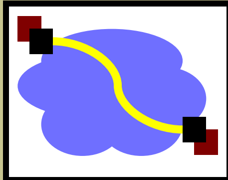

15-440 Distributed Systems

Lecture 06 – Remote Procedure Calls

Thursday, September 13th, 2018

Announcements

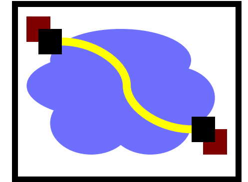

- P0 Due today (Thursday 9/13)
  - How is everyone doing on it? :-)
- P1 Released Friday 9/14
  - Dates: Recitation (9/19), Part A (10/6), Part B (10/16)
  - Group Project – teams of 2. Do you have a partner?
  - 6-week drop deadline 10/8 => discuss with partner!
- No office hours Today – P0 due
- Office Hour Protocol – Piazza Survey responses

Building up to today

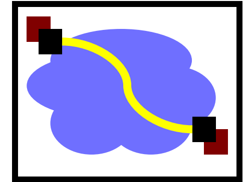

- Abstractions for *communication*
  - example: TCP masks some of the pain of communicating across unreliable IP
- Abstractions for *computation*

Splitting computation across the network

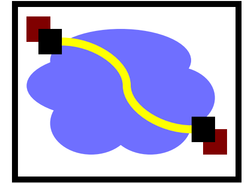

- We've looked at primitives for computation and for communication.
- Today, we'll put them together

  Key question:

What programming abstractions work well to split work among multiple networked computers?

Common communication pattern

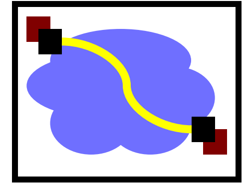

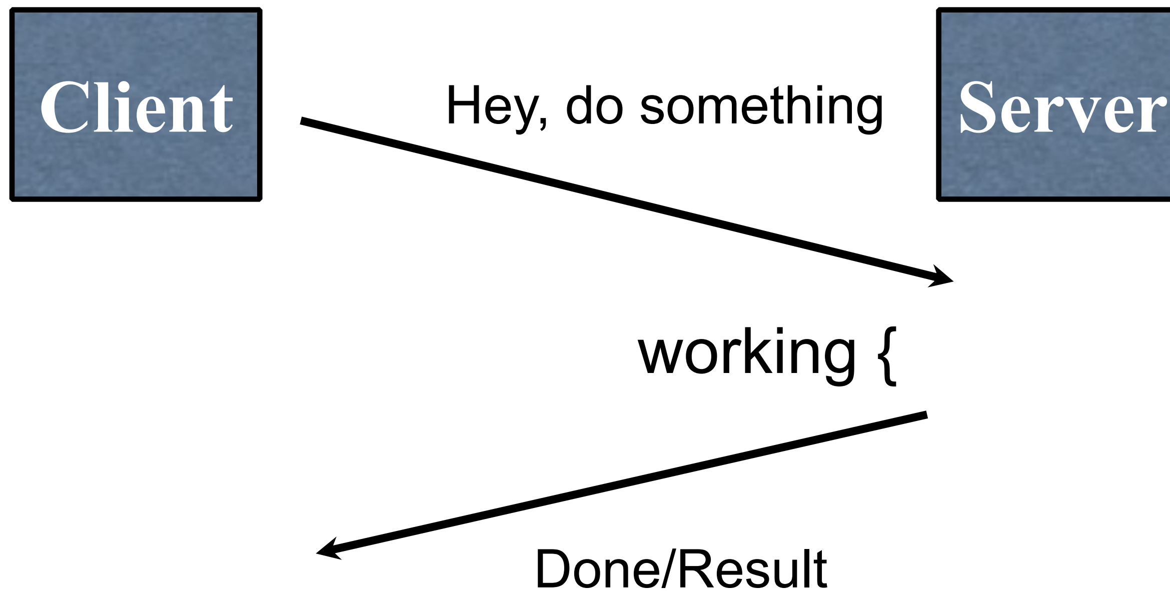

Writing it by hand...

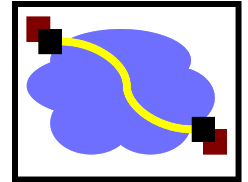

- E.g., if you had to write a, say, password cracker

```
struct foomsg {
  u_int32_t len;
}

send_foo(char *contents) {
   int msglen = sizeof(struct foomsg) + strlen(contents);
   char buf = malloc(msglen);
   struct foomsg *fm = (struct foomsg *)buf;
   fm->len = htonl(strlen(contents));
   memcpy(buf + sizeof(struct foomsg),
          contents,
          strlen(contents));
   write(outsock, buf, msglen);
}
```

Then wait for response, etc.

