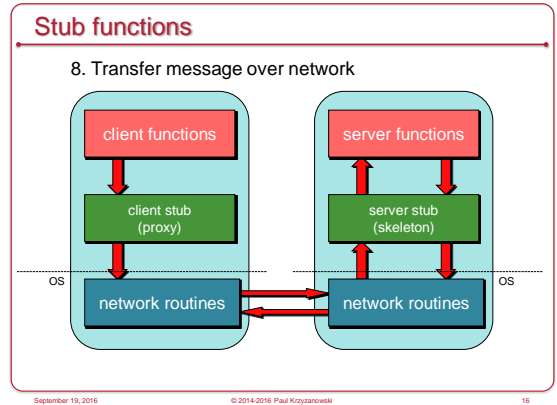
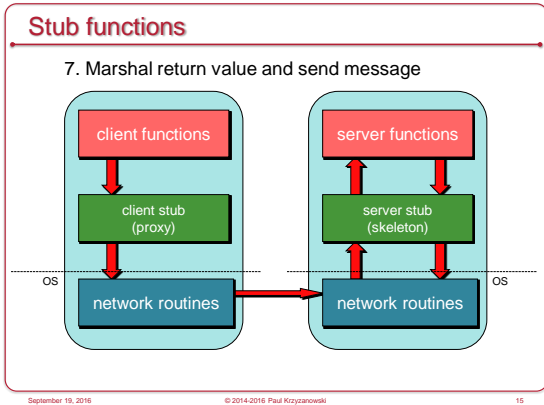
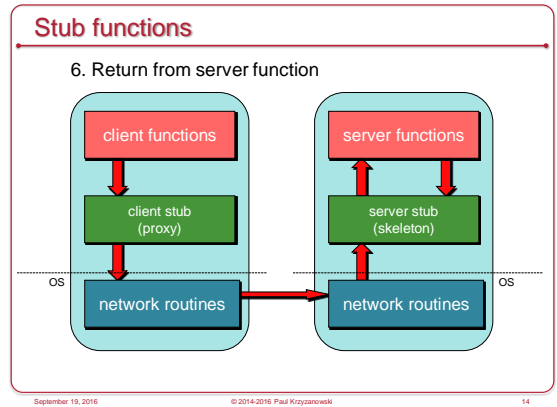
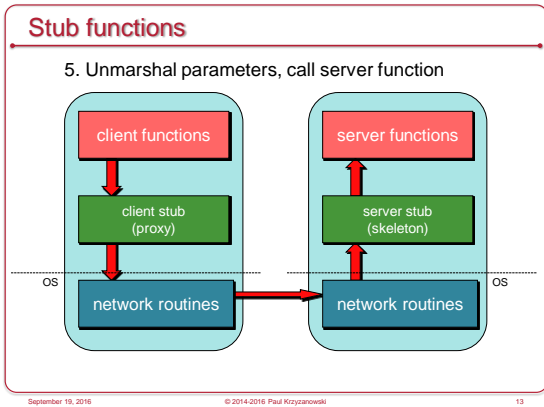
3. Network message sent to server

September 19, 2016 © 2014-2016 Paul Krzyzanowski 11

Stub functions

4. Receive message: send it to server stub

September 19, 2016 © 2014-2016 Paul Krzyzanowski 12



A client proxy looks like the remote function

- Client stub has the same interface as the remote function
- Looks & feels like the remote function to the programmer
 - But its function is to
 - Marshal parameters
 - Send the message
 - Wait for a response from the server
 - Unmarshal the response & return the appropriate data
 - Generate exceptions if problems arise

September 19, 2016

© 2014-2016 Paul Krzyzanowski

19

A server stub contains two parts

- **Dispatcher – the listener**
 - Receives client requests
 - Identifies appropriate function (method)
- **Skeleton – the unmarshaller & caller**
 - Unmarshals parameters
 - Calls the local server procedure
 - Marshals the response & sends it back to the dispatcher
- All this is invisible to the programmer
 - The programmer doesn't deal with any of this
 - Dispatcher + Skeleton may be integrated
 - Depends on implementation

