

#### Communication in Distributed Systems: RPC

#### Corso di Sistemi Distribuiti e Cloud Computing A.A. 2024/25

Valeria Cardellini

Laurea Magistrale in Ingegneria Informatica

#### Communication in distributed systems

- Based on message passing
  - Send and receive messages
- To allow for message passing, parties must agree on many low-level communication details
  - How many volts to signal a 0 bit and how many for a 1 bit?
  - How many bits for an integer?
  - How does the receiver know which is the last bit of the message?
  - How can the receiver find if a message has been corrupted and what to do then?
- ... we don't care about them

#### Basic networking model and its adaptation

- We already know the solution: divide network communication in layers
  - The well known ISO/OSI reference model
  - We don't care about low-level details: for distributed systems, the lowest-level interface is that of the network layer
- Adapted layering scheme for distributed systems



Valeria Cardellini - SDCC 2024/25

### Middleware layer

- Middleware provides common services and protocols that can be used by many different distributed applications and systems
  - General-purpose
  - Application-independent
- Some examples of middleware services and protocols:
  - Communication: remote procedures/methods, queue messages, multicasting
  - Naming: to allow easy sharing of resources
  - Security: to allow applications to communicate securely
  - Distributed consensus, including distributed commit
  - Distributed locking: to protect a shared resource against simultaneous access by a collection of distributed processes (e.g., multiple clients update a file in shared storage)
  - Data consistency

#### Types of communication

- Let's distinguish
  - Persistency
    - Transient versus persistent communication

#### - Synchronization

• Synchronous vs. asynchronous communication

#### - Time dependence

• Discrete vs. streaming communication

Valeria Cardellini - SDCC 2024/25

#### Persistent vs. transient communication

- Persistent communication
  - Message is stored by middleware as long as it takes to deliver it to receiver
  - Sender does not need to continue execution after submitting message
  - Receiver does not need to be executing when message is submitted

#### Transient communication

- Message is stored by middleware only as long as sender and receiver are executing: sender and receiver have to be active at time of communication
- If delivery is not possible, message is discarded
- Transport-layer example: routers store and forward, but discard if forward is not possible

#### Synchronous communication...

- Once message has been submitted, sender is blocked until operation is completed
- Send and receive are *blocking* operations
- Places for synchronization:
  - 1. At request submission
  - 2. At request delivery
  - 3. After request processing



Valeria Cardellini - SDCC 2024/25

#### ... vs asynchronous communication

- Once message has been submitted, sender continues its processing: message is temporarily stored by middleware until it is transmitted
- Send is *non-blocking*, receive can be blocking or nonblocking

#### Discrete vs. streaming communication

- Discrete communication
  - Each message forms a complete unit of information
- Streaming communication
  - Involves sending multiple messages, in temporal relationship or related to each other by sending order, which is needed to reconstruct complete information
  - E.g., audio, video

## Combining communication types

- Combining persistence and synchronization
- a) Persistent asynchronous communication
  - E.g., email, Teams chat, message-oriented middleware
- b) Persistent synchronous communication
  - Sender is blocked until message is delivered to receiver



(a) Persistent asynchronous communication

(b) Persistent synchronous communication

#### Combining communication types

- Combining persistence and synchronization
- c) Transient asynchronous communication
  - Sender does not wait but message can be lost if receiver is unreachable (e.g., UDP)



(c)Asynchronous communication

Valeria Cardellini - SDCC 2024/25

10

### Combining communication types

- Alternatives for transient synchronous communication
  - d) Receipt-based synchronous: sender is blocked until message is in receiver space (e.g., asynchronous RPC)
  - e) Delivery-based synchronous: sender is blocked until message is delivered to receiver





(d) Receipt-based synchronous communication

(e) Delivery-based synchronous communication

#### Combining communication types

- Alternatives for transient synchronous communication
  - f) Response-based synchronous: sender is blocked until it receives a reply message from receiver (e.g., synchronous RPC)



Valeria Cardellini - SDCC 2024/25

12

#### What happens when a client calls a RPC



• Which failures can happen?

#### Failures during communication

- Different communication failures between sender (client) and receiver (server): what can go wrong?
- 1. Request or reply message is lost or delayed, connection is reset
  - Pitfall: Network is reliable
- 2. Server crashes
  - a) before performing service
  - b) after performing service
  - Client cannot distinguish between a) and b) b)
- 3. Client crashes after sending request



Valeria Cardellini - SDCC 2024/25

14

#### Failure semantics during communication

- In a DS, which is the semantics of communication in the presence of failures?
  - May-be semantics
  - At-least-once semantics
  - At-most-once semantics
  - Exactly-once semantics
- Failure semantics applies both to service processing (e.g., RPC) and message delivery (e.g., MOM)
  - Let's focus on service processing for now
    - Sender -> client, receiver -> server

More guarantees, more complexity

#### Failure semantics in a nutshell

- Maybe: no guarantee that service has been executed or not by server
- At-least-once: service, if executed, has been executed at least once (perhaps more than once)
  - If client does not receive the reply: service *might* have been executed (or it might not, or it might have been executed more than once)
- At-most-once: service, if executed, has been executed at most once
  - If client receives the reply, it has been processed by server only once
- Exactly-once: service has been executed precisely once

Valeria Cardellini - SDCC 2024/25

16

#### Basic mechanisms for implementation

- 3 basic mechanisms to implement failure semantics
- 1. Client side: Request Retry (RR1)
  - Client keeps trying until it gets a reply or is confident about server failure after a certain number of failed retries
- 2. Server side: Duplicate Filtering (DF)
  - Server discards any duplicate request from same client
- 3. Server side: Result Retransmit (RR2)
  - Server keeps result (reply) so that it can be retransmitted without being computed again when server receives duplicate requests

- No guarantee that service has been executed or not by server
- No action is taken to ensure reliable communication: no mechanism (among RR1, DF, RR2) is used
- E.g., best-effort in UDP



#### **At-least-once semantics**

- Service, if executed, has been executed at least once
  - Perhaps more than once, because of request duplication due to retransmissions
- Client uses RR1, server uses neither DF nor RR2
- Upon reply receipt, client does not know how many times its request has been processed by server: client does not know about server status
  - Server may have executed the service but crashed before sending the reply: when timeout expires, client resends the request, server processes it again and sends the reply to client
- Suitable for idempotent services
  - Idempotent service: when applied multiple times to same input, service produces same output as if it were applied only once

• E.g., get(key) or put(key, value) in KV store Valeria Cardellini - SDCC 2024/25

#### At-least-once semantics



Valeria Cardellini - SDCC 2024/25

20

#### At-most-once semantics

- Service, if executed, has been executed at most once
  - Client knows that, *if it* receives the reply, it has been processed by server *only once*
  - In case of failure, no information (at-most-once: response has been calculated at most once, but possibly also none)
- All basic mechanisms (RR1, DF, RR2) are used
  - Client retransmits request when timeout expires
  - Server maintains some state to identify duplicate requests and avoid processing the same request more than once
- Suitable also for non-idempotent services
- No constraint on consequent actions
  - No strict coordination between client and server: in case of failure, client does not know if server run the service, while server ignores if client knows that the service run

#### At-most-once semantics



Valeria Cardellini - SDCC 2024/25

22

#### At-most-once semantics: implementation

- Server detects duplicate requests and returns saved reply instead of re-running service handler()
- How to detect duplicate request?
  - Client includes a unique ID (xid) with each request and uses same xid when retransmitting request
     Server:
- Issues to address
  - How to ensure unique xid?
  - Server must eventually discard info about saved responses: when is discard safe?
    - Use sliding windows and sequence numbers
    - Discard information older than maximum lifetime
  - How to handle duplicate requests while original one is still executing? Add a pending flag for each executing RPC

Idempotent Receiver pattern https://martinfowler.com/articles/patterns-of-

#### distributed-systems/idempotent-receiver.html

```
Server:
if seen[xid]
r = old[xid]
else
r = handler()
old[xid] = r
seen[xid] = true
```

- Strongest and hardest guarantee to implement, especially in large-scale DS
- Requires full agreement on interaction between client and server
  - Service is run only once or not run at all: all-or-nothing semantics
    - If everything goes well: service runs only once, duplicates are found
    - If something goes wrong: client or server knows if service has run (once - all) or not (never - nothing)
- Semantics with concordant knowledge of each other's state and without constraint on maximum duration of interaction protocol between client and server
  - No constraint on maximum duration: barely practical in real systems!