Today's Lecture

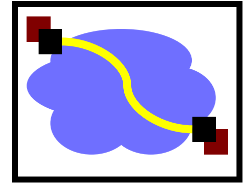

- RPC overview
- RPC challenges
- RPC other stuff

RPC – Remote Procedure Call

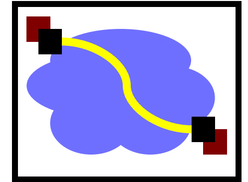

- A type of client/server communication
- Attempts to make remote procedure calls look like local ones

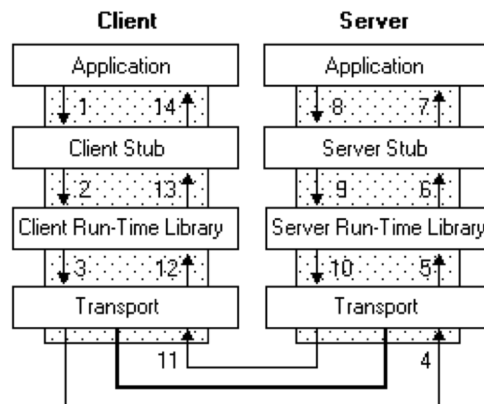

figure from Microsoft MSDN

```
{ ...
   foo()
}
void foo() {
  invoke_remote_foo()
}
```

Go Example

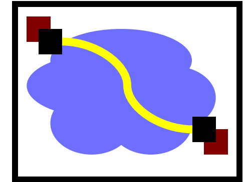

- Client first dials the server

```
client, err := rpc.DialHTTP("tcp", serverAddress + ":1234")
if err != nil { log.Fatal("dialing:", err) }
```

- Then it can make a remote 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)
```

Server Side

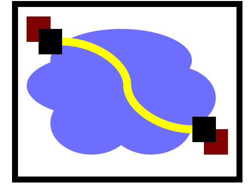

- Basic RPC code:

```
package server
type Args struct { A, B int }
type Quotient struct { Quo, Rem int }
type Arith int
func (t *Arith) Multiply(args *Args, reply *int) error {
    *reply = args.A * args.B
    return nil }
func (t *Arith) Divide(args *Args, quo *Quotient) error {
    if args.B == 0 { return errors.New("divide by zero") }
    quo.Quo = args.A / args.B
    quo.Rem = args.A % args.B
    return nil }
```

- The server then calls (for HTTP service):

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

RPC Goals

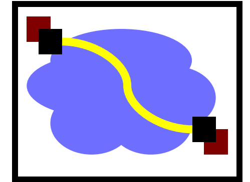

- Ease of programming
- Hide complexity
- Automates task of implementing distributed computation
- Familiar model for programmers (just make a function call)

Historical note: Seems obvious in retrospect, but RPC was only invented in the '80s. See Birrell & Nelson, "Implementing Remote Procedure Call" ... or
Bruce Nelson, Ph.D. Thesis, Carnegie Mellon University: Remote Procedure Call., 1981 :)

Conventional Procedure Call

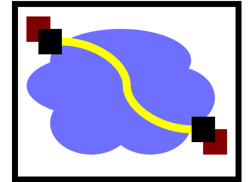

Stack pointer

Main program's local variables

(a)

Main program's local variables
nbytes
buf
fd
return address
read's local variables

(b)

count = read(fd, buf, nbytes)

- (a) Parameter passing in a local procedure call: the stack before the call to read
- (b) The stack while the called procedure – read(fd, buf, nbytes) - is active.

Remote procedure call

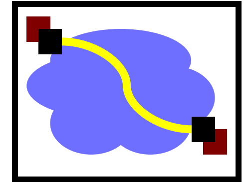

- A remote procedure call makes a call to a remote service look like a local call
  - RPC makes transparent whether server is local or remote
  - RPC allows applications to become distributed transparently
  - RPC makes architecture of remote machine transparent

…But it is not always simple

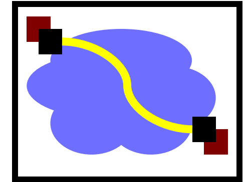

- Calling and called procedures run on different machines, with different address spaces
  - And perhaps different environments .. or operating systems ..
- Must convert to local representation of data
- Machines and network can fail

Stubs: obtaining transparency

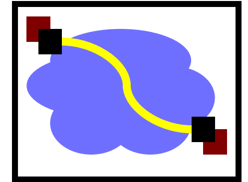

- Compiler generates from API stubs for a procedure on the client and server
- Client stub
  - **Marshals** arguments into machine-independent format
  - Sends request to server
  - Waits for response
  - **Unmarshals** result and returns to caller
- Server stub
  - **Unmarshals** arguments and builds stack frame
  - Calls procedure
  - Server stub **marshals** results and sends reply

Writing it by hand - (again…)

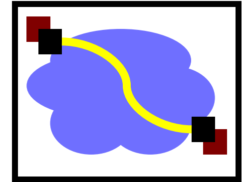

- E.g., if you had to write a, say, password cracker

