

IPC and Threads

- 7A Inter-Process Communication
- 3T Threads

Inter-Process Communication

- the exchange of data between processes
- Goals
 - simplicity
 - convenience
 - generality
 - efficiency
 - security/privacy
 - robustness and reliability
- some of these turn out to be contradictory

IPC, Threads, Races, Critical Sections

2

OS Support For IPC

- Wide range of semantics
 - may appear to be another kind of file
 - may involve very different APIs
 - provide more powerful semantics
 - more accurately reflect complex realities
- Connection establishment mediated by the OS
 - to ensure authentication and authorization
- Data exchange mediated by the OS
 - to protect processes from one-another
 - to ensure data integrity and authenticity

Typical IPC Operations

- channel creation and destruction
- write/send/put
 - insert data into the channel
- read/receive/get
 - extract data from the channel
- channel content query
 - how much data is currently in the channel
- connection establishment and query
 - control connection of one channel end to another
 - who are end-points, what is status of connections

IPC: messages vs streams

- streams
 - a continuous stream of bytes
 - read or write few or many bytes at a time
 - write and read buffer sizes are unrelated
 - stream may contain app-specific record delimiters
- Messages (aka datagrams)
 - a sequence of distinct messages
 - each message has its own length (subject to limits)
 - message is typically read/written as a unit
 - delivery of a message is typically all-or-nothing

IPC: flow-control

- queued messages consume system resources
 - buffered in the OS until the receiver asks for them
- many things can increase required buffer space
 - fast sender, non-responsive receiver
- must be a way to limit required buffer space
 - back-pressure: block sender or refuse message
 - receiver side: drop connection or messages
 - this is usually handled by network protocols
- mechanisms to report stifle/flush to sender

IPC: reliability and robustness

- reliable delivery (e.g. TCP vs UDP)
 - networks can lose requests and responses
- a sent message may not be processed
 - receiver invalid, dead, or not responding
- When do we tell the sender "OK"?
 - queued locally? added to receivers input queue?
 - receiver has read? receiver has acknowledged?
- how persistent is system in attempting to deliver?
 - retransmission, alternate routes, back-up servers, ...
- do channel/contents survive receiver restarts?
 - can new server instance pick up where the old left off?

Simplicity: pipelines

- data flows through a series of programs
 - ls | grep | sort | mail
 - macro processor | compiler | assembler
- data is a simple byte stream
 - buffered in the operating system
 - no need for intermediate temporary files
- there are no security/privacy/trust issues
 - all under control of a single user
- error conditions
 - input: End of File output: SIGPIPE

IPC, Threads, Races, Critical Sections

8

Generality: sockets

- connections between addresses/ports
 - connect/listen/accept
 - lookup: registry, DNS, service discovery protocols
- many data options
 - reliable (TCP) or best effort data-grams (UDP)
 - streams, messages, remote procedure calls, ...
- complex flow control and error handling
 - retransmissions, timeouts, node failures
 - possibility of reconnection or fail-over
- trust/security/privacy/integrity
 - we have a whole lecture on this subject

IPC, Threads, Races, Critical Sections

9

Ludicrous Speed – Shared Memory

- shared read/write memory segments
 - *mmap(2)* into multiple address spaces
 - any process can create/map shared segments
 - perhaps locked-in physical memory
 - applications maintain circular buffers
 - data transferred w/ordinary instructions
 - OS is not involved in data transfer
 - notifications can be done w/system calls
- simplicity, ease of use ... your kidding, right?
- reliability, security ... caveat emptor!
- generality ... locals only!

IPC, Threads, Races, Critical Sections

10

IPC: synchronous and asynchronous

- synchronous operations
 - writes block until message sent/delivered/received
 - reads block until a new message is available
 - easy for programmers, but no parallelism (w/o threads)
- asynchronous operations
 - writes return when system accepts message
 - no confirmation of transmission/delivery/reception
 - requires auxiliary mechanism to learn of errors
 - reads return promptly if no message available
 - requires auxiliary mechanism to learn of new messages
 - often involves "wait for any of these" (e.g. poll/select)

a brief history of threads

- processes are very expensive
 - to create: they own resources
 - to dispatch: they have address spaces
- different processes are very distinct
 - they cannot share the same address space
 - they cannot (usually) share resources
- not all programs require strong separation
 - cooperating parallel threads of execution
 - all are trusted, part of a single program

IPC, Threads, Races, Critical Sections

12

What is a thread?