#### Exactly-once semantics: more mechanisms

- Server-side basic mechanisms (RR1, DF, RR2) are not enough
- Need more mechanisms to tolerate serverside faults
  - Transparent server replication
  - Write-ahead logging
  - Recovery
    - Mechanisms to recover from whatever state failed server left behind and begin processing from a safe point
    - *Snapshot*: captures a consistent state of server (more complex if server is distributed)
    - *State checkpointing*: saves a snapshot of server state on persistent storage

### Write-ahead logging (WAL) pattern

- aka Commit log
- Goal
  - Provide durability guarantee by persisting every state change as command to append-only log
- How
  - Each state change is stored as log entry in file on disk and log is appended sequentially
  - File can be read on every restart and state can be recovered by replaying all log entries



Valeria Cardellini - SDCC 2024/25

26

### Summing up failure semantics

- At-least once and at-most once semantics are feasible and widely used in DS
- We often choose the lesser of two evils, i.e., at-leastonce semantics
  - Also easier to scale

Distributed systems are all about trade-offs!

#### **Distributed application programming**

- You know explicit network programming
  - Operating system constructs based on socket API and explicit management of message exchange
  - Used in most network applications (e.g. web browser, web server)
  - X Distribution is not transparent and requires developer effort
- How to increase abstraction level of distributed programming? By means of communication middleware between OS and applications
  - Hide complexity of underlying hw and sw layers
  - Free developer from automatable tasks
  - Improve software quality by reusing known, correct, and efficient solutions

Valeria Cardellini - SDCC 2024/25

28

#### **Distributed application programming**

- Implicit network programming
  - Language-level constructs

#### - Remote Procedure Call (RPC)

 Distributed app is realized through procedure calls, but caller (client) and callee (server) are located on remote machines and communication details are hidden to developer

#### - Remote method invocation (Java RMI)

• Distributed app in Java is realized by invoking object methods running on remote machines



#### Remote Procedure Call (RPC)

- Idea (by Birrel and Nelson, 1984): use client/server model to call procedures executed on other machines
  - Process on machine A calls procedure on machine B
  - Calling process on A is suspended
  - Called procedure is execute on B
  - Input and output parameters are transported into messages
  - No message passing is visible to developer



Valeria Cardellini - SDCC 2024/25

 Used in many distributed systems, including cloud computing ones

Why RPC

 Developed and employed in many languages and frameworks, among which:



#### Local procedure call

- Example of local procedure call: newlist = append(data, dbList)
- Caller pushes to stack input parameters (data, dbList) and return address
- When callee returns, control is back to caller



How to make RPC look as much as possible as a local procedure call?

Valeria Cardellini - SDCC 2024/25

32

#### **RPC:** architecture

- Solution: create proxies (aka stubs) to make it appear that the call is local
- Client side: client stub exposes service's interface
   Client calls client stub that manages all the details
- Server side: server stub receives request and dispatches it calling the local procedure
- · Goal: distribution transparency
  - Stubs are automatically generated
  - Developer focuses on application logic



#### **RPC:** basic steps

- 1. On client side, client calls a local procedure, called client stub
- Client stub packs request message and call local OS
  - Client stub marshals parameters: converts parameters from local to common format and packages them into a message
- Client OS sends request message to remote OS
- 4. Remote OS delivers request message to server stub
- 5. On server side, server stub unpacks request message and calls server as it was a local procedure
  - Server stub unmarshals parameters: extracts parameters from message and converts them from common to local format
- 6. Server executes local call and returns result to server stub
- 7. Server stubs packs reply message (marshals return value(s)) and calls OS
- 8. Server OS sends reply message to client OS
- 9. Client OS delivers message to client stub
- 10. Client stub unpacks reply message (unmarshals return value(s)) and returns result to client 34

Valeria Cardellini - SDCC 2024/25

#### **RPC:** example

RPC doit(a,b)



3. Message is sent across the network

#### RPC: what is needed

- Exchange messages, so to make it appear to developer that procedure call is local; we need to:
  - Identify request and reply messages, remote procedure
  - Pass parameters
- Manage data heterogeneity
  - Which data? Parameters, return value(s)
  - Marshaling vs. serialization:
    - *Marshaling*: bundle parameters into a form that can be reconstructed (unmarshaled) by another process
    - **Serialization**: convert object into a sequence of bytes that can be sent over a network; serialization is used in marshaling
- Handle failures due to distribution
  - During communication
  - User errors

Valeria Cardellini - SDCC 2024/25

36

#### **RPC: challenges**

- Implementation challenges for RPC
- 1. Manage heterogeneity in data representation
  - In addition, client and server must agree on transport protocol for message passing: TPC, UDP, both?
- 2. Perform parameter passing by reference
  - Client and server run on different machines with their own address space
- 3. Define failure semantics
  - Local procedure call: exactly-once
  - Remote procedure call: at-least-once or at-most-once (in most cases)
- 4. Bind client to server, i.e., locate server endpoint

- Client and server may use different data representations
  - E.g., byte ordering (little endian vs big endian), data size, padding, …
  - RPC needs to define the details of how RPC messages are sent on the wire
- Alternatives (general, not only RPC) to handle heterogeneity in data representation:
  - 1. Specify encoding within message itself
  - 2. Let sender convert data into receiver encoding
  - 3. Convert data into common encoding agreed between parties
    - Sender: converts from local (i.e., native) to common
    - Receiver: converts from common to local
  - 4. Let an intermediary convert between different encodings

Data heterogeneity

- Let's compare alternatives #2 and #3, assuming *N* distributed components
- #2: each component knows all conversion functions
   ✓ Faster conversion
   X Higher number of conversion functions: N\*(N-1)
- #3: all components agree on common encoding for data representation and each component knows how to convert from local to common format and vice versa
  - X Slower conversion
  - $\checkmark$  Lower number of conversion functions: 2\*(*N*-1)
- #3: standard choice in RPC systems

#### Data heterogeneity: patterns

- Patterns to implement alternatives #3 and #4
- Proxy
  - Goal: support access (and location) transparency
  - Manage access to an object using another proxy object
    - Proxy is created in local address space to represent remote object and exposes same interface of remote object
- Broker
  - Goal: separate and encapsulate communication details from its functionality
  - Enable components to interact without handling remote concerns by themselves
  - Locate server for client, hide communication details, etc.
- Proxy (aka stub): standard choice in RPC systems
  - Who automatically generates stubs?