September 19, 2016

© 2014-2016 Paul Krzyzanowski

20

RPC Benefits

- RPC gives us a procedure call interface
- Writing applications is simplified
 - RPC hides all network code into stub functions
 - Application programmers don't have to worry about details
 - Sockets, port numbers, byte ordering
- Where is RPC in the OSI model?
 - Layer 5: Session layer: Connection management
 - Layer 6: Presentation: Marshaling/data representation
 - Uses the transport layer (4) for communication (TCP/UDP)

September 19, 2016

© 2014-2016 Paul Krzyzanowski

21

RPC has challenges

September 19, 2016

© 2014-2016 Paul Krzyzanowski

22

Parameter passing

Pass by value

- Easy: just copy data to network message

Pass by reference

- Makes no sense without shared memory

September 19, 2016

© 2014-2016 Paul Krzyzanowski

23

Pass by reference?

1. Copy items referenced to message buffer
2. Ship them over
3. Unmarshal data at server
4. Pass *local* pointer to server stub function
5. Send new values back

To support complex structures

- Copy structure into pointerless representation
- Transmit
- Reconstruct structure with local pointers on server

September 19, 2016

© 2014-2016 Paul Krzyzanowski

24

Where to bind?

Need to locate host and correct server process

September 19, 2016

© 2014-2016 Paul Krzyzanowski

25

Where to bind? – Solution 1

Maintain a centralized DB that can locate a host that provides a particular service

(Birrell & Nelson's 1984 proposal)

Challenges:

- Who administers this?
- What is the scope of administration?
- What if the same services run on different machines (e.g., file systems)?

September 19, 2016

© 2014-2016 Paul Krzyzanowski

26

Where to bind? – Solution 2

A server on each host maintains a DB of *locally* provided services

September 19, 2016

© 2014-2016 Paul Krzyzanowski

27

Transport protocol

TCP or UDP? Which one should we use?

- Some implementations may offer only one (e.g. TCP)
- Most support several
 - Allow programmer (or end user) to choose at runtime

September 19, 2016

© 2014-2016 Paul Krzyzanowski

28

When things go wrong

- Local procedure calls do not fail
 - If they core dump, entire process dies
- More opportunities for error with RPC
- Transparency breaks here
 - Applications should be prepared to deal with RPC failure

September 19, 2016

© 2014-2016 Paul Krzyzanowski

29

When things go wrong

- Semantics of remote procedure calls
 - Local procedure call: *exactly once*
- A remote procedure call may be called:
 - 0 times:
 - server crashed or server process died before executing server code
 - 1 time:
 - everything worked well, as expected
 - 1 or more times: excess latency or lost reply from server and client retransmission

September 19, 2016

© 2014-2016 Paul Krzyzanowski

30

RPC semantics

- Most RPC systems will offer either:
 - **at least once** semantics
 - or **at most once** semantics
- Understand application:
 - **idempotent** functions: may be run any number of times without harm
 - **non-idempotent** functions: those with side-effects
- Try to design your application to be idempotent
 - Not always easy!
 - Store transaction IDs, previous return data, etc.

More issues

Performance

- RPC is slower ... a lot slower (why?)

Security

- messages may be visible over network – do we need to hide them?
- Authenticate client?
- Authenticate server?

Programming with RPC

Language support

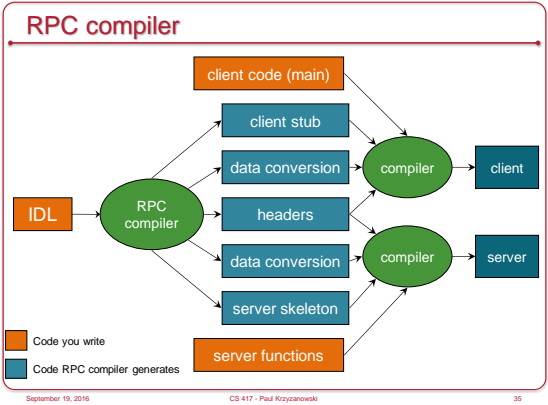
- Many programming languages have no language-level concept of remote procedure calls (C, C++, Java <J2SE 5.0, ...)
- These compilers will not automatically generate client and server stubs
- Some languages have support that enables RPC (Java, Python, Haskell, Go, Erlang)
- But we may need to deal with heterogeneous environments (e.g., Java communicating via XML)