```
struct foomsg {
  u_int32_t len;
}

send_foo(char *contents) {
   int msglen = sizeof(struct foomsg) + strlen(contents);
   char buf = malloc(msglen);
   struct foomsg *fm = (struct foomsg *)buf;
   fm->len = htonl(strlen(contents));
   memcpy(buf + sizeof(struct foomsg),
          contents,
          strlen(contents));
   write(outsock, buf, msglen);
}
```

Then wait for response, etc.

Marshaling and Unmarshaling

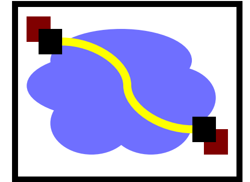

- (From example) htonl() -- “host to network-byte-order, long”.
  - network-byte-order (big-endian) standardized to deal with cross-platform variance
- Note how we arbitrarily decided to send the string by sending its length followed by L bytes of the string? That’s marshalling, too.
- Floating point...
- Nested structures? (Design question for the RPC system - do you support them?)
- Complex datastructures? (Some RPC systems let you send lists and maps as first-order objects)

Endian

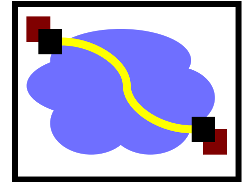

3 0	2 0	1 0	0 5
7 L	6 L	5 I	4 J

(a)

0 5	1 0	2 0	3 0
4 J	5 I	6 L	7 L

(b)

0 0	1 0	2 0	3 5
4 L	5 L	6 I	7 J

(c)

a) Original message on x86 (Little Endian)
b) The message after receipt on the SPARC (Big Endian)
c) The message after being inverted. The little numbers in boxes indicate the address of each byte

*Tanenbaum 4.2, https://en.wikipedia.org/wiki/Endianness

“stubs” and IDLs

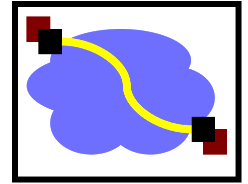

- RPC stubs do the work of marshaling and unmarshaling data
- But how do they know how to do it?
- Typically: Write a description of the function signature using an IDL -- interface definition language.
  - Lots of these. Some look like C, some look like XML, ... details don’t matter much.

Passing Value Parameters (1)

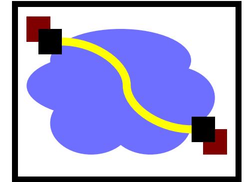

- The steps involved in a doing a remote computation through RPC.

Passing Reference Parameters

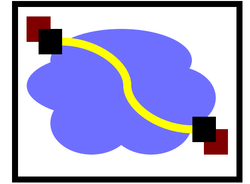

- Replace with pass by copy/restore
- Need to know size of data to copy
  - Difficult in some programming languages
- Solves the problem only partially
  - What about data structures containing pointers?
  - Access to memory in general?

Two styles of RPC implementation