Valeria Cardellini - SDCC 2024/25

40

#### Parameter passing techniques

- · Call by value
  - Parameter value is copied in a local isolated storage (usually stack)
  - Callee acts on copied data and changes will not affect caller
- Call by reference
  - Reference (pointer) to parameter is copied into stack
  - Callee acts directly on caller data
- Call by copy-restore
  - A somehow special case of call by reference: data is copied into caller stack; when procedure returns, updated content is copied back (restored)
  - Available in few programming languages (e.g., Ada, Fortran)

#### **RPC** parameter passing

- A reference is a memory address
  - Valid only in its context (local machine)
  - We need a pointer-less representation
- Solution: simulate call by reference by using call by copy-restore
  - Client stub copies the pointed data in the request message and sends the message to server stub
  - Server stub acts on copy, using the address space of the receiver host
  - If the copy is modified, it will be then restored by client stub overwriting the original data
  - Size of data to be copied should be known
  - What happens if data contains a pointer?

Valeria Cardellini - SDCC 2024/25

#### Semantics of remote call/method

- Exactly once semantics is costly: most RPC systems implement weaker semantics
- At-least-once semantics: if client receives reply from server, then remote call has been executed at least once by server
- At-most-once semantics: if client receives reply from server, then remote call/method has been executed at most once by server

- Binding: how to locate the server endpoint, including the proper process (port or transport address) on it
  - In principle: can be static or dynamic
- Static binding
  - Binding is known at design time: server address and other info (e.g., port) are hard-coded
  - Easy and no overhead, but lacks transparency and flexibility
- Dynamic binding
  - At run-time
  - Increased overhead, but gains transparency and flexibility
    - E.g., we can redirect requests in case of server replication
  - Try to limit overhead

44

#### Server binding: dynamic

- Two phases in client/server relationship
- Naming: static phase before execution
  - Client specifies to whom it wants to be connected, using a unique name that identifies the service
  - Unique names are associated with operations or abstract interfaces and binding is made to the specific service interface
- Addressing: dynamic phase during execution
  - Server effectively binds to client when client invokes service
  - Depending on middleware implementation, multiple replica servers can be looked for

- Addressing can be explicit or implicit
  - Explicit addressing: client sends request using broadcast or multicast, waiting only for first reply
  - Implicit addressing: there is a name server (aka binder, directory service, registry service) that registers services and manages a binding table
    - Service lookup, registration, update, and deletion



- Addressing frequency
  - Each procedure call requires addressing
  - To reduce cost, binding result is cached and re-used

46

#### More issues: Synchronous vs. asynchronous RPC

- Synchronous RPC: strict request-reply behavior
  - RPC call blocks client that waits for server reply
- Some RPC middleware supports asynchronous RPC
  - Client continues without waiting for server reply
  - Server can reply as soon as request is received and execute procedure later



#### Asynchronous RPC

Call local procedure Time ->

Server

- Is RPC truly transparent? Can we really just treat remote procedure calls as local procedure calls?
  - Performance, failures, concurrent requests, replication, migration, …
- Performance
  - RPC is slower ... a lot slower: why?
  - Local call: maybe 10 cycles =  $\sim$ 3 ns
  - RPC: 0.1-1 ms on a LAN => ~100K slower
    - Major source of overhead: context switching, copies, interprocess communication
    - · In WAN: can easily be millions of times slower

48

#### More issues: transparency

- Failures
  - Different failures can occur
    - Client cannot locate server
    - Lost request messages
    - Server crashes
    - Lost reply messages
    - Client crashes

- Authenticate client? Authenticate server?
  - Is client sending messages to correct server or to impostor?
  - Is server accepting messages only from legitimate clients?
     Can server identify user at client side?
- Messages may be visible over network
  - Messages may be sniffed (and modified) while they traverse the network: can we encrypt them?
  - Have messages been accidentally corrupted or truncated while on network?
- RPC protocol may be subject to replay attacks
  - Can a malicious host capture a message and retransmit it at a later time?

50

### Programming with RPC

- Language support
  - Some programming languages have no language-level concept of remote procedure calls (e.g., C, C++)
    - Their compilers will not automatically generate stubs
  - Some languages directly support RPC (Java, Python, Haskell, Go, Erlang)
    - But we may need to deal with heterogeneous environments (e.g., Java service communicating with Python service)
- Common solution
  - Interface Definition Language (IDL): describes remote procedures
  - Separate IDL compiler generates client and server stubs

#### Interface Definition Language (IDL)

- IDL allows developer to specify remote procedure interfaces (*names, parameters, return values*) in a machine-independent way
- IDL compiler uses these interfaces to generate client and server stubs
  - Code to marshal
  - Code to unmarshal
  - Network transport code
- An IDL looks similar to function prototypes

Valeria Cardellini - SDCC 2024/25



- Sun RPC
- Java RMI
- Python RPyC
- Go
- gRPC

54

#### **RPC** implementation: Sun RPC

- First-generation RPC
- Created by Sun (now Oracle): Sun RPC
  - RFC 1831 (1995), RFC 5531 (2009)
  - Remains in use mostly because of NFS (Network File System)
- Interfaces defined in an IDL called XDR

- Sun RPC uses XDR (eXternal Data Representation) as IDL to address data heterogeneity
  - Standard to describe and encode machine-independent data (RFC 4506)
  - IDL compiler is rpcgen
- XDR provides built-in conversion functions for:
  - Predefined primitive types, e.g., xdr\_int()
  - Predefined structured types, e.g., xdr\_string()
- XDR is a binary format using *implicit typing* 
  - Implicit typing: only values are transmitted, not data types or parameter info

56

### Sun RPC: define RPC program using XDR

- Two descriptive parts written in XDR and grouped in a file with extension .x
  - 1. Definition: specifics of procedures (services) to identify procedures and their parameters' data types
  - 2. XDR definitions: definitions of parameters' data types (if not built-in)
- Our Sun RPC example: calculate square of integer number

#### Sun RPC example: define remote procedure

```
struct square_in { /* input (argument) */ square.x
long arg1;
};
struct square_out { /* output (result) */
long res1;
};
program SQUARE_PROG {
  version SQUARE_VERS {
  square_out SQUAREPROC(square_in) = 1; /* procedure number = 1 */
  } = 1; /* version number */
} = 0x31230000; /* program number */
```

- Define remote procedure SQUAREPROC
  - Each procedure has only one input parameter and one output parameter
  - Identifiers are written in uppercase
  - Each procedure is associated with a procedure number which is unique within RPC program (e.g., 1)

Valeria Cardellini - SDCC 2024/25

#### Sun RPC: how to implement RPC program

- Developer codes:
  - Client program: implements main() and logic needed to find remote procedure and bind to it (example: square\_client.c)
  - Server program: implements remote procedures provided by RPC server (example: square\_server.c)
- Note: developer does not write server-side main()

- Who calls remote procedure on server side?

