

Administrivia

- (I guess there's not any!)

Slide 1

Semaphores — Recap

- Idea — define ADT that will be easier to use for interprocess communication/synchronization, maybe we can implement without (as much) busy-waiting.
- Definition as ADT:
 - “Value” — non-negative integer.
 - Two operations, both atomic:
 - * up (V) — add one to value.
 - * down (P) — block until value is nonzero, then subtract one.
- How does this relate to operating systems?
 - Process abstraction (and its use within the o/s) means we have to solve “synchronization problems”.
 - Solution should somehow be part of o/s.

Slide 3

Minute Essay From Last Lecture

- Alleged joke (from some random Usenet person):
 - A man's P should exceed his V else what's a sema for?
- Do you understand this?
(P is down, V is up — if not more P's than V's, no point in having a semaphore?)
- Anything else unclear?

Slide 2

Implementing Semaphores

- We want to define:
 - Data structure to represent a semaphore.
 - Functions `up` and `down`.
- `up` and `down` should work the way we said, and we'd like to do as little busy-waiting as possible.

Slide 4

Implementing Semaphores, Continued

- Idea — represent semaphore as integer plus queue of waiting processes (represented as, e.g., process IDs).
- Then how should this work ...

Slide 5

Monitors

- History — Hoare (1975) and Brinch Hansen (1975).
- Idea — combine synchronization and object-oriented paradigm.
- A monitor consists of
 - Data for a shared object (and initial values).
 - Procedures — only one at a time can run (e.g., whole procedure is a critical region).
- “Condition variable” ADT allows us to wait for specified conditions (e.g., buffer not empty):
 - Value — queue of suspended processes.
 - Operations:
 - * Wait — suspend execution (and release mutual exclusion).
 - * Signal — if there are processes suspended, allow *one* to continue. (if not, signal is “lost”).

Slide 7

Implementing Semaphores, Continued

- Variables — integer value, queue of process IDs queue.

```

down() {
    bool zero;
    enter_cr();
    zero = (value == 0);
    if (!zero)
        value -= 1;
    else
        enqueue(current_process, queue);
    leave_cr();
    if (zero)
        block(); // mark current process blocked
}

up() {
    process p = null;
    enter_cr();
    if (empty(queue))
        value += 1;
    else
        p = dequeue(queue);
    leave_cr();
    if (p != null)
        unblock(p); // mark p runnable
}

```

- `enter_cr()`, `leave_cr()` mostly like before; see p. 113.

Slide 6

Bounded Buffer Problem, Revisited

- Define a `bounded_buffer` monitor with a queue and insert and remove procedures.

- Shared variables:

```
bounded_buffer B(N);
```

Pseudocode for producers:

```
while (true) {
    item = generate();
    B.insert(item);
}
```

Pseudocode for consumers:

```
while (true) {
    B.remove(item);
    use(item);
}
```

Slide 8

Bounded-Buffer Monitor

• Data:

```

buffer B(N); // N is a constant, buffer initially empty
int count = 0;
condition full;
condition empty;

insert(item itm) {          remove(item &itm) {
  while (count == N)        while (count == 0)
    wait(full);             wait(empty);
  put(itm, B);              itm = get(B);
  count += 1;               signal(full);
  signal(empty);           }
}

```

Slide 9

Message Passing

- Previous synchronization mechanisms all involve shared variables, okay in some circumstances but not very feasible in others (e.g., multiple-processor system without shared memory).
- Idea of message passing — each process has a unique ID; two basic operations:
 - Send — specify destination ID, data to send (message).
 - Receive — specify source ID, buffer to hold received data. Usually some way to let source ID be “any”.

Slide 11

Implementing Monitors

- Requires compiler support, so more difficult to implement than (e.g.) semaphores.
- Java's methods for thread synchronization are based on monitors:
 - Data for monitor is instance variables (data for class).
 - Procedures for monitor are *synchronized* methods — mutual exclusion provided by implicit object lock.
 - `wait`, `notify`, `notifyAll` methods.
 - No condition variables, but above methods provide more or less equivalent functionality.

Slide 10

Message Passing, Continued

- Exact specifications can vary, but typical assumptions include:
 - Sending a message never blocks a process (more difficult to implement but easier to work with).
 - Receiving a message blocks a process until there is a message to receive.
 - All messages sent are eventually available to receive (can be non-trivial to implement).
 - Messages from process A to process B arrive in the order in which they were sent.

Slide 12

Implementing Message Passing

- On a machine with no physically shared memory (e.g., multicomputer), must send messages across interconnection network.
- On a machine with physically shared memory, can either copy (from address space to address space) or somehow be clever.
(Why would you want to do this? programming model is in some ways simpler, doesn't require memory shared among processes.)

Slide 13

Mutual Exclusion With Message-Passing (1)

- Idea — have "master process" (centralized control).

Pseudocode for client process: <pre> while (true) { send(master, "request"); receive(master, &msg); // assume 'token' do_cr(); send(master, "token"); do_non_cr(); } </pre>	Pseudocode for master process: <pre> bool have_token = true; queue waitQ; while (true) { receive(ANY, &msg); if (msg == "request") { if (have_token) { send(msg.sender, "token"); have_token = false; } else enqueue(sender, waitQ); } else { // assume "token" if (empty(waitQ)) have_token = true; else { p = dequeue(waitQ); send(p, "token"); } } } </pre>
---	--

Slide 15

Mutual Exclusion, Revisited

- How to solve mutual exclusion problem with message passing?
- Several approaches based on idea of a single "token"; process must "have the token" to enter its critical region.
(I.e., desired invariant is "only one token in the system, and if a process is in its critical region it has the token.")
- One such approach — a "master process" that all other processes communicate with; simple but can be a bottleneck.
- Another such approach — ring of "server processes", one for each "client process", token circulates.

Slide 14

Mutual Exclusion With Message-Passing (2)

- Idea — ring of servers, one for each client.

Pseudocode for client process: <pre> while (true) { send(my_server, "request"); receive(my_server, &msg); // assume "token" do_cr(); send(my_server, "token"); do_non_cr(); } </pre>	Pseudocode for server process: <pre> bool need_token = false; if (my_id == first) send(next_server, "token"); while (true) { receive(ANY, &msg); if (msg == "request") need_token = true; else { // assume "token" if (msg.sender == my_client) { need_token = false; send(next_server, "token"); } else if (need_token) send(my_client, "token"); else send(next_server, "token"); } } </pre>
--	--

Slide 16

Minute Essay

- Which of the synchronization mechanisms we've talked about (semaphores, monitors, message passing) do you think you would prefer to use? Why?