INTER PROCESS AND THREAD COMMUNICATION: DESIGN SPACE

Instructor: Prasun Dewan (FB 150, dewan@unc.edu)



POINTS VS. DESIGN SPACE



20



WHY IPC DESIGN SPACE?



Common mechanism and issues

Abstract model to cover all of them

Implementation of the model

SCOPE



Distributed program

\rightarrow Inter-process coupling

→ inter-process information communication

→ Communication among lightweight or heavyweight processes



UNIFYING BASIS?





INTRA PROCESS SHARED MEMORY



Write sends info, read receives it

e.g.: Shared object could be result of matrix communication

How does information consumer get notified about new information?

Thread synchronization mechanisms (semaphores, conditions, ...)

Receiving thread must wait

Threads in different processes?

INTER PROCESS SHARED MEMORY



Shared file/database/memory

If not on same computer, then need process around it

Need some other mechanism to communicate with the process

SHARED MEMORY PROS AND CONS

Familiar model

Not sufficient when threads on different computers

Global variables considered bad

Threads other than the writer and reader can access them

Receiving thread must wait for or poll for information (unless some notification mechanism added)

Inter-process notification method?

SOFTWARE INTERRUPTS (SIGNALS)



Information receiving process registers software interrupt number and handlers with OS

register (interupt#, handler)

Information sender interrupts (signals) receiving process by naming process and handler#

interrupt (process id, interrupt#)

OS calls software interrupt handler in process, interrupting its current activity, just as hardware calls hardware-interrupt handler in CPU.

New thread created?

Uses stack of current thread

OS-PROCESS COMMUNICATION



User wishes to kill process (CTRL-C)

An alarm set by process goes on

Some limit such as file size or virtual time expired

SOFTWARE INTERRUPTS PRO AND CONS

React to OS Event

Familiar interrupt model for IPC

Communicates only event id, not parameters of the event

Assumes processes share an operating system thus not suitable when processes on different computers

Process ids (usually) make no sense on other computers

Receiver cannot delay processing information (Dual of shared memory problem)

No way of queuing signals

MESSAGE PASSING

Alternative to Shared Memory and Software Interrupts



Traditionally: RPC not considered message passing

Common mechanism and issues

Abstract model to cover all of them

GIPC (Generalized IPC) based on abstract model



UNIFYING BASIS?





Message Queue Process/ Thread Processs/ Thread Message Queue Processes / threads send each other message in port queue (of varying size)

Port

(Mailbox)

Process/

Thread

message²

message¹

message³

Receiver gets message from queue

Issues?







ISSUES IN MESSAGE PASSING



RELIABLE VS. UN-RELIABLE



Reliable: every message sent to a port is receivable

Noise, disconnection can cause unreliability

When reliability needed, programmer does not need to implement it

Reliable not always needed

Spam, load averages, telepointers

Reliable is less space efficient (redundancy) and time efficient (reliability algorithm)

Networking handles this issue

IN-ORDER VS. NOT IN-ORDER DELIVERY



Order: message received in order in which they are sent

Different routes can take different times

When order needed, programmer does not need to implement it

Order not always needed

Commuting operations

In-order is less space efficient (seq #) and time efficient (in-order algorithm)

Networking handles this issue, but relevant to distributed computing as we see later

ACCESS/ROUTING



SIMPLEX BOUND PORT



DUPLEX BOUND PORT



SIMPLEX INPUT PORT



A single receiver, called server

Arbitrary number of senders, called clients

Set of simplex ports?

Messages from all clients queued together (not a collection of simplex bound ports)

Example?

Print Server

Few examples where no information comes back from server

(REPLYING) DUPLEX INPUT PORT



Like simplex equivalent except server can also send messages to clients, each with separate queue

In replying duplex input port, server can only reply back to messages, cannot initiate sends

Security, simpler API

Example of replying?

Http Server, File Server (open, close)

Example of general?

Session manager, relayer

Dual of input port?

MESSAGE PASSING: COMMON BASIS (REVIEW)



DUPLEX BOUND PORT (REVIEW)



(REPLYING) DUPLEX INPUT PORT (REVIEW)



Like simplex equivalent except server can also send messages to clients, each with separate queue

In replying duplex input port, server can only reply back to messages, cannot initiate sends

Security, simpler API

Example of replying?

Http Server, File Server (open, close)

Example of general?

Session manager, relayer

Dual of input port?