Common solution

- Interface Definition Language (IDL): describes remote procedures
- Separate compiler that generate stubs (pre-compiler)

Interface Definition Language (IDL)

- Allow programmer to specify remote procedure interfaces (names, parameters, return values)
- Pre-compiler can use this to generate client and server stubs
 - Marshaling code
 - Unmarshaling code
 - Network transport routines
 - Conform to defined interface
- An IDL looks similar to function prototypes



Writing the program

- Client code has to be modified
 - Initialize RPC-related options
 - Identify transport type
 - Locate server/service
 - Handle failure of remote procedure calls
- Server functions
 - Generally need little or no modification

RPC API

What kind of services does an RPC system need?

- **Name service operations**
 - Export/lookup of binding information (ports, machines)
 - Support dynamic ports
- **Binding operations**
 - Establish client/server communications using appropriate protocol (establish endpoints)
- **Endpoint operations**
 - Listen for requests, export endpoint to name server (often the main program on the server)

September 19, 2016

© 2014-2016 Paul Krzyzanowski

37

RPC API

What kind of services does an RPC system need?

- **Security operations**
 - Authenticate client/server
- **Internationalization operations (maybe)**
- **Marshaling/data conversion operations**
- **Stub memory management**
 - Dealing with "reference" data, temporary buffers
- **Program ID operations**
 - Allow applications to access IDs of RPC interfaces
 - Can you pass references to remote functions to other processes?

September 19, 2016

© 2014-2016 Paul Krzyzanowski

38

Sending data over the network

September 19, 2016

© 2014-2016 Paul Krzyzanowski

39

We need a stream of bytes

```
struct item {
    char name[64];
    unsigned long id;
    int number_in_stock;
    float rating;
    double price;
} scratcher = {
    "Bear Claw Black Telescopic Back Scratcher",
    00120,
    332,
    4.6,
    5.99
}
42 65 61 72 20 43 6c 61 77 20 42 6c 61 63 6b 20 54 ...
```

September 19, 2016

© 2014-2016 Paul Krzyzanowski

40

Representing data

No such thing as
incompatibility problems on local system

Remote machine may have:

- Different byte ordering
- Different sizes of integers and other types
- Different floating point representations
- Different character sets
- Alignment requirements

September 19, 2016

© 2014-2016 Paul Krzyzanowski

41

Representing data

IP (headers) forced all to use **big endian** byte ordering for 16- and 32-bit values

Big endian: Most significant byte in low memory
– SPARC < V9, Motorola 680x0, older PowerPC ← IP headers use big endian

Little endian: Most significant byte in high memory
– Intel/AMD IA-32, x64

Bi-endian: Processor may operate in either mode
– ARM, PowerPC, MIPS, SPARC V9, IA-64 (Intel Itanium)

```
main() {
    unsigned int n;
    char *a = (char *) &n;

    n = 0x11223344;
    printf("%02x, %02x, %02x, %02x\n",
           a[0], a[1], a[2], a[3]);
}
```

Output on an Intel:
44, 33, 22, 11

Output on a PowerPC:
11, 22, 33, 44

September 19, 2016

© 2014-2016 Paul Krzyzanowski

42

Representing data: serialization

Need standard encoding to enable communication between heterogeneous systems

- **Serialization**

- Convert data into a pointerless format: *an array of bytes*

- **Examples**

- XDR (eXternal Data Representation), used by ONC RPC
 - JSON (JavaScript Object Notation)
 - W3C XML Schema Language
 - ASN.1 (ISO Abstract Syntax Notation)
 - Google Protocol Buffers

September 19, 2016

© 2014-2016 Paul Krzyzanowski

43

Serializing data

- **Implicit typing**

- only values are transmitted, not data types or parameter info
 - e.g., ONC XDR (RFC 4506)

- **Explicit typing**

- Type is transmitted with each value
 - e.g., ISO's ASN.1, XML, protocol buffers, JSON

Marshaling vs. serialization – almost synonymous:

Serialization: converting an object into a sequence of bytes that can be sent over a network

Marshaling: bundling parameters into a form that can be reconstructed (unmarshaled) by another process. May include object ID or other state. Marshaling uses serialization.

