

Clocks

Three uses:

1. Timer reset on long-running programs; used by scheduler to switch processes.
2. Time-of-day and interval calculations for time-accounting
3. Wake-up of sleeping processes ~~and timer~~ as on alarm clock.

Possibilities

1. Use interval timer / CPU for 1,
integrate " " for 2,
simulate with " " for 3

Difficulties:

- a. must send in an initialization time.
- b. will lose ~3 mic. / day in breakage from lost ticks.
- c. Any given process may drift away from another - creating time accounting problems.
- d. Simulation of wakeup requires considerable overhead coding.

2. A synchronized time-of-day register / processor, read only, guaranteed identical in all processors. Tick every $15.625 \mu s$, just like interval timer.

Problems: Suppose clock is manually reset by guest?
Still doing wake-up simulation.

3. Wake-up clock in ~~GIC~~ GIC allows stopping simulation on a backup.
(Stopping means wake up at nearest timer reset.)

Idea: Send a time-interval + data word (processor) at any time.

When time-interval expires, send back data word and an interrupt; can store any number of requests.

Satisfactory: One time-interval at a time, no data word.
Program must handle interrupts + stacking.

Need: ~~the~~ ability to kill earlier requests to slip in new one.

VI CPU Time Keeping

1. There is an Interval Timer per processor named $CLOCK(I)$
2. Each processor maintains a cell $STOP_TIME(I)$ which is the time-of-day of the next ~~trap~~ clock trap. This cell is always updated simultaneously with the clock of the processor.

3.

3. The exact time of day is then

$$TIME_NOW = STOP_TIME(I) - CLOCK(I)$$

even if the clock trap is being inhibited.

4. For accounting purposes, the cell $START_TIME(I)$ is always the time that accounting was last done for this processor. It is updated by storing $TIME_NOW$, computed as in 3) above.

Time used is $\rightarrow TIME_USED = TIME_NOW - START_TIME(I)$

(The same technique is used in term accounting)

Alarm Clock Time-Keeping

1. Master process keeps track of all wake-up requests.

call $WAKE_UP(INTERVAL)$

puts an entry into a sorted table to wake this

process at $TIME_NOW + INTERVAL$. It blocks the process.

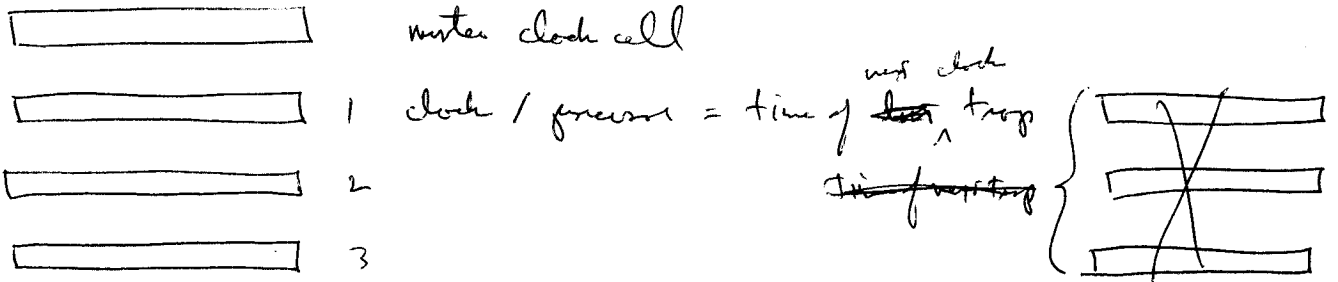
2. On every interrupt, every processor will

call $AWAKE(VU_interval, wake_interval, sleeper)$
(wake_interval)

which returns $WAKE_TIME - TIME_NOW$ of the hard sleeping
~~time~~ process in the list.

3. If wake-interval is smaller than the assigned quantum of the process about to be run, the interval timer is set for the shorter period, the difference is stored in ~~over-time(D)~~ $OVER_TIME(D)$, and the controlling process for a clock-fault is set to be the sleeper.

4. If not, $over_time$ is set to zero, and the controlling process for a clock-fault is set to be the CPU-process.



for any processor,

$$TIME_{NBW} = TIME \text{ OF NEXT TRAP} + \epsilon (\text{Interval time})$$

rule: at every trap

1. Compute $TIME_{NBW}$
2. Compare with master clock.

$TIME_{NBW} \geq \text{master clock}$

reset master clock to

$TIME_{NBW}$

$TIME_{NBW} < \text{master clock}$

reset ~~processor clock to~~ ^{Time of next trap to} master clock to

master clock + ϵ (Interval time)

19

Is Interval Timer still 24 bits?

this will overflow in 16 sec. if increment
every μsec .