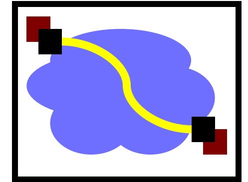

- Shallow integration. Must use lots of library calls to set things up:
  - How to format data
  - Registering which functions are available and how they are invoked.
- Deep integration.
  - Data formatting done based on type declarations
  - (Almost) all public methods of object are registered.
- Go is the latter.

Today's Lecture

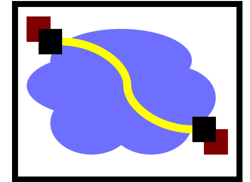

- RPC overview
- RPC challenges
- RPC other stuff

RPC vs. LPC

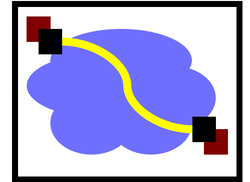

- 3 properties of distributed computing that make achieving transparency difficult:
  - Memory access
  - Partial failures
  - Latency

RPC failures

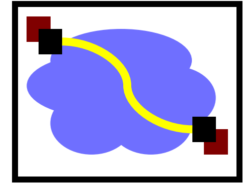

- Request from cli → srv lost
- Reply from srv → cli lost
- Server crashes after receiving request
- Client crashes after sending request

Partial failures

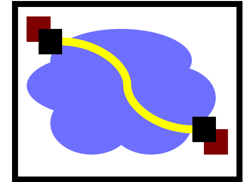

- In local computing:
  - if machine fails, application fails
- In distributed computing:
  - if a machine fails, part of application fails
  - one cannot tell the difference between a machine failure and network failure
- How to make partial failures transparent to client?

Strawman solution

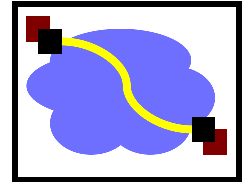

- Make remote behavior identical to local behavior:
  - Every partial failure results in complete failure
    - You abort and reboot the whole system
  - You wait patiently until system is repaired

- Problems with this solution:
  - Many catastrophic failures
  - Clients block for long periods
    - System might not be able to recover

Real solution: break transparency

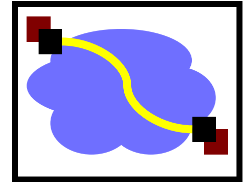

- Possible semantics for RPC:
  - Exactly-once
    - Impossible in practice
  - At least once:
    - Only for idempotent operations
  - At most once
    - Zero, don't know, or once
  - Zero or once
    - Transactional semantics

Exactly-Once?

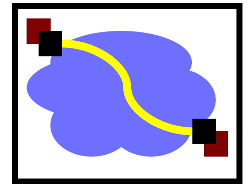

- Sorry – impossible to do in *general*.
- Imagine that message triggers an external physical thing (say, a robot fires a nerf dart at the professor)
- The robot could crash immediately before or after firing and lose its state. Don't know which one happened. Can, however, make this window very small.