Let's first consider standard local procedure

```
#include <stdio.h>
                                              square local.c
#include <stdlib.h>
struct square in { /* input (argument) */
  long arg;
};
struct square out { /* output (result) */
  long res;
};
typedef struct square in square in;
typedef struct square_out square_out;
square out *squareproc(square in *inp) {
  static square out out;
  out.res = inp->arg * inp->arg;
  return(&out);
}
```

```
Valeria Cardellini - SDCC 2024/25
```

60

#### Example: local procedure

```
int main(int argc, char **argv) {
  square_in in;
  square_out *outp;

  if (argc != 2) {
    printf("usage: %s <integer-value>\n", argv[0]);
    exit(1);
  }
  in.arg = atol(argv[1]);

  outp = squareproc(&in);
  printf("result: %ld\n", outp->res);
  exit(0);
}
```

• Local procedure (continue)

#### Which changes in case of remote procedure?

#### Sun RPC example: remote procedure

Remote procedure is similar to local one

```
#include <stdio.h> Server.c
#include <rpc/rpc.h>
#include "square.h" /* generated by rpcgen */
square_out *squareproc_1_svc(square_in *inp, struct svc_req
    *rqstp) {
    static square_out out;
    out.res1 = inp->arg1 * inp->arg1;
    return(&out);
}
• Notes:
```

- Input and output parameters use pointers
- Output parameter must be pointer to static variable (i.e., global memory allocation) so that pointed area exists when procedure returns
- Name of RPC procedure changes slightly (add \_ suffixed by version number and \_svc, e.g., \_1\_svc, all in lowercase)

```
Valeria Cardellini - SDCC 2024/25
```

```
62
```

#### Sun RPC example: client

```
    Run client with remote hostname and integer value; it

  calls remote procedure
#include <stdio.h>
                                                   client.c
#include <rpc/rpc.h>
#include "square.h" /* generated by rpcgen */
int main(int argc, char **argv) {
  CLIENT *clnt;
  char *host;
  square_in in;
  square out *result;
   if (argc != 3) {
    printf("usage: client <hostname> <integer-value>\n");
    exit(1);
                     CLIENT *clnt create(char *host, unsigned long prog,
                     unsigned number vers, char *proto)
  }
  host = argv[1]:
  clnt = clnt_create(host, SQUARE_PROG, SQUARE_VERS, "tcp");
```

```
if (clnt == NULL) {
    clnt_pcreateerror(host);
    exit(1);
    }
    in.arg1 = atol(argv[2]);
    if ((result = squareproc_1(&in, clnt)) == NULL) {
        printf("%s", clnt_sperror(clnt, argv[1]));
        exit(1);
    }
    printf("result: %ld\n", result->res1);
    exit(0);
}
```

64

### Sun RPC example: client

- **clnt\_create()**: creates client transport manager to handle communication with remote server
  - TPC or UDP, default timeout for request retransmission
- Client must know:
  - Remote server hostname
  - Info to call remote procedure: program name (SQUARE\_PROG), version number (1) and procedure name (square\_proc)
- To call remote procedure:
  - Procedure name changes slightly: add \_ followed by version number and write name in lowercase
  - Two input parameters:
    - · Effective input parameter plus client transport manager
  - Client gets pointer to result
    - To identify failed RPC: NULL return
- Handling of failures that may occur during remote call
  - clnt\_pcreateerror() and clnt\_perror()

- RPC program contains multiple remote procedures
  - Versioning support
  - Each procedure has one input and one output parameter
  - Call by copy-restore
- Transport independent
  - Transport protocol can be selected at run-time
- Mutual exclusion guaranteed by server
  - Default: no concurrency on server side
- Synchronous client: blocked until server replies
- At-least-once semantics
  - Request retransmission when timeout expires
- Security? Authentication mechanism added later with Secure RPC

66

#### Sun RPC: server binding

- Client stub needs to know port number: how?
- · Server stub registers RPC program
  - Each procedure is identified by: program number, procedure number, version number
- Where? In *port map* 
  - Dynamic table of RPC services on that host machine
  - port map is managed by *port mapper* (rpcbind): one per host, listens on port 111
  - Client stub contacts port mapper to find out port number and then sends request message to server stub

#### List RPC programs on a given host

>\$ rpcinfo -p

program		vers	proto	port	
	100000	4	tcp	111	rpcbind
	100000	4	udp	111	rpcbind
Γ	824377344	1	udp	59528	
	824377344	1	tcp	49311	

Our RPC program SQUAREPROC

- 824377344 (= 0x31230000) is the program number in square.x
- Server supports both TCP and UDP: transport-protocol independent

Valeria Cardellini - SDCC 2024/25

68

#### SUN RPC: development process



#### What goes on: server side

- Let's analyze server stub code (square\_svc.c)
- In main() server stub creates a socket and binds any available local port to it
- Calls svc\_register (RPC library function)
  - To register procedures with port mapper
    - Associates the specified program and version number pair with the specified dispatch routine
- Then waits for requests by calling svc\_run (RPC library function)
  - svc\_run invokes specific service procedures in response to RPC call messages



Valeria Cardellini - SDCC 2024/25

#### What goes on: client side

- When we start client program, clnt\_create contacts port mapper on server side to find port for that interface
  - Early binding: done once, not per each procedure call



- Client stub (square\_clnt.c) manages communication
  - Request timeout
  - Marshaling from local representation to XDR format and unmarshaling from XDR format to local representation

- Second-generation RPC
- Java RMI (Remote Method Invocation): RPC in Java
- Extends RPC to distributed objects
  - Allows to develop distributed applications in Java where an object on one JVM invokes methods on an object in another JVM
  - Goal: access transparency, but distribution transparency is still not full



#### **RMI:** basics

- Recall Java separation between definition (interface) and implementation (class)
- Idea: logical separation between interface and object allows for their physical separation
- Remote interface: specifies set of methods to be invoked remotely
- *Remote object*: instance of a class that implements a remote interface

- But internal state of remote object is not distributed!



- *Remote method invocation*: invoke methods of a remote interface on a remote object
  - Goal: keep same syntax as local invocation
  - How to achieve it?
- Once again, proxy pattern: client-side *stub* and server-side *skeleton* to hide distributed nature of application



## **RMI:** serialization/deserialization

- No need for external data representation
- Serialization/deserialization directly supported by Java
  - Thanks to Java bytecode, no need to (un)marshal, but data is (de)serialized using language-level features
- Serialization: converts object that is passed as parameter into byte stream
  - writeObject on output stream
- Deserialization: decodes byte stream and builds copy of original object
  - readObject from input stream
- Stub and skeleton use serialization/deserialization to exchange parameters and return values between different JVMs

#### Marshaling vs serialization

- · Loosely synonymous but semantically different
- Marshaling: stub converts local data into network data (using encoding/decoding routines) and packages network data for transmission
- Serialization: object state is converted into byte stream, which can be converted back into object copy
- Difference becomes noticeable for objects
  - Object codebase has also to be marshaled
  - Java serialization relies on codebase being present at receiver
- Different programming languages either make or don't make the distinction between the two
  - E.g., in Python (Pickle module) marshaling and serialization are considered the same, but not in Java

Valeria Cardellini - SDCC 2024/25

76

#### RMI: stub and skeleton interaction

- Client obtains object remote reference (i.e., stub instance) through *RMI registry*
  - Name server that relates objects with names
  - Server registers its remote objects on RMI registry so that they can be looked up
  - Client contacts registry to look an object up by name
- Client invokes remote method on stub
  - Syntax identical to local invocation
- Stub serializes data needed to invoke method (method ID and parameters) and sends them to skeleton in a message
- Skeleton receives message, deserializes received data, invokes method, serializes return value, and sends it to stub in a message
- Stub receives message, deserializes return value, and returns it to client
   Valeria Cardellini - SDCC 2024/25

### RMI example: echo remote interface

- Remote interface extends
   Remote
  - Remote identifies interfaces whose methods may be invoked from non-local JVMs
- Remote method
  - throws RemoteException
    - To handle communication failure or protocol error
    - · Remote method invocation is not fully transparent
  - passes parameters
    - by value in case of primitive data types (int, char, ...) or objects that implement java.io.Serializable interface: serialization/deserialization managed by stub/skeleton

}

 by reference in case of Remote objects: "reference" is not a pointer, it is a data structure: {IP address, port, time, object #, interface of remote object }