SIMPLEX OUTPUT PORT



DUPLEX OUTPUT PORT



Like simplex output port, except servers can send messages to client

Example?

SETI computation

Few examples of one client with many servers

SIMPLEX FREE PORT



Arbitrary number of senders, called clients, to a single message queue

Arbitrary number of receivers, called servers, sharing a single queue of received messages

Examples?

Load distribution among servers

Shared pool of print servers



DUPLEX FREE PORT



Like simplex free port, except servers can send messages to client

Unlike servers, clients do not share a common message queue

Example?

Web search server

General scalabale solution

Less efficient when clients on different computers

Which machine has the queue?

QUEUE LOCATION IN INPUT PORT AND MESSAGE LATENCY



DUPLEX OUTPUT PORT



DUPLEX FREE PORT



OUTPUT/FREE PORT




SESSION PORT



ACCESS/ROUTING



DISTINCTION BASED ON ACCESS?

	Xinu	
	Pipes	
9	Sockets	
	NIO	
	RMI	



XINU

send (<thread_id>, <int expression>)

int receive ()

int recvclr ()



XINU: SIMPLEX INPUT PORT



Each thread (process) has a single built in input port

A la each Java monitor having a single built in condition

Less flexible

No need for port creation operations

Thread id is the port id

Not duplex, but sender can be sent back a message on its input port



PIPES: SIMPLEX BOUND PORT?



DUPLEX FREE PORT



Some process calls pipe() to create the port

Children of pipe creator can send and receive messages

Usually one child writes and another reads

SOCKET/NIO ACCESS: MANY KINDS OF SOCKETS



SERVER SOCKET: DUPLEX INPUT PORT



Server socket is used to create regular sockets

Represent a non-stream input port handling multiple clients

Each client connects to it to create a dedicated stream duplex port

STREAM SOCKET: DUPLEX BOUND PORT?





REGULAR STREAM SOCKET: DUPLEX/FREE PORT



Regular socket represents a duplex port

Children processes share descriptors, so actually free port

DATAGRAM SERVER SOCKET: SIMPLEX INPUT PORT



No need to create ServerSocket

Each client also creates a simplex input port to receive replies

Address of client simplex port sent with each message

RMI ACCESS?



RMI: REPLYING DUPLEX INPUT PORT



Can imagine simplex, bound, output, free port semantics for RMI

ISSUES IN MESSAGE PASSING



RMI vs. Xinu?



SEND SEMANTICS?



REMOTE ASSIGNMENT (DATA COMMUNICATION)



RMI VS. XINU? (REVIEW)



SEND SEMANTICS? (REVIEW)



REMOTE ASSIGNMENT (REVIEW)



REMOTE PROCEDURE CALL (RPC)



VARIATIONS OF REMOTE ASSIGNMENT

Socket/NIO Stream

Socket/NIO Datagram

Socket Object Stream



BYTE DATA COMMUNICATION



BYTE COMMUNICATION COMMUNICATION



BLOCK COMMUNICATION



OBJECT DATA COMMUNICATION



REMOTE ASSIGNMENT VS. PROCEDURE CALL



SIMULATING REMOTE PROCEDURE CALL



Goal is to call a method with params

Send expressions encoding method and parameters

These are assigned to corresponding remote variables

Side effect of assignment is to call method with parameters

Efficiency?

Several simulating calls for one simulated call – less efficient if OS involved

SIMULATING REMOTE ASSIGNMENT



Goal is to assign some expression to a buffer variable

Define special assign method with single formal parameter

Call this method with the expression to be assigned to the buffer variable

Side effect of method call is to assign actual parameter to buffer variable

Efficiency?

More expensive RPC to do RA.

REMOTE PROCEDURE CALL VS. REMOTE ASSIGNMENT

Suitable when requests must be serviced

E.g. join a session

Can simulate remote assignment

Simulation awkward and not efficient

Remote assignment suitable when data must be sent

E.g. Next command in simulation

Remote assignment can implement RPC

However, simulation awkward and not as efficient

OPERATIONS?



LANGUAGE/COMPILER SUPPORT: TYPE CHECKING AND SPECIAL SYNTAX







REPLY?

send <port>(<expr>)

send <port>(<actual parms>)

reply (<expr>)

reply (<actual parms>)

Syntax and semantics?

