

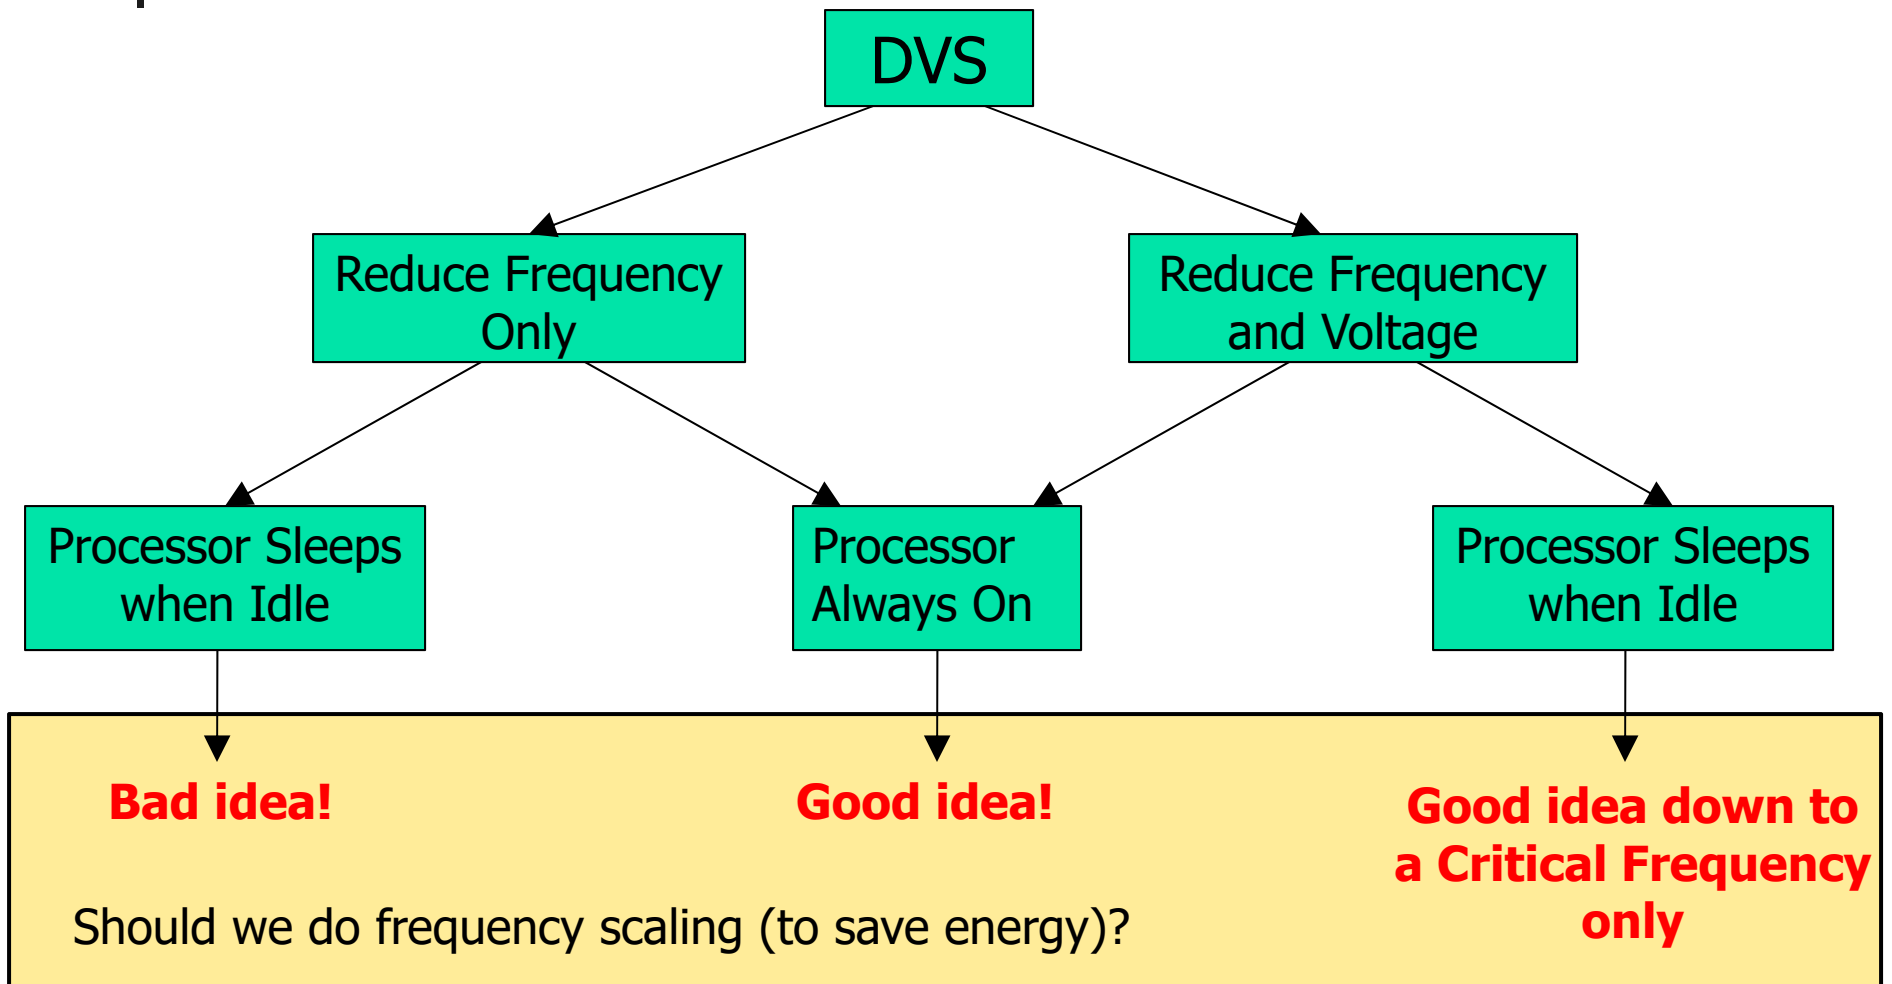


Energy

Continued

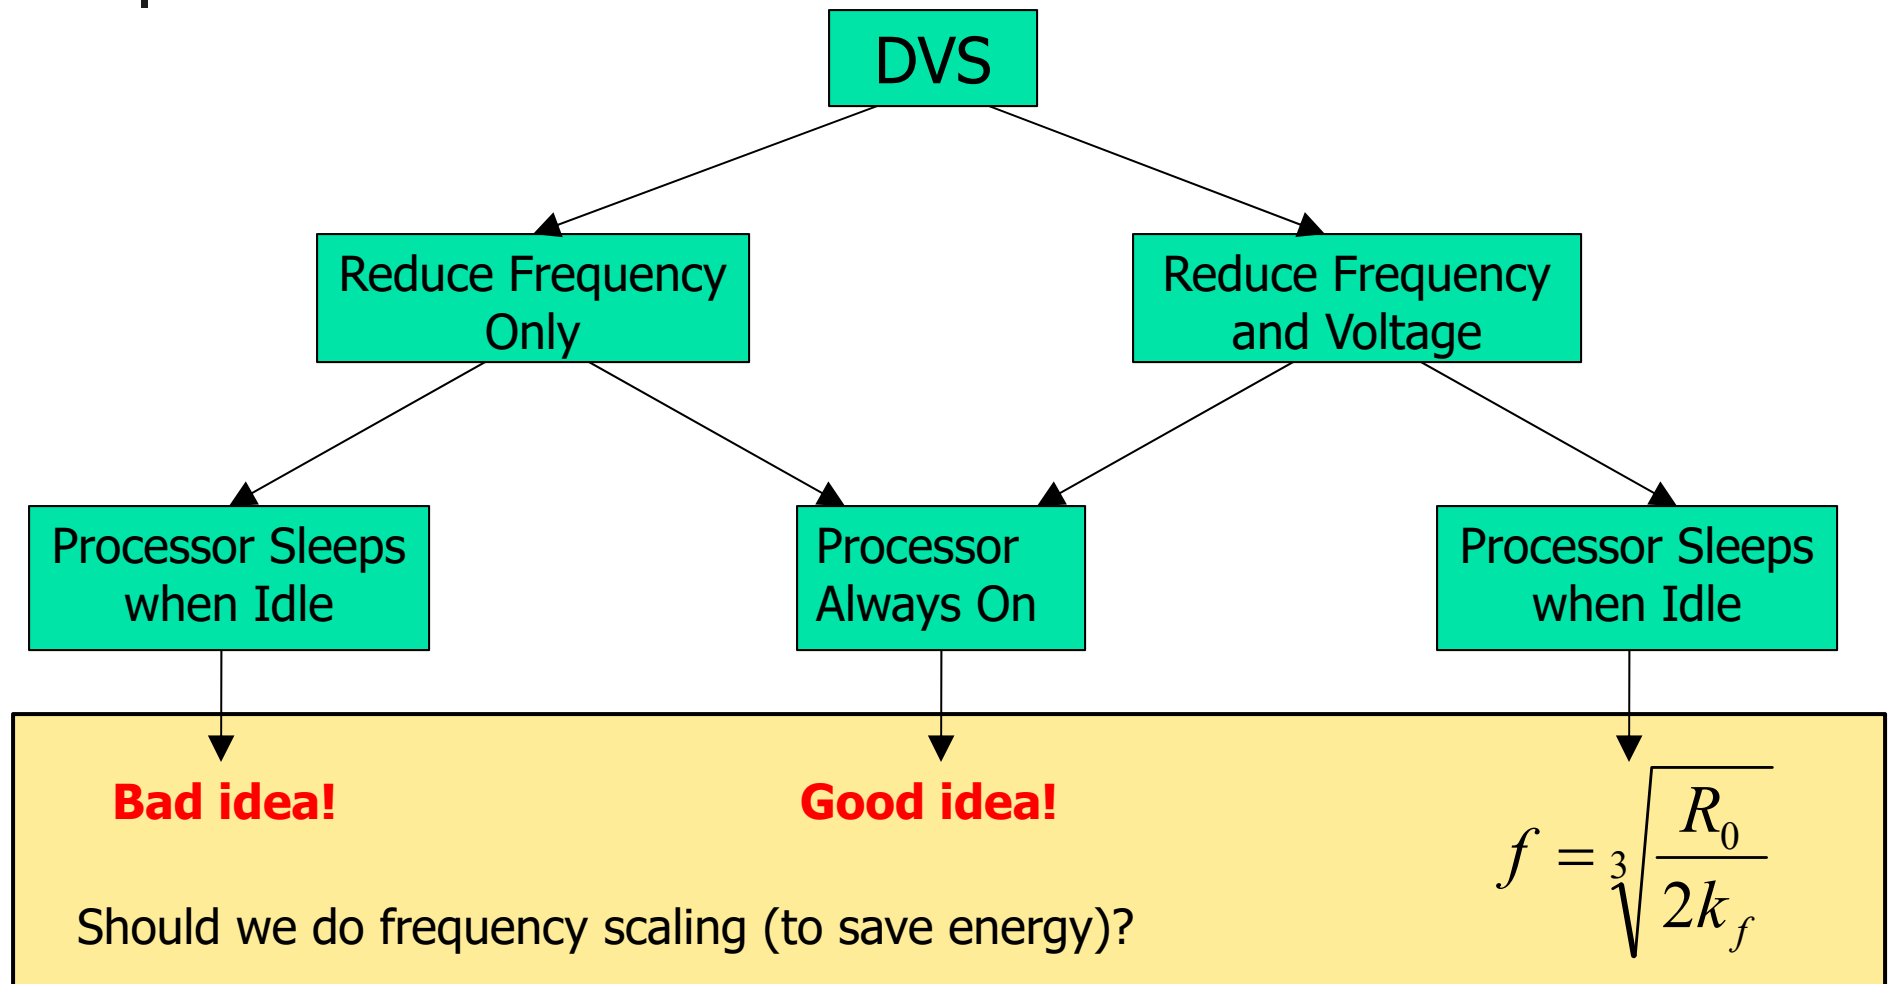


Recap





Recap





Advanced Configuration and Power Interface (ACPI)

- Defines different power saving states in a platform-independent manner
- The standard was originally developed by Intel, Microsoft, and Toshiba (in 1996), then later joined by HP, and Phoenix.
- The latest version is "Revision 6.3," published by UEFI (Jan 2019).



Global States

- **G0:** *working*
- **G1:** *Sleeping and hibernation* (several degrees available)
- **G2:**, *Soft Off*. almost the same as G3 *Mechanical Off*, except that the power supply still supplies power, at a minimum, to the power button to allow wakeup. A full reboot is required.
- **G3**, *Mechanical Off*. The computer's power has been totally removed via a mechanical switch.

Processor Performance States (P-States)



- **P0** max power and frequency
- **P1** less than P0, voltage/frequency scaled
- **P2** less than P1, voltage/frequency scaled
- ...
- **Pn** less than $P(n-1)$, voltage/frequency scaled



Processor “Sleep” States (C-states)

- **C0**: is the operating state.
- **C1** (often known as *Halt*): is a state where the processor is not executing instructions, but can return to an executing state instantaneously. All ACPI-conformant processors must support this power state.
- **C2** (often known as *Stop-Clock*): is a state where the processor maintains all software-visible state, but may take longer to wake up. This processor state is optional.
- **C3** (often known as *Sleep*) is a state where the processor does not need to keep its cache, but maintains other state. This processor state is optional.



Turning Processors Off

The Cost of Wakeup

- Energy expended on wakeup, E_{wake}
- To sleep or not to sleep?



Turning Processors Off

The Cost of Wakeup

- Energy expended on wakeup, E_{wake}
- To sleep or not to sleep?

