













<u>Client</u>	<u>Server</u> socket – create socket bind – assign address, port listen – listen for clients
socket - creat bind* - assign	te socket n address (optional) unect to listening socket
	accept – accept connection
Both can read Both can call c	and write from the connection.
	Convictor © 2015 Thomas W Doonnest Redside

## Go's Interface is not too different

```
Server:
     ln, err := net.Listen("tcp", ":8080")
     if err != nil {
          // handle error
     }
     for {
          conn, err := ln.Accept()
          if err != nil {
               // handle error
          }
          go handleConnection(conn)
     }
                                             Copyright © 2015 Thomas W. Doeppner, Rodrigo Fonseca.
CS 138
                                 VI–9
                                                       Some content from David Andersen.
```

## Go's Interface is not too different

```
Client:
```

```
conn, err := net.Dial("tcp", "google.com:80")
if err != nil {
    // handle error
}
fmt.Fprintf(conn, "GET / HTTP/1.0\r\n\r\n")
status, err :=
bufio.NewReader(conn).ReadString('\n')
// ...
CS 138 VI-10
Copyright@2015 Thomas W. Deepprer, Rodrige Fonseca.
Some content from David Andersen.
```











## <section-header><section-header><section-header><section-header><code-block><code-block><code-block><page-footer><page-footer></code></code></code>









The basic theory of operation of RPC is pretty straightforward. But, to understand *remote* procedure calls, let's first make sure that we understand *local* procedure calls. The client (or caller) supplies some number of arguments and invokes the procedure. The server (or callee) receives the invocation and gets a (shallow) copy of the arguments (other languages, such as C++, provide other argument-passing modes, but copying is all that is provided in C). In the usual implementation, the callee's copy of the arguments have been placed on the runtime stack by the caller—the callee code knows exactly where to find them. When the call completes, a return value may be supplied by the callee to the caller. Some of the arguments might be out arguments—changes to their value are reflected back to the caller. This is handled in C indirectly—the actual argument, passed by copying, is a pointer to some value. The callee follows the pointer and modifies the value.



Now suppose that the client and server are on separate machines. As much as possible, we would like remote procedure calling to look and behave like local procedure calling. Furthermore, we would like to use the same languages and compilers for the remote case as in the local case. But how do we make this work? A remote call is very different from a local call. For example, in the local call, the caller simply puts the arguments on the runtime stack and expects the callee to find them there. In C, the callee returns data through out arguments by following a pointer into the space of the caller. These techniques simply don't work in the remote case.



The solution is to use *stub procedures*: the client places a call to something that has the name of the desired procedure, but is actually a proxy for it, known as the *client-side stub*. This proxy gathers together all of the arguments (actually, just the in and in-out arguments) and packages them into a message that it sends to the server. The server has a corresponding *server-side stub* that receives the invocation message, pulls out the arguments, and calls the actual (remote) procedure. When this procedure returns, returned data is packaged by the server-side stub into another message, which is transmitted back to the client-side stub, which pulls out the data and returns it to the original caller. From the points of view of the caller and callee procedure, the entire process appears to be a local procedure call—they behave no differently for the remote case.



We will see later how these are generated.



Remote Procedure Calls (1)			
	•	A remote procedure call occurs in the following steps:	
	1.	The client procedure calls the client stub in the normal way.	
	2.	The client stub builds a message and calls the local operating system.	
	3.	The client's OS sends the message to the remote OS.	
	4.	The remote OS gives the message to the server stub.	
	5.	The server stub unpacks the parameters and calls the server.	
	Co	ontinued	
CS 138		VI–25 Content from David Andersen	









So, what are the differences between local and distributed, and can we create a proper illusion?






















Consider this slide first with the assumption that RPC is layered on UDP. Thus, since the response acts as the acknowledgement, there is uncertainty as to whether the request was handled by the server.

Does this uncertainty go away if RPC is layered on TCP? If you consider the possibility that the TCP connection might be lost, perhaps due to a transient network problem, the answer is clearly no. For example, suppose the TCP connection is lost just after the server receives the request. With no connection, the server cannot send a response, so the client is uncertain about what happened.