Real solution: break transparency

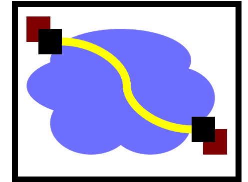

- **At-least-once**: Just keep retrying on client side until you get a response.
  - Server just processes requests as normal, doesn't remember anything. Simple! (as long as idempotent)
- **At-most-once**: Server might get same request twice...
  - Must re-send previous reply and not process request (implies: keep cache of handled requests/responses)
  - Must be able to identify requests
  - Strawman: remember all RPC IDs handled.
    - → Ugh! Requires infinite memory.
  - Real: Keep sliding window of valid RPC IDs, have client number them sequentially.

Implementation Concerns

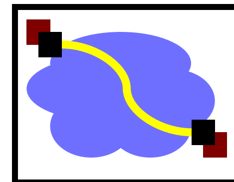

- As a general library, performance is often a big concern for RPC systems
- Major source of overhead: copies and marshaling/unmarshaling overhead
- Zero-copy tricks:
  - Representation: Send on the wire in native format and indicate that format with a bit/byte beforehand. What does this do? Think about sending uint32 between two little-endian machines
  - Scatter-gather writes (writev() and friends)

Dealing with Environmental Differences

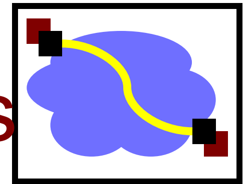

- If my function does: read(foo, ...)
- Can I make it look like it was really a local procedure call??
- Maybe!
  - Distributed filesystem...
- But what about address space?
  - This is called distributed shared memory
  - People have kind of given up on it - it turns out often better to admit that you are doing things remotely

Today's Lecture

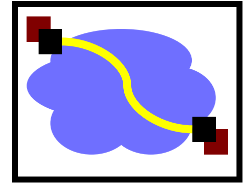

- RPC overview
- RPC challenges
- RPC other stuff

Asynchronous RPC (1)

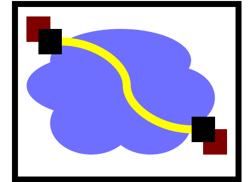

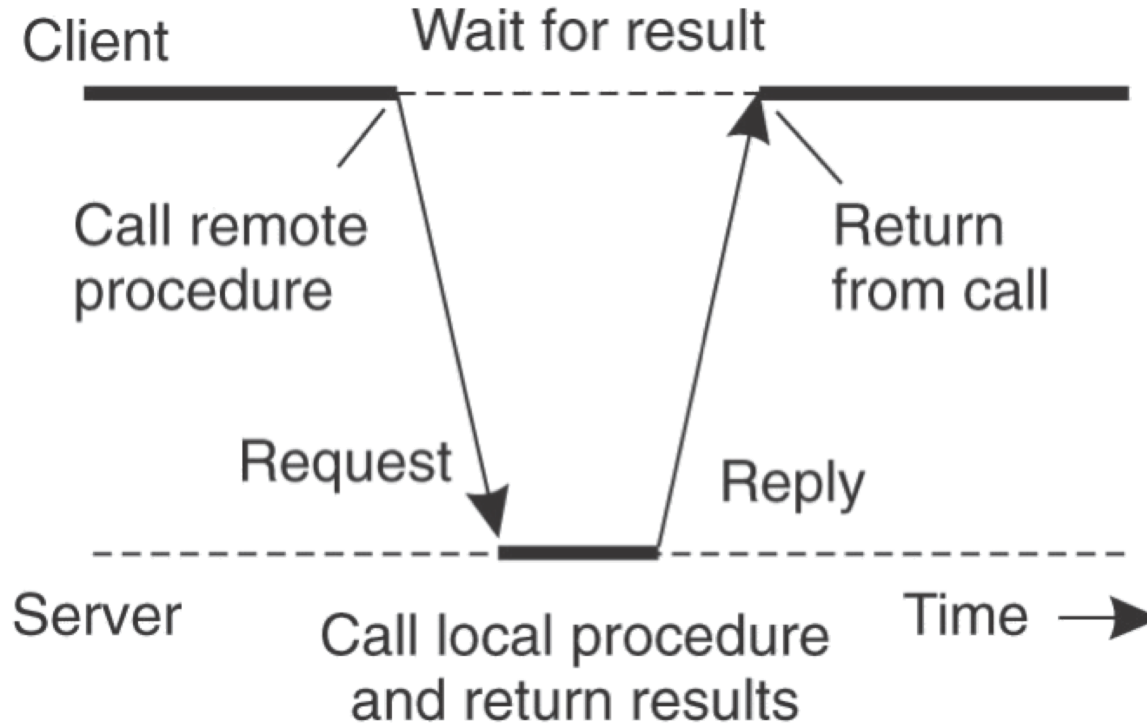

- The interaction between client and server in a traditional synchronous RPC.

Asynchronous RPC (2)

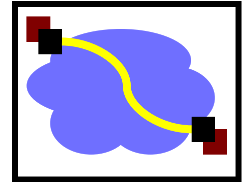

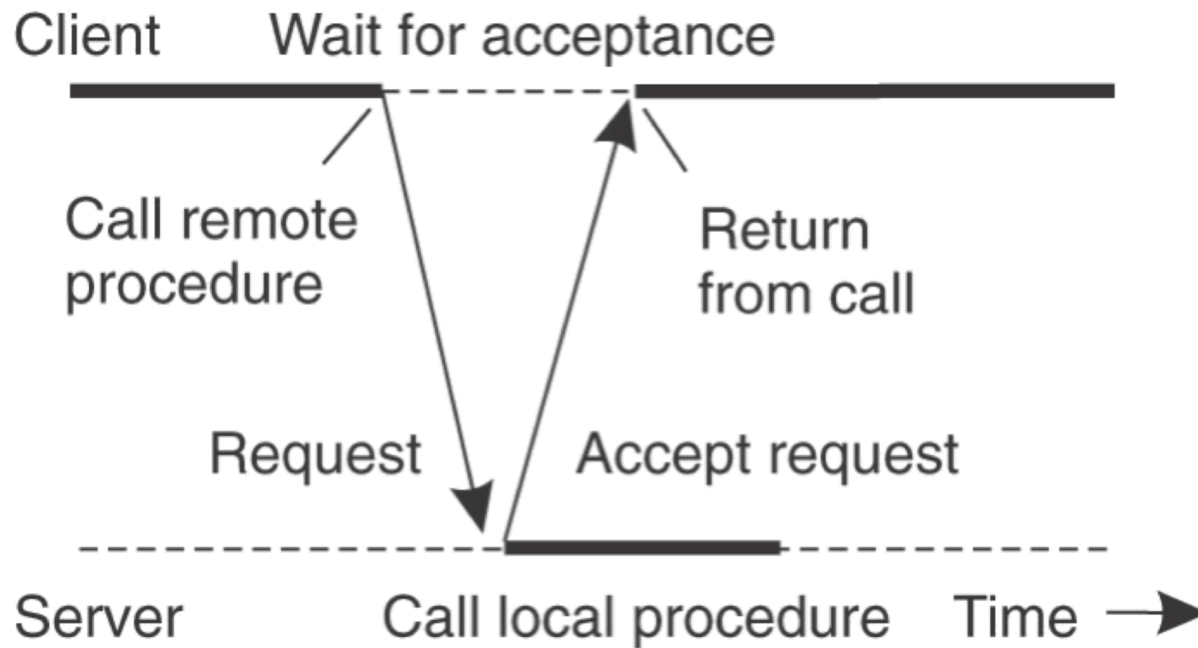

- The interaction using asynchronous RPC.

Asynchronous RPC (3)

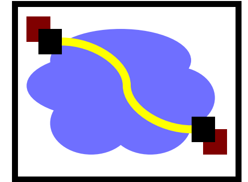

- A client and server interacting through two asynchronous RPCs.

Go Example

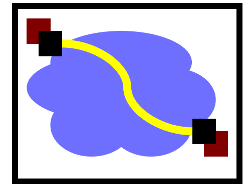

- Client first dials the server