Valeria Cardellini - SDCC 2024/25

RMI example: echo server

}

- Class implements remote interface
  - Extends UnicastRemoteObject
  - super() calls class constructor
     UnicastRemoteObject which allows server to wait for requests and serve them
  - Implements remote method, which throws remote exception

```
import java.rmi.RemoteException;
```

import java.rmi.Remote;

public interface EchoInterface
 extends Remote{
 String getEcho(String echo)

throws RemoteException;

public class EchoRMIServer
 extends UnicastRemoteObject
 implements EchoInterface {

```
// Constructor
public EchoRMIServer()
 throws RemoteException {
   super();
```

```
// Implement remote method
public String getEcho(String echo)
  throws RemoteException {
    return echo;
  }
```

#### RMI example: echo server

```
public static void main(String[] args) {
                                      final int REGISTRYPORT = 1099;
                                      String registryHost = "localhost";
                                      String serviceName = "EchoService";
                                      try {

    In main server object

                                        // Create an instance of echo server
  instance is created
                                        EchoRMIServer serverRMI = new EchoRMIServer();
                                        // Bind remote object with RMI registry
• RMI registry is created
                                        Registry registry =

    Server registers by name its

                                          LocateRegistry.createRegistry(REGISTRYPORT);
  remote objects with RMI
                                      registry.bind(serviceName, serverRMI);
                                      }
  registry
                                      catch (Exception e) {

    RMI registry and server run

                                        System.err.println("EchoRMIServer exception: ");
       on same host
                                        e.printStackTrace();
                                      }
                                    }
                                  }
```

Valeria Cardellini - SDCC 2024/25

80

#### RMI example: echo client

•	Remote reference is obtained by <b>lookup</b> on server's RMI registry public class EchoRMIClient // Start RMI client public static void main(String[] args)
•	Client contacts RMI registry to look up the remote object using its name (registered by server) { bufferedReader stdIn = new BufferedReader( new InputStreamReader(System.in)); try
	- Registry returns a reference to remote object { // Connect to remote RMI service EchoInterface serverRMI = (EchoInterface)
•	Client invokes remote method - Synchronous blocking call System.out.print("Message? "); message = stdIn.readLine(); // Invoke remote service echo = serverRMI.getEcho(message); System.out.println("Echo: "+echo+"\n"); } catch (Exception e)
	<pre>{e.printStackTrace(); System.exit(1); }</pre>
	}

}

- Automatic generation of stub and skeleton
- Synchronous blocking calls
- At-most-once semantics
- Concurrency: remote method can be invoked concurrently by multiple clients

From Java RMI specification: "Since remote method invocation on the same remote object may execute concurrently, a remote object implementation needs to make sure its implementation is thread-safe" https://docs.oracle.com/en/java/javase/23/docs/specs/rmi/arch.html

- Define remote method as synchronized to protect it

 Dynamic class loading: since class definitions are required to serialize/deserialize objects passed as parameters, RMI provides dynamic class loading

Valeria Cardellini - SDCC 2024/25

#### **RMI:** architecture

• RMI system consists of 3 layers: *stub/skeleton* layer, *remote reference* layer, and *transport* layer



- Various RPC implementations
  - Pyro <u>https://pyro4.readthedocs.io</u>, RPyC, ZeroRPC
     <u>https://www.zerorpc.io</u>
- What helps Python achieve transparency
  - Inspection of live objects through inspect module <u>https://docs.python.org/3/library/inspect.html</u>
    - Examine class contents, retrieve method's source code, and extract function's argument list
- General idea of implementing RPC in Python
  - Create connection using RPC object
  - Invoke remote methods using that object
- Let's analyze RPyC (Remote Python Call)
   https://rpyc.readthedocs.io

84

#### **RPyC:** features

- Transparent RPC interface
  - Based on proxy pattern (once again)
  - No interface definition file, IDL compiler, name server for binding and lookup, transport service
  - Access to remote objects as if they were local
- Symmetric operations
  - Client and server can invoke RPCs on each other (e.g., server can invoke callbacks on client)
- Synchronous and asynchronous calls
- Secure
  - Capability-based security model
  - Integrates with TLS/SSL, SSH

#### **RPyC: server and client**

- Server
  - Multiple implementations, including ThreadedServer (thread for each connection)
    - Plus OneShotServer, ThreadPoolServer, ForkingServer
  - Binds to default port (18812, 18821 for SSL) or developer provides host's IP address and port within code
- Client
  - Connects to server
  - Performs remote operations through modules property, which exposes server module's namespace

Valeria Cardellini - SDCC 2024/25

86

#### **RPyC:** passing data

- By value
  - Simple types (immutable objects: string, int, tuple)
    - Sent directly to remote side
- By reference
  - Object: reference (object name) to object is passed
    - Remote server contacts client to access attributes and invoke methods on object
    - · Changes will be reflected onto actual object
  - Enables passing of location-sensitive objects, like files or other OS resources
    - Remote process can write to stdout of local process by getting its sys.stdout
  - Implementation: netrefs = transparent object proxies
    - Local objects that forward all operations to corresponding remote objects
    - They make remote objects look and feel like local objects

- Client creates a local proxy object for remote module
  - Allows for transparent access
  - Reference wrapped in special object called proxy that looks like the actual object
  - Any operation on proxy is delivered to target
  - Client is unaware of this
- Synchronous and asynchronous calls
  - Synchronous: client blocks and waits for return value
  - Asynchronous: immediate return, notification when complete
  - Calls can be made asynchronous by wrapping proxy with an asynchronous wrapper

88

#### **RPyC: services and security**

- RPyC is built around services
  - Each connections' end exposes a service that is responsible for the set of supported remote operations (aka policy)
- Services are classes that derive from rpyc.core.service.Service and define exposed methods
  - Methods whose names begin with exposed\_ or use @rpyc.exposed decorator
  - All exposed members of a service class are available to the other side

#### RPyC example: calculator server

```
import rpyc
class CalculatorService(rpyc.Service):
    def exposed_add(self, a, b):
        return a + b
    def exposed_sub(self, a, b):
        return a - b
    def exposed_mul(self, a, b):
        return a * b
    def exposed_div(self, a, b):
        return a / b
    def exposed_fib(self,n):
       seq = []
        a, b = 0, 1
        while a < n:
            seq.append(a)
            a, b = b, a+b
        return seq
    def foo(self):
        print("foo")
if name == " main ":
    from rpyc.utils.server import ThreadedServer
    t = ThreadedServer(CalculatorService, port=18861)
    print('Service started on port 18861')
    t.start()
```

90

#### **RPyC** example: calculator client

#### import rpyc

```
conn = rpyc.connect("localhost", 18861)
x = conn.root.add(4,7)
assert x == 11
print(conn.root.fib(1000))
# print(conn.root.div(1,0))
To remote party, service is
exposed as root object of
connection (conn.root).
This root object is a network
reference (netref) to service
instance living in server
process
```

### RPyC: async operation

- Key feature of RPyC
  - Client starts request and continues rather than blocking
  - Gets AsyncResult object that will eventually hold result
- Asynchronous behavior must be *explicitly* enabled: to turn the invocation of a remote method (or any callable object) asynchronous, wrap it with async\_()
- Then, client can
  - test AsyncResult object for completion using ready
  - wait for completion using wait()
  - get result using value
  - set timeout for result using set\_expiry()
  - register a callback function to be invoked when result arrives using add\_callback()
- No guarantee on execution order for async requests

See <u>https://rpyc.readthedocs.io/en/latest/docs/async.html</u> Valeria Cardellini - SDCC 2024/25