- strictly a unit of execution/scheduling
 - each thread has its own stack, PC, registers
- multiple threads can run in a process
 - they all share the same code and data space
 - they all have access to the same resources
 - this makes the cheaper to create and run
- sharing the CPU between multiple threads
 - user level threads (w/voluntary yielding)
 - scheduled system threads (w/preemption)

When to use processes

- running multiple distinct programs
- creation/destruction are rare events
- running agents with distinct privileges
- limited interactions and shared resources
- prevent interference between processes
- firewall one from failures of the other

Using Multiple Processes: cc

```
# shell script to implement the cc command
cpp $1.c | cc1 | ccopt > $1.s
as $1.s
ld /lib/crt0.o $1.o /lib/libc.so
mv a.out $1
rm $1.s $1.o
```

When to use threads

- parallel activities in a single program
- frequent creation and destruction
- all can run with same privileges
- they operate on the same (live) data
- they need to share resources
- they exchange many messages/events
- no need to protect from each other

Using Multiple Threads: telnet

```
netfd = get_telnet_connection(host);
pthread_create(&tid, NULL, writer, netfd);
reader(netfd);
pthread_join(tid, &status);
...
reader( fd ) { int cnt; char buf[100];
    while( cnt = read(0, buf, sizeof (buf) > 0 )
        write(fd, buf, cnt);
}
writer( fd ) { int cnt; char buf[100];
    while( cnt = read(fd, buf, sizeof (buf) > 0 )
        write(1, buf, cnt);
}
```

OS vs user-mode implementation

- Does OS schedule threads or processes?
- Advantages of Kernel implemented threads
 - multiple threads can truly run in parallel
 - one thread blocking does not block others
 - OS can enforce priorities and preemption
 - OS can provide atomic sleep/wakeup/signals
- Advantages of library implemented threads
 - fewer system calls
 - faster context switches
 - ability to tailor semantics to application needs

Thread state and thread stacks

- each thread has its own registers, PS, PC
- each thread must have its own stack area
- max size specified when thread is created
 - a process can contain many threads
 - they cannot all grow towards a single hole
 - thread creator must know max required stack size
 - stack space must be reclaimed when thread exits
- procedure linkage conventions are unchanged

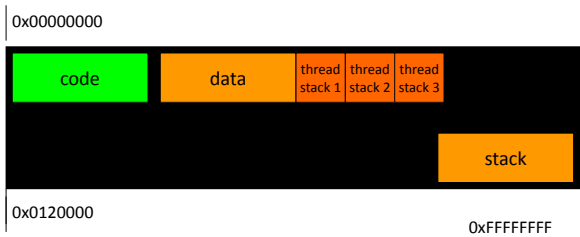
UNIX stack space management



Data segment starts at page boundary after code segment
Stack segment starts at high end of address space
Unix extends stack automatically as program needs more.

Data segment grows up; Stack segment grows down
Both grow towards the hole in the middle. They are not allowed to meet.

Thread Stack Allocation



Thread Safety - Reentrancy

- thread-safe routines must be reentrant
 - any routine can be called by multiple threads
 - concurrent or interspersed execution
 - signals can also cause reentrancy
- state cannot be saved in static variables
 - e.g. errno ... getting around C scalar returns
 - e.g. optarg ... implicit session state
- transient state can be safely allocated on stack
- persistent session state must be client-owned
 - open returns a descriptor
 - descriptor is passed to all subsequent operations

Thread Safety – Shared Data/Events

- threads operate in a single address space
 - automatic (stack) locals are thread-private
 - storage (from thread-safe malloc) can be private
 - read-only data causes no problems
 - shared read/write data can cause problems
- signals are sent to processes
 - delivered to first available thread
 - chosen thread may not have been expecting it
- a call to *exit(2)* terminates all threads

Reading and Assignments

Reading:

- Arpaci C25 ... introduction to synchronization
- Arpaci C26 ... threads and races
- Arpaci C27.3-4 ... synchronization APIs
- Arpaci C28 ... locking and implementing locks

Projects:

- start looking at project 2A (races and synchronization)

Supplementary Slides

IPC: communication fan-out

- point-to-point/unicast (1->1)
 - channel carries traffic from one sender to one receiver
- multi-cast (1->N)
 - messages are sent to specified receivers or group
- broadcast (1->N)
 - messages are sent to all receivers in a community
- publish/subscribe (N->M)
 - messages are distributed/filtered based on content
 - routing can be at sender, receiver, and in-between

IPC: in-band vs. out-of-band

- in-band messages
 - messages delivered in same order as sent
 - message n+1 won't be seen till after message n
- out-of-band messages
 - messages that leap ahead of queued traffic
 - often used to announce errors or cancel requests
 - use priority to “cut” ahead in the queue
 - priority must be honored on each link in the path
 - deliver them over a separate channel
 - a separate message channel, or perhaps a signal

half way: mail boxes, named pipes

- client/server rendezvous point
 - a name corresponds to a service
 - a server awaits client connections
 - once open, it may be as simple as a pipe
 - OS may authenticate message sender
- limited fail-over capability
 - if server dies, another can take its place
 - but what about in-progress requests?
- client/server must be on same system

IPC examples: UNIX sockets

- more powerful than pipes
 - can be bound to various protocols
 - tcp ... reliable stream, network protocol
 - udp ... unreliable datagrams, network protocol
 - unix ... named pipes
 - more versatile connection options
 - connect, listen, accept, broadcast, multicast
- both stream and message semantics
 - read/write ... synchronous stream
 - send/recv ... synchronous datagrams
- socket is destroyed when creator dies

IPC examples: mail boxes

- named message queues
 - associated with a particular receiving process
 - any process can send messages to any mailbox
- additional semantics vary with implementations
 - trusted identification of sending process
 - synchronous and asynchronous options
 - confirmation of delivery (or receipt)
 - contents of queue may survive a kill and restart
- messages typically buffered in the OS
 - some flow control is usually provided