

CSE 486/586 Distributed Systems Remote Procedure Call

Steve Ko
Computer Sciences and Engineering
University at Buffalo

CSE 486/586


Recap

- Paxos phase 1
 - A proposer sends a prepare message.
 - Acceptors reply with the highest-numbered proposal
- Paxos phase 2:
 - The proposer waits for a majority of acceptors.
 - The proposer chooses the value from the highest-numbered proposal.
 - Upon receiving a new proposal, acceptors either:
 - » Accept it
 - » Or, reject it if there was another prepare request with N' higher than N, and it replied to it (due to the promise in phase 1).

CSE 486/586

2

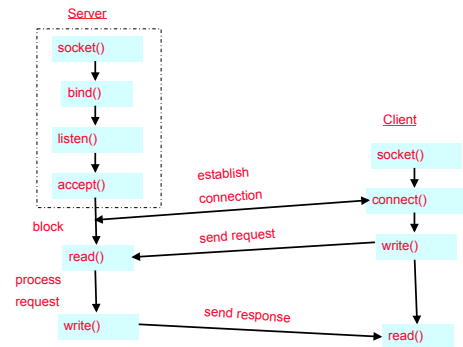
Recall?



CSE 486/586

3

Socket API



CSE 486/586

4

What's Wrong with Socket API?

- Low-level read/write
- Communication oriented
- Same sequence of calls, repeated many times
- Etc, etc...
- Not programmer friendly

CSE 486/586

5

Another Abstraction

- RPC (Remote Procedure Call)
 - Goal: it should appear that the programmer is calling a local function
 - Mechanism to enable function calls between different processes
 - First proposed in the 80's
- Examples
 - Sun RPC
 - Java RMI
 - CORBA
- Other examples that borrow the idea
 - XML-RPC
 - Android Bound Services with AIDL
 - Google Protocol Buffers

CSE 486/586

6

How Do You Generate Stubs?

- Ever heard of C/C++, Java, Python syntax for RPC?
 - None!
- Language compilers don't generate client and server stubs.
- **Common solution:** use a separate language and a pre-compiler

CSE 486/586

13

Interface Definition Language (IDL)

- Allow programmers to express remote procedures, e.g., names, parameters, and return values.
- Pre-compilers take this and generate stubs, marshalling/unmarshalling mechanisms.
- Similar to writing function definitions

CSE 486/586

14

Example: SUN XDR

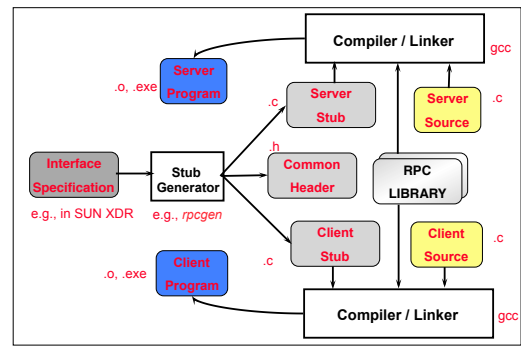
```

const MAX = 1000;
typedef int FileIdentifier;
typedef int FilePointer;
typedef int Length;
struct Data {
    int length;
    char buffer[MAX];
};
struct writeargs {
    FileIdentifier f;
    FilePointer position;
    Data data;
};
struct readargs {
    FileIdentifier f;
    FilePointer position;
    Length length;
};
program FILEREADWRITE {
    version VERSION {
        void WRITE(writeargs)=1;
        Data READ(readargs)=2;
    }= 9999;
}
    
```

CSE 486/586

15

Stub Generation



CSE 486/586

16

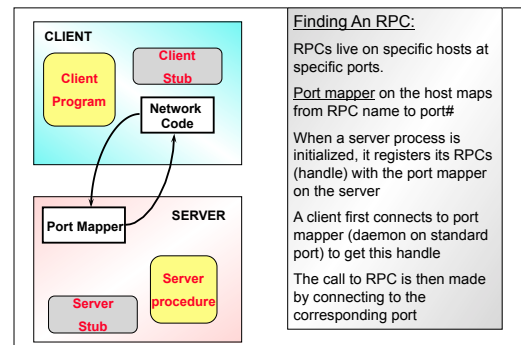
How Do You Find the Server Process?