A procedure is *idempotent* if the effect of executing it twice in a row is the same as executing it just once. With such procedures, the client may repeatedly send a request until it finally gets a response. If an RPC protocol depends on such retries, it is said to have *atleast-once semantics* — clients are assured that, after all the retries, the remote procedure is executed at least once.



Not everything is idempotent! If we have non-idempotent procedures, then RPC requests should not be blindly retried, but instead should be sent just once. RPC protocols that do this are said to have *at-most-once semantics*. DCE RPC guarantees at-most-once semantics by default, though a remote procedure may be declared (in its IDL description) to be idempotent, in which case calls are done using at-least-once semantics.



The server might keep track of what operations it has already performed and what the responses were. If it gets a repeat of a previous request, it merely repeats the original response.



If the server crashed and no longer has its history information, it can respond by raising an exception at the client, indicating that it has no knowledge as to whether the operation has taken place. But it guarantees that it hasn't taken place more than once.

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 reply	At-most-once

This table is Figure 5.9 of the text. Note the distinction made between "maybe" and "at-most-once" semantics.

















	Example				
typedef stru int double long lou char	uct { comp1; comp2; ng  comp3[6]; *annotation;				
typedef struver value_t value_t } list_t;	uct { element; * next;				
<b>char</b> add( <b>in</b> <b>char</b> remove <b>list_t</b> query	t key, <b>value_t</b> value); e( <b>int</b> key, <b>value_t</b> valu ( <b>int</b> key);	e);			
:S 138	VI–54	Copyright © 2015 Thomas W. Doeppner, Rodrigo Fonseca. Some content from David Andersen.			

Here's an example of a C declaration for a collection of simple database procedures.



Here's a specification of our database interface written in the XDR language. This can now be compiled by *rpcgen* into a pair of stubs, server and client. See http://docs.oracle.com/cd/E19120-01/open.solaris/816-1435/rpcgenpguide-24243/index.html for a description of how to use rpcgen.



XDR is based on a number of primitive types, whose representation is fixed. We won't go into exactly what these representations are, see RFC 1832 (http://www.faqs.org/rfcs/rfc1832.html) for details.



One builds more complicated data types from the primitives by using the XDR's structured type constructors. We mention the most important ones on the slide; we see them in use in the next several slides.

Generate	ed Header File	
struct value {	<b>typedef list</b> list_t;	
<pre>int comp1;</pre>		
double comp2;	<pre>struct add_1_argument {</pre>	
<b>int64_t</b> comp3[6];	int key;	
<b>char</b> *annotation;	<b>value_t</b> value;	
};	};	
typedef struct value value;	<b>typedef struct add_1_argument</b> add_1_argument;	
<b>typedef value</b> value_t;		
	<pre>struct remove_1_argument {</pre>	
<b>struct</b> list {	<b>int</b> key;	
<b>value_t</b> element;	<b>value_t</b> value;	
<pre>struct list *next;</pre>	};	
};	typedef struct remove_1_argument	
typedef struct list list;	remove_1_argument;	
S 138	VI–58 Copyright © 2015 Thomas W. Doeppner, Rodrigo F Some content from David An	

٦

After running *rpcgen* on our XDR description, we get our stubs, code for marshalling and unmarshalling, and a common header file, shown here.



Here the client places calls to each of the procedures in our remote program.





Now we re-do the example using DCE RPC.



DCE RPC uses an augmented C syntax known as *interface definition language* (IDL) to express its interfaces. By compiling a description written in IDL (using a special IDL compiler), one automatically produces the client- and server-side stubs.

Here we have the IDL description for the simple database example of the previous slide. It starts with the declaration of a data type that is that of the somewhat complex items stored in the database. Following this are the specifications of the three (remote) procedures that clients may call—one for adding values to the database, one for removing values, and one for issuing queries. Values are entered into the database in association with keys. The query operation returns all values, up to an indicated maximum number, that share the given key.