92

### RPyC: async operation and events

- Events can be implemented as asynchronous callbacks
  - Server produces events which are consumed by client
- Let's analyze FileMonitor example
  - Server periodically monitors file on client machine using os.stat() to detect changes
  - Whenever file is changed, server sends event to client (invoking async callback) and provides old and new stat results
  - At client side, incoming events are served by background thread



#### RPyC example: FileMonitor server

```
import rpyc
    import os
    import time
    from threading import Thread
    class FileMonitorService(rpyc.Service):
        class exposed_FileMonitor(object): # expose class
            def __init__(self, filename, callback, interval = 1):
                self.filename = filename
                self.interval = interval
                self.last_stat = None
                self.callback = rpyc.async_(callback) # create async callback
                self.active = True
                self.thread = Thread(target = self.work)
                self.thread.start()
            def exposed_stop(self): # this method has to be exposed too
                self.active = False
                self.thread.join()
            def work(self):
                while self.active:
                    stat = os.stat(self.filename) # get status of monitored file
                    if self.last_stat is not None and self.last_stat != stat:
                         self.callback(self.last_stat, stat) # notify client of change
                    self.last_stat = stat
                    time.sleep(self.interval) # monitor file status periodically
    if __name__ == "_
                      _main__":
        from rpyc.utils.server import ThreadedServer
        # spawn a thread for each connection
        ThreadedServer(FileMonitorService, port = 18871).start()
Valeria Cardellini - SDCC 2024/25
```

```
94
```

#### **RPyC example: FileMonitor client**

```
import rpvc
   import time
   import os
   # File to monitor
   filename = "/tmp/floop.bloop"
   if os.path.exists(filename):
        os.remove(filename)
   # Open file to monitor
   f = open(filename, "wb", buffering=0)
   conn = rpyc.connect("localhost", 18871)
   # Start background thread to process incoming events
   bgsrv = rpyc.BgServingThread(conn)
   # Callback function, obviusly written on client side
   def on_file_changed(oldstat, newstat):
       print("\nfile changed")
       print(f"
                   old stat: {oldstat}")
                    new stat: {newstat}")
       print(f"
   # Create file monitor and pass filename and callback function as arguments
   filemon = conn.root.FileMonitor(filename, on_file_changed)
   # Wait for file monitor to have a look at file
   time.sleep(2)
Valeria Cardellini - SDCC 2024/25
```

#### **RPyC example: FileMonitor client**

```
# Change file size and be notified by file monitor about change
print("change file size")
f.write(b"loop")
# Wait for file monitor to have a look at changed file
time.sleep(2)
# Change file size again
print("change size again")
f.write(b"groop")
time.sleep(2)
f.close()
filemon.stop()
bgsrv.stop()  # stop background thread
conn.close()
```

Valeria Cardellini - SDCC 2024/25

96

#### **RPC** in Go

- Let's introduce the Go programming language http://www.ce.uniroma2.it/courses/sdcc2425/slides/Go.pdf
- What about RPC in Go?

- Go standard library supports RPC right out-of-the-box
  - Package net/rpc <u>https://pkg.go.dev/net/rpc</u>
  - Provides access to the exported methods of an object across a network
- TCP or HTTP as "transport" protocols
- Requirements for server RPC methods
  - Method type and method are exported (capital letter)
  - Only two arguments, both exported
  - Second argument is a pointer to a reply struct that stores the corresponding data
  - An error is always returned

func (t \*T) MethodName(argType T1, replyType \*T2) error

Valeria Cardellini - SDCC 2024/25

98

#### **RPC** in Go: server

- On server side
  - Create a TPC server (or an HTTP server) to receive data
  - Use Register (or RegisterName): register an object, making it visible as a service

func (server \*Server) Register(rcvr any) error
func RegisterName(name string, rcvr any) error

- One input parameter, which is the interface: any is an alias for interface{}
- It publishes the methods that are part of the given interface on the RPC server and allows them to be called by clients connecting to the service
- Use Listen to announce on the local network address

 Use Accept to accept connections on the listener and serve requests for each incoming connection
 func (server \*Server) Accept(lis net.Listener)

- Accept blocks; if the server wishes to do other work as well, it should call this in a goroutine (go statement)
- Can also use HTTP handler for RPC messages
  - · See example on course site

Valeria Cardellini - SDCC 2024/25

100

#### RPC in Go: client

- On client side
  - Use **Dial** to connect to RPC server at the specified network address (and port)
- func Dial(network, address string) (\*Client, error)
  - Use DialHTTP for HTTP connection
  - Use Call to call synchronous RPC: Call waits for the remote call to complete

```
func (client *Client) Call(serviceMethod string, args
any, reply any) error
```

 Use Go to call asynchronous RPC: Go invokes the call asynchronously and signals completion using the Call structure's Done channel

```
func (client *Client) Go(serviceMethod string, args
any, reply any, done chan *Call) *Call
```

#### • On client side

```
Valeria Cardellini - SDCC 2024/25
```

102

#### RPC in Go example: calculator

- Let's consider a simple RPC calculator with two functions: multiply and divide two integers
- Code available on course site

#### RPC in Go: synchronous call

- Need some setup in advance of this...
- Call makes blocking RPC call
- Call invokes the remote function, waits for it to complete, and returns its error status

```
// Synchronous call
args := &server.Args{7,8}
var reply int
err = client.Call("Arith.Multiply", args, &reply)
if err != nil {
    log.Fatal("arith error:", err)
}
fmt.Printf("Arith: %d*%d=%d", args.A, args.B, reply)
func (client *Client) Call(serviceMethod string,
args any, reply any) error
```

```
Valeria Cardellini - SDCC 2024/25
```

104

#### RPC in Go: asynchronous call

- How to make asynchronous RPC? Go uses a channel as parameter to retrieve RPC reply when the call is complete
- Done channel will signal when the call is complete by returning the same object of Call
  - If Done is nil, Go will allocate a new channel

```
// Asynchronous call
quotient := new(Quotient)
divCall := client.Go("Arith.Divide", args, quotient, nil)
divCall = <- divCall.Done
// check errors, print, etc.</pre>
```

```
func (client *Client) Go(serviceMethod string, args any,
reply any, done chan *Call) *Call
```

• For Go internal implementation, see https://go.dev/src/net/rpc/client.go?s=8029:8135 - L284

- Go's RPC is at-most-once
- Client opens TCP connection and writes request
  - TCP may retransmit but server's TCP receiver will filter out duplicates internally, with sequence numbers
  - No request retry in Go's RPC code (i.e. will not create a second TCP connection)
- However, client returns error if it does not get reply
  - Perhaps after TCP timeout
  - Perhaps server did not see request
  - Perhaps server processed request but server or network failed before reply came back

106

### RPC in Go: marshaling and unmarshaling

- By default Go uses package encoding/gob for parameters marshaling (encode) and unmarshaling (decode) <u>https://pkg.go.dev/encoding/gob</u>
  - Package gob manages streams of gobs (Go binary values) exchanged between Encoder (transmitter) and Decoder (receiver)
  - A stream of gobs is *self-describing*: each data item in the stream is preceded by a specification of its type, expressed in terms of a small set of predefined types; pointers are not transmitted, but values they point to are transmitted
  - Basic usage: create encoder, transmit values, receive them with decoder
  - Requires that RPC client and server are both written in Go

#### RPC in Go: marshaling and unmarshaling

- Alternatives to gob
- 1. net/rpc/jsonrpc package in Go's standard library https://pkg.go.dev/net/rpc/jsonrpc
  - Implements a JSON-RPC 1.0 ClientCodec and ServerCodec for rpc package
- 2. gRPC

108

### RPC in Go example: Word count

- Let's consider a basic word count: given a document, count the occurrences of each word in that document
   – Tokenize and put words in a hash map
- Code available on course site

#### "hello I am good hello bye Bye"

Result: map[am:1 bye:3 good:1 hello:2 i:1]

#### From word count to MapReduce

- Word count: Tokenize and put words in a hash map
- How do we parallelize this?
- Partition document into *n* partitions
- Build *n* hash maps, one for each partition
- Merge the *n* hash maps (by key)
- How do you do this in a distributed environment?