Last sender is implicit port

Could reply multiple times according to this definition


ISSUES IN MESSAGE PASSING



SYNCHRONOUS VS. ASYNCHRONOUS

operation(<parms>)

read(file)

send(loadAvgPort, 1.2)

Synchronous: Operation invoker waits until the operation finishes

Asynchronous: Operation invoker does not wait until completion

Some other operation (e.g. software interrupt) needed to wait for result or completion status



SYNCHRONOUS VS. ASYNCHRONOUS VS. BLOCKING OPERATIONS

operation(<parms>)

read(file)

send(loadAvgPort, 1.2)

Blocking: Operation invoker waits, unblocking possibly before, until, or after operation completion

Synchronous is always blocking

Blocking is not always synchronous

Logical blocking times?

SYNCHRONOUS VS. ASYNCHRONOUS VS. BLOCKING OPERATIONS (REVIEW)

operation(<parms>)

read(file)

send(loadAvgPort, 1.2)

Blocking: Operation invoker waits, unblocking possibly before, until, or after operation completion

Synchronous is always blocking

Blocking is not always synchronous

Logical blocking times?

BLOCKING (LOGICAL) TIMES? A LA BINDING TIME

Binding time: When is some property bound to an entity?

E.g. when variable bound to to an address, value

Program writing, compilation, link, load, runtime

Tied to phases in the lifetime of a program

Logical times for blocking?

Phases in the lifetime of a message?

Communication pipe line?



BLOCKING TIMES PROS

Operation started	No waiting, most concurrency (without using extra thread)	Sending load average
Message in Source System Buffer	Waits until buffer available, prevents flooding	Sending tele-pointers
Message in Destination System Buffer	Sender knows message did not get lost in network	Sending load average on unreliable port
Destination thread/process starts operation	Sender knows receiving process did not fail or ignore message	Remote animation started
Destination thread/process finishes operation	Sender knows operation finished	Reservation made by airline, file received by Dropbox

LATE BLOCKING AND BUFFERING



USING SENDER BUFFER

Async NIO has direct buffer to prevent copying

Programmer can use select to determine when buffer is available



SYNCHRONOUS VS. ASYNCHRONOUS RECEIVE (SEMANTICS)

receive <port>(<var>)

receive <type> <port>(<formal
param declarations>) {<body>}

Synchronous: Operation invoker waits until the operation finishes

Synchronous receive: Receive blocks until remote assignment or procedure call finishes

Asynchronous: Operation invoker initiates operation and does not wait until completion

Asynchronous receive: Receive provides buffers to receive <expression> or <method> call

Some other operation (e.g. software interrupt) needed to wait for result or completion status



BLOCKING TIMES IN SOCKETS

socket.connect(
 new InetSocketAddress(host,port));

Block until server accepts connection to server socket

Socket socket =
 serverSocket.accept();

Block until next client tries to contact the server socket

outputStream = socket.getOutputStream(); outputStream.write(buf, offset, length);

Block until in system buffer

```
inputStream = socket.getInputStream();
```

```
int retVal =
```

```
inputStream.read(buf, offset, length);
```

Block until <= length >=1 bytes received

ASYNCHRONOUS RPC?

Asynchronous RPC semantics?

Starting remote thread

Example?

A remote animation



RENDEZVOUS IN SYNC RPC SEND, RECEIVE



SYNCHRONOUS VS. ASYNCHRONOUS RECEIVE (PROS/CONS)

Sync: no need to have separate operation to determine when operation finishes

Sometimes receiver does not need notification, shared memory model

Alpha/beta search optimization parameters, when receiver next looks at them, the value may have been set

Receiver can block on a port on which a message does not arrive



WAITING ON MULTIPLE RECEIPTS

receive <port1>(<var>)

receive <type1>
<port1>(<formal params>
{<body>}

receive <port1>(<var>)

receive <type1>
<port2>(<formal params>)
{<body>}

A single thread/process can expect messages on multiple ports

Synchronous receive \rightarrow fork a thread for each receive (inefficient)

Asynchronous receive \rightarrow a single thread can wait on multiple receives

New construct for single thread and sync receive?

ABSTRACT SYNC SELECT WITH SYNC RECEIVE

select
<pre>receive <port<sup>1></port<sup></pre>
<pre>receive <port<sup>n></port<sup></pre>
end

Select operation waits until a matching send arrives for one of the receives

It completes when the receive completes

If more than one matching send?