One of the purposes of IDL is to overcome various shortcomings of the C syntax for declaring procedures. Among the issues are:

- Which arguments are input arguments, which are output arguments, and which are both? There is no way to determine this from standard C syntax. With IDL, we have new attributes, enclosed in square brackets, that identify the use of the arguments.
- What is an integer? C has three signed integer types: *char, short, and long.* One can also declare something to be an *int,* but, depending on the architecture, it will either be a *short* or a *long.* (Of course, we must at some point deal with 64-bit architectures. IDL has a 64-bit integer type called *hyper.*) *Short* and *long* are pretty straightforward, but what is a *char*? Its name certainly implies some ambiguity. To eliminate this ambiguity, in IDL, if one wants an 8-bit signed integer, one calls it a *small.*



- If we pass an array, how big is it? From our knowledge of the program, we can see that one of the arguments is the size of the two arrays, but how would the IDL compiler know this? It must know the size of the arrays, so that it can determine how much data to pass to the server (and how much data the server should pass to the client for output arguments). In the example here, we use the *size\_is* attribute to notify the IDL compiler that one of the arguments is the size of the array.
- What does it mean to pass a pointer argument? For example, if one of the arguments is declared *char* \* what would that mean? We certainly don't want to follow a pointer on the server back to the client, so we should pass the data pointed to along with the pointer. But how much should be passed? Is a *char* \* really a pointer to a single character? Is it a pointer to a null-terminated string? Is it a pointer to a counted array of bytes? The IDL syntax allows us to distinguish these cases (and supply whatever additional information is necessary).
- When we pass strings of characters, which character set are we referring to? ASCII is only acceptable (and just barely) in English-speaking countries (for which a seven-bit character set suffices). In the example we specify ISO\_LATIN\_1, an eight-bit character set, which is suitable for most of Europe and the Americas, but for little of Asia, which needs a sixteen-bit (at least) character set.

There are alternatives to this approach to handling the data types of parameters. One popular approach, used in a early RPC protocol developed at Xerox in the '70s and now popularized in Microsoft's .Net and other recent systems, is, rather than have linked-in stubs that "know" the types of parameters, to send the typing information along with the parameters. Thus, rather than simply sending the integer "6", what would be sent is "INT 6". This is currently being done in conjunction with XML.



To represent an array, we need to include its length.



Marshalling pointers is sometimes pretty simple: one simply transmits the target of the pointer, rather than the pointer. Unmarshalling depends on whether the receiver is a callee (i.e, the pointer is an input parameter) or is a caller (the pointer is an output parameter). For a callee, the pointed-to item is copied into storage on the receiver's stack in the server-side stub's frame; the stub passes a pointer to the item to the remote procedure. For a caller, the marshaled item is copied into the original target of the pointer; the pointer itself doesn't change.



Three situations can complicate the marshalling of pointers. The first is when a pointer contains a null value: since it's pointing at nothing, there's nothing to send! The second is when two different pointers point to the same location (this is known as *aliasing*). It's not enough to send the value of what the pointer points to: the reconstructed pointers on the receiver must also point to the same location. Lastly, what if the pointer points to a data structure containing another pointer?



Marshalling unrestricted pointers, i.e., pointers that might be null, might be aliased, or might point to data structures containing other pointers, requires that one send a representation of how the data structure is organized. One such representation is illustrated in the slide: what the pointers point to is represented as an array and the pointers are represented as indices of the array.



## Maintaining Client State on Servers

```
interface trees {
   typedef [context_handle] void *tree_t;
   void create (
       [in] long value,
       [out] tree_t pine
   );
   void insert (
       [in] long value,
       [in, out] tree_t pine
   );
}
CS 138 VI-69
Copyright © 2015 Thomas W. Deepmer. Rodrigo Fonseca.
   Some content from David Andersene
```

Rather than pass a tree back and forth between client and server, it might make more sense to leave the tree on the server and have the client merely send the server requests to perform operations on it. The interface shown here has two procedures—one to create and initialize a tree, and another to add new items to it. From the client's point of view, the tree is represented as an opaque pointer of type *context handle*. This is created implicitly (on both client and server) via the use of the *out* parameter of the *create* procedure. The client holds onto this handle; whenever it uses the handle with the *insert* procedure, it is converted on the server side to point to whatever the pointer pointed to that the server originally returned via the *out* parameter of the *create* procedure.

If the server crashes, then, from the client's point of view, the context handle becomes useless. The client will be notified of a server failure if it tries to use the context handle after the server is known to have crashed.

If the client crashes, then the server might want to be notified, especially if the client is the only one interested in the tree represented by the context handle. In the event of a crash, the server runtime will clean up its state. If the server is interested, it can define a cleanup (or *rundown*) procedure. The name of the cleanup procedure in this case would be *tree\_t\_rundown*, which will be called with the server-side pointer (to the tree) as the argument. The server, in this example, would free the storage that had been allocated for the tree.