```
client, err := rpc.DialHTTP("tcp", serverAddress + ":1234")
if err != nil { log.Fatal("dialing:", err) }
```

- Then it can make a remote 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.
```

Binding a Client to a Server

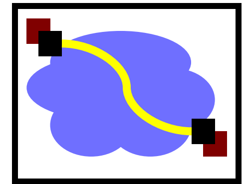

- Registration of a server makes it possible for a client to locate the server and bind to it
- Server location is done in two steps:
  - Locate the server's machine.
  - Locate the server on that machine.

Directory machine

Directory server

3. Look up server

2. Register service

Client machine

Server machine

Client

5. Do RPC

Server

1. Register end point

4. Ask for end point

DCE daemon

End point table

Server Side

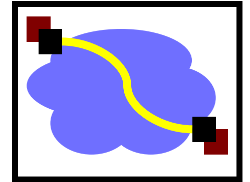

- Basic RPC code:

```
package server
type Args struct { A, B int }
type Quotient struct { Quo, Rem int }
type Arith int
func (t *Arith) Multiply(args *Args, reply *int) error {
    *reply = args.A * args.B
    return nil }
func (t *Arith) Divide(args *Args, quo *Quotient) error {
    if args.B == 0 { return errors.New("divide by zero") }
    quo.Quo = args.A / args.B
    quo.Rem = args.A % args.B
    return nil }
```

- The server then calls (for HTTP service):

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

Using RPC

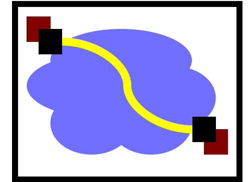

- How about a distributed bitcoin miner using RPC

- Three classes of agents:
  1. Request client. Submits cracking request to server. Waits until server responds.

  2. Worker. Initially a client. Sends join request to server. Now it should reverse role & become a server. Then it can receive requests from main server to attempt cracking over limited range.

  3. Server. Orchestrates whole thing. Maintains collection of workers. When receive request from client, split into smaller jobs over limited ranges. Farm these out to workers. When finds bitcoin, or exhausts complete range, respond to request client.

Using RPC

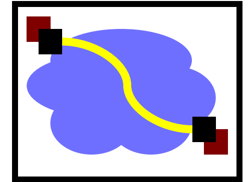

- Request→Server→Response:
  - Classic synchronous RPC
- Worker→Server.
  - Synch RPC, but no return value.
  - "I'm a worker and I'm listening for you on host XXX, port YYY."
- Server→Worker.
  - Synch RPC?
  - No that would be a bad idea. Better be Asynch.
  - Otherwise, it would have to block while worker does its work, which misses the whole point of having many workers.

Important Lessons

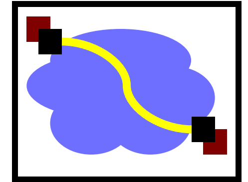

- Remote procedure calls
  - Simple way to pass control and data
  - Elegant transparent way to distribute application
  - Not the only way…
- Hard to provide true transparency
  - Failures
  - Performance
  - Memory access
  - Etc.
- Application writers have to decide how to deal with partial failures
  - Consider: E-commerce application vs. game