- Not to sleep (for time t):

$$E_{no-sleep} = (k_v V^2 f + R_0) t$$

- To sleep (for time t) then wake up:

$$E_{sleep} = P_{sleep} t + E_{wake}$$



Turning Processors Off

The Cost of Wakeup

- Energy expended on wakeup, E_{wake}
- To sleep or not to sleep?

- Not to sleep (for time t):

$$E_{no-sleep} = (k_v V^2 f + R_0) t$$

- To sleep (for time t) then wake up:

$$E_{sleep} = P_{sleep} t + E_{wake}$$

- To save energy by sleeping: $E_{sleep} < E_{no-sleep}$

$$t > \frac{E_{wake}}{k_v V^2 f + R_0 - P_{sleep}}$$

Turning Processors Off

The Cost of Wakeup

- Energy expended on wakeup, E_{wake}
- To sleep or not to sleep?

- Not to sleep (for time t):

$$E_{no-sleep} = (k_v V^2 f + R_0) t$$

- To sleep (for time t) then wake up:

$$E_{sleep} = P_{sleep} t + E_{wake}$$

- To save energy by sleeping: $E_{sleep} < E_{no-sleep}$

$$t > \frac{E_{wake}}{k_v V^2 f + R_0 - P_{sleep}}$$

Minimum sleep interval



Dynamic Power Management

- DPM refers to turning devices off (or putting them in deep sleep modes)
- Device wakeup has a cost that imposes a minimum sleep interval (a breakeven time)
- DPM must maximize power savings due to sleep while maintaining schedulability

DPM and the Problem with Work-conserving Scheduling

- Example:

Task 1 (C=2, P=12)



Task 2 (C=1, P=16)



DPM and the Problem with Work-conserving Scheduling

- Example:

Task 1 (C=2, P=12)



Task 2 (C=1, P=16)



—————
Minimum sleep period

DPM and the Problem with Work-conserving Scheduling

- Example:

Task 1 (C=2, P=12)



Task 2 (C=1, P=16)



Minimum sleep period

DPM and the Problem with Work-conserving Scheduling

- No opportunity to sleep ☹️

Task 1 (C=2, P=12)



Task 2 (C=1, P=16)




Minimum sleep period

DPM and the Problem with Work-conserving Scheduling

- Must batch! 😊

Task 1 (C=2, P=12)




Task 2 (C=1, P=16)




Minimum sleep period

A Schedulability Question: How to Analyze Schedules with Sleep Periods?



Task 1 (C=2, P=12)



Task 2 (C=1, P=16)



Minimum sleep period

A Schedulability Question:

How to Analyze Schedules with Sleep Periods?

- Option 1: Treat sleep periods like a periodic task. Use the Liu and Layland utilization bound for schedulability. Problems?

Task 1 (C=2, P=12)



Task 2 (C=1, P=16)



Task 3 (C=11, P=16)



A Schedulability Question:

How to Analyze Schedules with Sleep Periods?

- Option 1: Treat sleep periods like a periodic task. Use the Liu and Layland utilization bound for schedulability. Problems?
 - Does not work because the “sleep task” cannot be preempted, whereas the rest of the tasks are preemptible. The utilization bound works only for fully preemptive scheduling.

Task 1 (C=2, P=12)



Task 2 (C=1, P=16)



Task 3 (C=11, P=16)



A Schedulability Question: How to Analyze Schedules with Sleep Periods?

- Option 2: Treat sleep periods like the *highest-priority* periodic task. Use the Liu and Layland utilization bound for schedulability. Problems?

Task 1 (C=2, P=12)



Task 2 (C=1, P=16)



Task 3 (C=11, P=16)



A Schedulability Question:

How to Analyze Schedules with Sleep Periods?

- Option 2: Treat sleep periods like the *highest-priority* periodic task. Use the Liu and Layland utilization bound for schedulability. Problems?
 - Does not work because the “sleep task” may need to have a larger period than the actual top-priority task, which contradicts rate-monotonic scheduling. The bound does not work.

Task 1 (C=2, P=12)



Task 2 (C=1, P=16)



Task 3 (C=11, P=16)



A Schedulability Question:

How to Analyze Schedules with Sleep Periods?

- Option 3: Treat sleep periods like the *highest-priority* periodic task. Use *exact response time analysis* for schedulability. Problems?

Task 3 (C=11, P=16)



Task 1 (C=2, P=12)



Task 2 (C=1, P=16)



Device Forbidden Regions

- Option 3: Treat sleep periods like the *highest-priority* periodic task. Use *exact response time analysis* for schedulability.

Problems?

- A Valid solution, but pessimistic.

(Called: Device Forbidden Regions. Published in RTAS 2008.)

Task 3 (C=11, P=16)



Task 1 (C=2, P=12)



Task 2 (C=1, P=16)