Pick one non deterministically



USUALLY RECEIVE ALL REQUESTS



Select is typically in a loop

Each receive is executed atomically



GUARDED RECEIVE



BOUNDED BUFFER THREAD

Bounded Buffer



THREAD/PROCESS VS. MONITOR





SHARED MEMORY/MESSAGE PASSING DUALITY



```
public synchronized void put (ElementType element)
 while (size >= MAX SIZE) {
   nonFull.condWait();
 buffer[nextIn] = element;
 nextIn = (nextIn + 1) % MAX SIZE;
  size++;
 nonEmpty.condBroadcastSignal();
public synchronized ElementType get() {
  while (size == 0) {
   nonEmpty.condWait();
  }
 ElementType retVal =
    (ElementType) buffer[nextOut];
  nextOut = (nextOut + 1) % MAX SIZE;
  size--;
 nonFull.condSignal()
 return retVal:
            put("hello");
                   Procedure
               Entry procedure
                Condition wait
              Signal and return
```

Entry procedure call

FLEXIBILITY COMPARISON

```
loop
  select
     when size < MAX_SIZE:
        receive void put(ElementType element) {
           buffer[nextIn] = element;
           nextIn = (nextIn + 1) % MAX SIZE;
           size++;
        }
     when size > 0:
        receive ElementType get() {
          ElementType retVal =
            (ElementType) buffer[nextOut];
          nextOut = (nextOut + 1) % MAX SIZE;
          size--;
          return retVal;
        }
```

```
public synchronized void put(ElementType element) {
  while (size >= MAX SIZE) {
    nonFull.condWait();
  }
  buffer[nextIn] = element;
  nextIn = (nextIn + 1) % MAX SIZE;
  size++;
  nonEmpty.condBroadcastSignal();
public synchronized ElementType get() {
  while (size == 0) {
    nonEmpty.condWait();
   }
  ElementType retVal =
    (ElementType) buffer[nextOut];
  nextOut = (nextOut + 1) % MAX SIZE;
  size--;
  nonFull.condSignal()
  return retVal;
```

Guard evaluated once at start

One signal, which is a return

Multiple waits in entry procedure

Multiple signals

Do we need the extra monitor flexibility?

Probability not



EXPLICIT SYNC VS. IMPLICIT RECEIVE (PROS, CONS)

Can be used to make sender wait when sync send

Necessary to simulate monitors

Sender has to wait

loop, select, receive needed

If other mechanisms available to make sender wait, then implicit receive better

Usually implicit receive for RPC, and explicit receive for remote assignment (data communication)

LOCATION OF COMMUNICATING THREADS?





LOCATION OF COMMUNICATING THREADS?





INTRA-PROCESS (ADDRESS SPACE) COMMUNICATION



Rationale? Intra-Process: Avoids global memory

INTRA-COMPUTER MESSAGE PASSING



Inter-Process: Cooperating processes

Is | more

INTER-COMPUTER MESSAGE PASSING



Inter-Computer → Inter-Process



LOCATION: APPLICATIONS

Intra-Process: Avoids global memory

Inter-Process: Cooperating processes

Is | more

Inter-Computer: Cooperating remote processes

File, web, ..., servers

More flexible and distant the message passing, the more complex the API and implementation

HOW TO NAME PORT?





INTRA-PROCESS (ADDRESS SPACE) COMMUNICATION



INTRA-COMPUTER MESSAGE PASSING



Similar to naming shared file

With different processes, on same OS, different resource (file) descriptors name common port These can be inherited from common parent process (like standard I/O) or they can have external name like a file name

GENERAL MESSAGE PASSING: EXTERNAL NAME





ALTERNATIVES

Can use external name for each message, specifying machine address each time

Datagram Socket

Not compatible if underlying communication mechanism requires handshake and thus connection

Access control needed on each message

External name may be bound to handle, which may result in handshake between two machines and authentication and access control at connection time

Stream Socket

An external name server keeps name to handle binding

Java RMI

CONNECTION VS. NON CONNECTION

Non connection based IPC can use any of these mechanisms

Connection based ipc uses handles, not involving external name server, representing specific connection that defines data associated with connection such as offset in stream and rights
IPC DESIGN SPACE

Shared Memory

Software Interrupts

Message Passing (Multiple Dimensions)

Reliable?	Buffer sizes?	Explicit receive?	Reply?
In-order?	Send blocking times	Receive Blocking times?	Reply blocking times?
Access?	Language support?	Select?	
RPC or RA		Location of communicating threads?	
Byte, Block or Object		Naming?	