This go example from the gorpc page: http://golang.org/pkg/net/rpc/

## **Server Startup**

```
arith := new(Arith)
rpc.Register(arith)
rpc.HandleHTTP()
1, e := net.Listen("tcp", ":1234")
if e != nil {
    log.Fatal("listen error:", e)
}
go http.Serve(1, nil)
```

CS 138

VI–71

Copyright © 2015 Thomas W. Doeppner, Rodrigo Fonseca. Some content from David Andersen.


## **Client Call**

```
// 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)
CS 138 VI-73
```

## **Client Call**

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

CS 138

VI-74

Copyright © 2015 Thomas W. Doeppner, Rodrigo Fonseca. Some content from David Andersen.



This example (extending through the next five slides) is taken from http://download.oracle.com/ javase/tutorial/rmi/overview.html (and the code is copyrighted by Oracle). The idea is that we are designing a "compute server" to which we can send an object representing a computation to be performed. The server performs the computation and sends back the result.

We start with a declaration of the server's interface. By extending Remote, it becomes the interface to a remote object, meaning that clients with references to the object can invoke its methods remotely. Note that all methods of remote objects must be declared as throwing *RemoteException*, which occurs when there is some sort of problem, such as a communication error.

The interface provides one method, that takes a parameterized type, Task<T> as an argument, and returns something of type T. Thus effectively the method has two parameters: Task<T> and T.



Here is the declaration of Task<T>. Note that it is not a remote object, but will be passed (by copying) to the Compute object.

```
RMI: the Server (1)
package engine;
import java.rmi.RemoteException;
import java.rmi.registry.LocateRegistry;
import java.rmi.registry.Registry;
import java.rmi.server.UnicastRemoteObject;
import compute.Compute;
import compute.Task;
public class ComputeEngine implements Compute {
    public ComputeEngine() { super(); }
    public <T> T executeTask(Task<T> t) {
         return t.execute();
     }
                                           Copyright © 2015 Thomas W. Doeppner, Rodrigo Fonseca.
CS 138
                                VI-77
                                                     Some content from David Andersen.
```

Now that we've seen the declaration of the interface, we look at the server itself, whose code begins on this slide. The class Compute Engine implements the Compute interface. Its constructor simply calls upon the superclass constructor. Its executeTask method simply invokes the execute method of the supplied Task<T> argument.



The main routine, which is static, after setting up a security manager (which restricts what remote invocations of the object can do) creates an instance of a *ComputeEngine* object and makes it available for clients to access. The first step in this is to set things up so that the object can receive invocations of its method from remote clients. The *exportObject* method of *UnicastRemoteObject* does this: the first argument is the remote interface being offered to clients and the second is the port on which to receive requests (0 means to use the default port).

Note that remote interfaces do not have constructors, thus clients must be given some external means for getting references to remote objects. This is done here by putting the name of the object in the local registry along with the reference to the remote object provided by *exportObject*. Thus clients can contact the registry and get the reference associated with the name.



Here we have code that's run by clients. As with servers, a security manager must be set up to enforce restrictions on what objects can do that are returned by servers. The client then gets the object reference from the server's registry (the name of the server is passed to the client code as an argument). It then creates an instance of an object that represents a computation to be done (in this case: compute the value of pi to a given number of digits). This object is passed to the remote object as an argument and the result of the computation is returned.

```
RMI: the Client's Compute
Object
package client;
import compute.Task;
import java.io.Serializable;
import java.math.BigDecimal;
public class Pi implements Task<BigDecimal>, Serializable {
    private final int digits;
    public Pi(int digits) {this.digits = digits;} // constructor
    // lots of stuff deleted ...
    public BigDecimal execute() {
         return computePi(digits);
    }
    // lots more stuff deleted \ldots
}
                                               Copyright © 2015 Thomas W. Doeppner, Rodrigo Fonseca.
CS 138
                                   VI-80
                                                          Some content from David Andersen.
```

Lastly we have the skeleton of the object that computes pi. Note that the object must implement *Serializable* so that it can be marshalled and unmarshalled.