- Solution 1
 - Central DB (the first solution proposed)
- Solution 2
 - Local DB with a well-known port (SUN RPC)

CSE 486/586

17

Local DB with Well-Known Port



CSE 486/586

18

How to Pass Parameters?

- Pass by value: no problem
 - Just copy the value
- What about pointers/references?
 - Need to copy the actual data as well
 - Marshall them at the client and unmarshall them at the server
 - Pass the local pointers/references
- What about complex data structures? struct, class, etc.
 - Need to have a platform independent way of representing data

CSE 486/586

19

External Data Representation

- Communication between two heterogeneous machines
 - Different byte ordering (big-endian & little-endian)
 - Different sizes of integers and other types
 - Different floating point representations
 - Different character sets
 - Alignment requirements
- Used in general contexts, not just in RPCs

CSE 486/586

20

Example: Google Protocol Buffers

- Goal: language- and platform-neutral way to specify and serialize data
- Provides syntax & pre-compiler (open-source)
 - Pre-compiler generates code to manipulate objects for a specific language, e.g. C++, Java, Python.
 - The runtime support applies a fast & sloppy compression algorithm.

```
message Book {
  required string title = 1;
  repeated string author = 2;
  optional BookStats statistics = 3;
  message BookStats {
    required int32 sales = 1;
  }
}
```

CSE 486/586

21

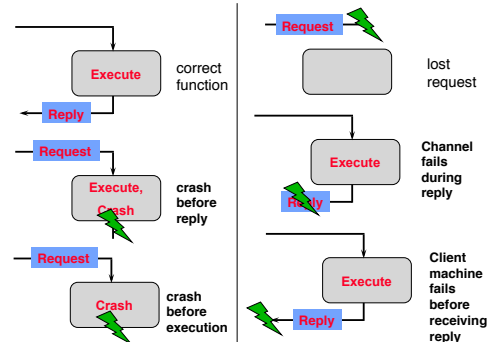
What About Failures?

- Local calls do not fail.
- Remote calls might fail.
- Programmers should deal with this.
 - No transparency here

CSE 486/586

22

Failure Modes of RPC



CSE 486/586

23

Invocation Semantics

- Local procedure call: **exactly-once**
- Remote procedure call:
 - 0 times: server crashed or server process died before executing server code
 - 1 time: everything worked well, as expected
 - 1 or more: excess latency or lost reply from server and client retransmission
- When do these make sense?
 - Idempotent functions: OK to run any number of times
 - Non-idempotent functions: cannot do it
- What we can offer
 - At least once
 - At most once

CSE 486/586

24

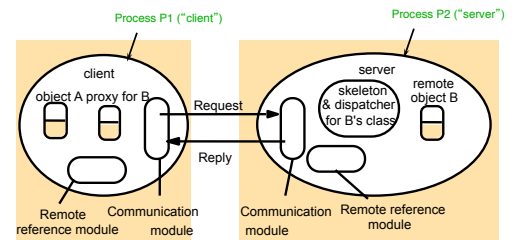
Invocation Semantics

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

CSE 486/586

25

Remote Method Invocation (RMI)



CSE 486/586

26

Summary

- RPC enables programmers to call functions in remote processes.
- IDL (Interface Definition Language) allows programmers to define remote procedure calls.
- Stubs are used to make it appear that the call is local.
- Semantics
 - Cannot provide exactly once
 - At least once
 - At most once
 - Depends on the application requirements

CSE 486/586

27

Acknowledgements

- These slides contain material developed and copyrighted by Indranil Gupta (UIUC).

CSE 486/586

28