Valeria Cardellini - SDCC 2024/25

110

#### From word count to MapReduce

 Partition input document into shards and assign shards to nodes



· Count words locally on each node



Valeria Cardellini - SDCC 2024/25

112

#### From word count to MapReduce

- · How to merge results computed locally?
- A first solution: send everything to one node What if data is too big? Too slow...
- The solution: partition key space among nodes (e.g, [a-e], [f-j], [k-p] ...)
- 1. Assign a key space to each node
- 2. Split local results by the key spaces
- 3. Fetch and merge results that correspond to the node's key space

- The solution: partition key space among nodes
  - 1. Assign a key space to each node
  - 2. Split local results by the key spaces



114

### From word count to MapReduce

- The solution: partition key space among nodes
  - 1. Assign a key space to each node
  - 2. Split local results by the key spaces



# The solution: partition key space among nodes

- 1. Assign a key space to each node
- 2. Split local results by the key spaces
- 3. Merge results received from other nodes



Valeria Cardellini - SDCC 2024/25

116

### MapReduce in a nutshell

- · Partition dataset into many shards
- Map stage: each node processes one or more shards locally
- Reduce stage: each node fetches and merges partial results from all other nodes
- But implementing MapReduce is not so easy...
  - Failures are common
  - Data skew causes unbalanced performance across cluster
  - System scale

#### Comparing RPC so far

- Let's compare RPC implementations
  - How do they differ in terms of distribution transparency?
  - Access transparency?
  - Location transparency?
  - Replication transparency?
    - Add server-side load balancer which acts as proxy between RPC clients and replicated RPC servers
  - Concurrency transparency?
  - Failure transparency?

Valeria Cardellini - SDCC 2024/25

118

#### Motivation for new RPC middleware

- Large-scale distributed applications composed of microservices
  - Microservices architecture: build sw application as a collection of *independent*, *autonomous* (developed, deployed, and scaled independently), *business capability– oriented*, and *loosely coupled* services
  - Multi-language (i.e., polyglot) development
  - Communication mainly structured as RPCs





- High-performance, open source universal RPC framework https://grpc.jo/
- Can run in any environment
  - Multi-language, multi-platform framework
- Main usage scenarios
  - Connect polyglot microservices that use request-reply communication style in and across data centers with pluggable support for load balancing, tracing, health checking and authentication
  - Connect devices, mobile apps and browsers to backend services
  - Generate efficient client libraries
- Developed by Google, now CNCF project
- Used by many companies and in many distributed systems

 E.g., Cisco, Cockroach Labs, Netflix, Square, etcd <u>https://etcd.io/</u> Valeria Cardellini - SDCC 2024/25

120

#### gRPC: main features

- HTTP/2 for transport
  - Bidirectional streaming, multiplexing, header compression

#### Protocol buffers as IDL •

- Simple service definition
- Automatic code generation
- Strict specification: prevents errors
- Plus authentication, flow control, blocking or nonblocking bindings, deadline/timeouts and cancellation
  - Deadline and timeouts allow clients to specify how long they are willing to wait for response

- Transport over HTTP/2
  - Basic idea of gRPC: treat RPCs as references to HTTP objects
- HTTP/2: major revision of HTTP that provides significant performance benefits over HTTP 1.x
- HTTP/2 in a nutshell
  - Binary framing layer: HTTP/2 request/response is divided into small messages and framed in binary format, making message transmission efficient



122

# gRPC: HTTP/2

#### • HTTP/2 in a nutshell

- From request/response messages to streams
  - Stream: bidirectional flow of bytes within an established connection, which may carry one or more messages
  - Message: complete sequence of frames that map to a logical request or response message
  - Frame: smallest unit of communication in HTTP/2, each containing a frame header, which at least identifies the stream to which the frame belongs
- Request/response multiplexing (usage of a single connection per client): allows for efficient use of TCP connections and avoids head-of-line blocking at HTTP level
- Native support for bidirectional streaming
- HTTP header compression: to reduce protocol overhead

See <a href="https://hpbn.co/http2/">https://hpbn.co/http2/</a>

- gRPC uses protocol buffers (aka protobufs) as:
  - IDL to define service interface: automatic generation of client stubs and abstract server classes
  - Message interchange format: gRPC messages are serialized using protocol buffers, thus resulting in small message payloads
- Based on proxy pattern: *stub* and *server*



Protocol buffers

- Google's mature open-source mechanism to serialize structured data
- Binary data representation
- Strongly typed

```
message Person {
   string name = 1;
   int32 id = 2;
   bool has_ponycopter = 3;
}
```

- Data types are structured as messages
  - message: small logical record of information containing a series of name-value pairs called *fields*
  - Fields have unique field numbers (e.g., string name = 1), used to identify fields in message binary format

#### Protocol buffers: example



#### gRPC: basic steps

- 1. Define service (collection of remote methods) and message types that are exchanged between client and service in .proto file using protobufs as IDL
- 2. Generate server and client code using protoc (protocol buffer compiler) in your preferred language(s) from proto definition
  - Go: compile manually using protoc command
  - Java: use build automation tools (e.g., Maven, Gradle)
- 3. Use gRPC API in your preferred language (e.g., Go, Java, Python) to write service client and server
  - Let's consider Go: gRPC-Go https://grpc.io/docs/languages/go

#### gRPC: helloworld example in Go

 See <a href="https://grpc.io/docs/languages/go/quickstart/">https://grpc.io/docs/languages/go/quickstart/</a> and <a href="https://github.com/grpc/grpc-go/tree/master/examples/helloworld/">https://github.com/grpc/grpc-go/tree/master/examples/helloworld/</a>

1. Define service (helloworld.proto file)

```
package helloworld;
```

```
// The greeting service definition.
service Greeter {
    // Sends a greeting
    rpc SayHello (HelloRequest) returns (HelloReply) {}
  }
  // The request message containing the user's name.
message HelloRequest {
    string name = 1;
  }
  // The response message containing the greetings
message HelloReply {
    string message = 1;
  Valeria dardellini - SDCC 2024/25
```

128

### gRPC: helloworld example in Go

- 2. Compile service definition:
- \$ protoc --go\_out=. --go\_opt=paths=source\_relative \
   --go-grpc\_out=. --go-grpc\_opt=paths=source\_relative \
   helloworld/helloworld.proto
- Automatically generated files:
  - helloworld.pb.go: contains protobuf code to populate, serialize, and retrieve request and response message types
  - helloworld\_grpc.pb.go: contains
    - interface type (or *stub*) for clients to call with methods defined in Helloworld service
    - interface type for servers to implement, also with methods defined in Helloworld service
    - uses message definitions given in helloworld.pb.go
  - Take a look at those files and see that
    - · message types have become Go structs
    - · RPC definitions have become Go interfaces