September 19, 2016

© 2014-2016 Paul Krzyzanowski

44

XML: eXtensible Markup Language

```
<ShoppingCart>
  <Items>
    <Item>
      <ItemID> 00120 </ItemID>
      <Item> Bear Claw Black Telescopic Back Scratcher </Item>
      <Price> 5.99 </Price>
    </Item>
    <Item>
      <ItemID> 00121 </ItemID>
      <Item> Scalp Massager </Item>
      <Price> 5.95 </Price>
    </Item>
  </Items>
</ShoppingCart>
```

Benefits:

- Human-readable
- Human-editable
- Interleaves structure with text (data)

September 19, 2016

© 2014-2016 Paul Krzyzanowski

45

JSON: JavaScript Object Notation

- **Lightweight** (relatively efficient) data interchange format
 - Introduced as the *"fat-free alternative to XML"*
 - Based on JavaScript
- **Human writeable and readable**
- **Self-describing** (explicitly typed)
- **Language independent**
- **Easy to parse**
- **Currently converters for 50+ languages**
- **Includes support for RPC invocation via JSON-RPC**

September 19, 2016

© 2014-2016 Paul Krzyzanowski

46

JSON Example

```
{ "menu": {
  "id": "file",
  "value": "File",
  "popup": {
    "menuitem": [
      { "value": "New", "onclick": "CreateNewDoc()" },
      { "value": "Open", "onclick": "OpenDoc()" },
      { "value": "Close", "onclick": "CloseDoc()" }
    ]
  }
}}
```

from json.org/example.html

September 19, 2016

© 2014-2016 Paul Krzyzanowski

47

Google Protocol Buffers

- **Efficient mechanism for serializing structured data**
 - Much simpler, smaller, and faster than XML
- **Language independent**
- **Define messages**
 - Each message is a set of names and types
- **Compile the messages to generate data access classes for your language**
- **Used extensively within Google. Currently over 48,000 different message types defined.**
 - Used both for RPC and for persistent storage

September 19, 2016

© 2014-2016 Paul Krzyzanowski

48

Example (from the Developer Guide)

```
message Person {
  required string name = 1;
  required int32 id = 2;
  optional string email = 3;

  enum PhoneType {
    MOBILE = 0;
    HOME = 1;
    WORK = 2;
  }

  message PhoneNumber {
    required string number = 1;
    optional PhoneType type = 2 [default = HOME];
  }

  repeated PhoneNumber phone = 4;
}
```

September 19, 2016

© 2014-2016 Paul Krzyzanowski

49

Example (from the Developer Guide)

```
Person person;
person.set_name("John Doe");
person.set_id(1234);
person.set_email("jdoe@example.com");
fstream output("myfile", ios::out | ios::binary);
person.SerializeToOstream(&output);
```

September 19, 2016

© 2014-2016 Paul Krzyzanowski

50

Efficiency example (from the Developer Guide)

```
<person>
  <name>John Doe</name>
  <email>jdoe@example.com</email>
</person>
```

XML version

```
person {
  name: "John Doe"
  email: "jdoe@example.com"
}
```

Text (uncompiled) protocol buffer

- Binary encoded message: ~28 bytes long, 100-200 ns to parse
- XML version: ≥69 bytes, 5,000-10,000 ns to parse

September 19, 2016

© 2014-2016 Paul Krzyzanowski

51

The End

September 19, 2016

© 2014-2016 Paul Krzyzanowski

52