#### gRPC: helloworld example in Go

```
Create server: composed of two parts
3.
       Implement service interface generated from service
   а.
       definition: the actual "work"
   func (s *server) SayHello(ctx context.Context,
      in *pb.HelloRequest) (*pb.HelloReply, error) {
   }
   b.
       Create and run gRPC server to listen for requests from
       clients and dispatch them to service implementation
   lis, err := net.Listen("tcp", port)
   if err != nil {
       log.Fatalf("failed to listen: %v", err)
   }
   s := grpc.NewServer()
   pb.RegisterGreeterServer(s, &server{})
   s.Serve(lis)
```

Valeria Cardellini - SDCC 2024/25

130

#### gRPC: helloworld example in Go

#### 4. Create client

 To call service methods, first create a *gRPC channel* to communicate with server using Dial

```
conn, err := grpc.Dial(address, opts...)
```

- We need a client stub to perform RPCs: get it using pb.NewGreeterClient provided by pb package (generated from .proto file)
- c := pb.NewGreeterClient(conn)
- Call service methods on client stub: create and populate a request protobul object (HelloRequest) and pass a context object which lets us change RPC's behavior if necessary (e.g., time-out/cancel RPC in flight)
- r, err := c.SayHello(ctx, &pb.HelloRequest{Name: name})

#### gRPC: update application components

- gRPC simplifies application updates
- If you want make changes to server API (e.g., adding new method)
  - 1. Update service definition (.proto file)
  - 2. Regenerate stubs so that client and server implementations can reflect changes
  - 3. Update server code to implement new method
  - 4. Update client code to call new method

Valeria Cardellini - SDCC 2024/25

```
132
```

#### gRPC: update helloworld example in Go

- Let's update gRPC service adding new method SayHelloAgain()
  - 1. Update .proto file
  - 2. Regenerate gRPC code using protoc
  - 3. Update server code to implement new method

```
func (s *server) SayHelloAgain(ctx context.Context, in
 *pb.HelloRequest) (*pb.HelloReply, error) {
    log.Printf("Received: %v", in.GetName())
    return &pb.HelloReply{Message: "Hello again " +
    in.GetName()}, nil
}
4. Update client code to call new method
r, err = c.SayHelloAgain(ctx, &pb.HelloRequest{Name: name})
if err != nil {
    log.Fatalf("could not greet: %v", err)
  }
  log.Printf("Greeting: %s", r.GetMessage())
```

#### gRPC: ProductInfo example

- Online retail having ProductInfo microservice which manages products and their information
- · Clients can add and retrieve products
- 1. Define service in .proto file (see slide 128)
- 2. Implement server and client in Go
- 3. Implement server and client in Java

Go code: Teams/course website Java code: <u>https://github.com/grpc-up-and-running/samples/tree/master/ch02</u> See chapter 2 of gRPC: Up and Running https://www.oreilly.com/library/view/grpc-up-and/9781492058328/

134

### gRPC: types of RPC methods

- gRPC supports 4 kinds of service methods
  - Defined in .proto file
  - See routeguide example <u>https://github.com/grpc/grpc-go/tree/master/examples/route\_guide</u>
- 1. Simple RPC: client sends request to server and waits for single response to come back (i.e., unary)
  - A normal function call

rpc SayHello (HelloRequest) returns (HelloReply) {}

- 2. Server-streaming RPC: client sends request to server and gets stream to read a sequence of messages back
  - Client reads from stream until there are no more messages
  - gRPC guarantees message ordering within individual RPC call
- rpc ListFeatures(Rectangle) returns (stream Feature) {}

#### gRPC: types of RPC methods

3. Client-streaming RPC: client writes sequence of messages and sends them to server, waits for server to read them and return its response

 gRPC guarantees message ordering within individual RPC call rpc RecordRoute(stream Point) returns (RouteSummary) {}

- 4. Bidirectional streaming RPC: both sides send sequence of messages using read-write stream (i.e., full duplex)
  - The two streams operate independently, so client and server can read and write messages as preferred (server can wait until it has received all client's messages before writing its messages, or server and client can play "ping-pong")
  - gRPC preserves message ordering in each stream

rpc RouteChat(stream RouteNote) returns (stream RouteNote) {}

Valeria Cardellini - SDCC 2024/25

136

#### gRPC: distribution transparency

- What about distribution transparency?
  - Access transparency?
    - The usual proxy pattern
  - Location transparency?
    - gRPC supports service discovery mechanisms to locate services dynamically
    - By default, DNS as name resolver
       <u>https://grpc.io/docs/guides/custom-name-resolution/</u>
  - Concurrency transparency?
    - Language dependent, for Go see <a href="https://github.com/grpc/grpc-go/blob/master/Documentation/concurrency.md">https://github.com/grpc/grpc-go/blob/master/Documentation/concurrency.md</a>
  - Replication transparency?
    - gRPC can integrate with load balancers to distribute client requests evenly across multiple server instances https://grpc.io/docs/guides/custom-load-balancing

#### gRPC: distribution transparency

- What about distribution transparency?
  - Failure transparency?
    - At-most-once by default
    - gRPC can support implementing retry strategies for transient failures, including *exponential backoff* <u>https://grpc.io/docs/guides/retry/</u>
    - gRPC supports request hedging: sending multiple copies of same request without waiting for response https://grpc.io/docs/guides/request-hedging/

See: Dean and Barroso, The Tail at Scale, 2013 https://research.google/pubs/the-tail-at-scale/



Valeria Cardellini - SDCC 2024/25

138

### gRPC: other features

- Security
  - gRPC supports SSL/TLS
  - How to use
    - Need to generate SSL/TLS certificates (for development only: self-signed certificates using OpenSSL)
    - On server side: configure gRPC server to use SSL/TLS certificates
    - On client side. configure gRPC client to trust server's SSL/TLS certificates
    - Example: <u>https://github.com/grpc/grpc-go/tree/master/examples/features/encryption</u>
- · Benchmarking and load testing
  - ghz tool https://ghz.sh

- Limited support in browsers
  - Cannot directly call a gRPC service from a browser
  - gRPC-Web can provide gRPC support to browser but limited features (only simple RPC and limited server streaming) <u>https://grpc.io/docs/platforms/web/basics/</u>
- Non-human readable format
  - Protobuf is efficient to send and receive, but its binary format is not human readable
  - Developers need additional tools to analyze protobuf payloads on the wire, write manual requests, and perform debugging
    - E.g., gRPC command-line tool or Wireshark https://wiki.wireshark.org/gRPC

140

#### References

- Chapter 4 of van Steen & Tanenbaum book
- Remote Procedure Calls <a href="https://www.linuxjournal.com/article/2204">https://www.linuxjournal.com/article/2204</a>
- Java Remote Method Invocation Specification
   <a href="https://docs.oracle.com/en/java/javase/23/docs/specs/rmi/">https://docs.oracle.com/en/java/javase/23/docs/specs/rmi/</a>
- Trail: RMI <u>https://docs.oracle.com/javase/tutorial/rmi/</u>
- RPyC tutorial <a href="https://rpyc.readthedocs.io/en/latest/tutorial.html">https://rpyc.readthedocs.io/en/latest/tutorial.html</a>
- RPyC documentation <u>https://rpyc.readthedocs.io</u>
- Go's package rpc <u>https://pkg.go.dev/net/rpc</u>
- RPC in the Go standard library, in Building Microservices with Go <u>https://subscription.packtpub.com/book/web-</u> development/9781786468666/1/ch01lvl1sec12/rpc-in-the-go-standardlibrary
- gRPC Up and Running <a href="https://www.oreilly.com/library/view/grpc-up-and/9781492058328/">https://www.oreilly.com/library/view/grpc-up-and/9781492058328/</a>