
Rechnerarchitektur SS 2017

Exercises: Energy and Power

Jian-Jia Chen

TU Dortmund

to be discussed on 30 May 2017
revisited on 31 May 2017

Exercise - Voltage/Frequency Scaling

Suppose that we have to design a multiprocessor system to improve the energy efficiency of an application that was originally run on one processor with 1 GHz. The power consumption function when activating M processors at frequency f is

$$P(f) = 120 + M \left(10^4 \left(\frac{f}{1\text{GHz}} \right)^3 + 60 \right) \text{ mWatt.}$$

Please apply Amdahl's law.

- 1 What is the best number of processors to minimize the power consumption of the activated processors with the same performance as the original platform running at 1 GHz? What is the resulting power consumption?

7, and power consumption is ≈ 744.09 mWatt.

Exercise - Voltage/Frequency Scaling (cont.)

The power consumption function when activating M processors at frequency f is

$$P(f) = 120 + M \left(10^4 \left(\frac{f}{1\text{GHz}} \right)^3 + 60 \right) \text{mWatt}.$$

- ② Suppose that the application has 5% of its execution time that cannot be parallelized. The rest 95% of its execution time can be fully parallelized. What is the best number of processors to minimize the power consumption of the activated processors with the same performance as the original platform running at 1 GHz? What is the resulting power consumption?

To maintain the same performance with M processors, the frequency should be set to $0.05 + \frac{0.95}{M}$ GHz.

$P(0.05 + \frac{0.95}{M})$ is minimized when M is 8, and the power consumption is ≈ 984.44 mWatt.

Wrong Setting of the Frequency!!

- Original platform: 1GHz for t time units; therefore, $t \times 1$ GHz cycles.
- New platform: f GHz for t time units; 5% with one processor and 95% with M processors

$$f \times t \times 0.05 + f \times t \times 0.95 \times M = t \times 1\text{GHz}$$
$$\Rightarrow f = \left(\frac{1}{0.05 + 0.95 \times M} \right) \text{GHz}$$

Wrong Setting of the Frequency!!

- Original platform: 1GHz for t time units; therefore, $t \times 1$ GHz cycles.
- New platform: f GHz for t time units; 5% with one processor and 95% with M processors

$$f \times t \times 0.05 + f \times t \times 0.95 \times M = t \times 1\text{GHz}$$
$$\Rightarrow f = \left(\frac{1}{0.05 + 0.95 \times M} \right) \text{GHz}$$

This is wrong because only $f \times t \times 0.05$ of the execution is executed sequentially and $f \times t \times 0.95 \times M$ of the execution is executed in parallel. That is, the sequential part is only $\frac{0.05}{0.05+0.95 \times M} < 0.05$ when $M > 2$. *This does not comply with Amdal's law!!!*

Correct Setting of the Frequency!!

- Original platform: 1GHz for t time units; therefore, $t \times 1$ GHz cycles.
 - Among them, 5% is executed sequentially.
 - Among them, 95% can be executed fully in parallel on M processors.
- New platform: M processors with f GHz for t time units

$$\frac{5\% \times t \times 1 \text{ GHz}}{f} + \frac{95\% \times t \times 1 \text{ GHz}}{fM} = t$$
$$\Rightarrow f = \left(0.05 + \frac{0.95}{M} \right) \text{ GHz}$$

Exercise - Voltage/Frequency Scaling (cont.)

The power consumption function when activating M processors at frequency f is

$$P(f) = 120 + M \left(10^4 \left(\frac{f}{1\text{GHz}} \right)^3 + 60 \right) \text{mWatt.}$$

- 3 Suppose that the application has $x\%$ of its execution time that cannot be parallelized. The rest $(100-x)\%$ of its execution time can be fully parallelized. What is the trend of the best number of processors to minimize the power consumption of the activated processors with respect to x when we would like to achieve the same performance as the original platform running at 1 GHz?


Hint: Try with $x=0, 5, 10, 15, 20, 50$.

x	0	5	10	15	20	25	30	35	40	45	50
M	7	8	8	7	6	5	4	4	3	2	2
power mW (approx.)	744	984	1367	1939	2702	3620	4646	5744	6780	7861	8677

Exercise - Dynamic Power Management (DPM)

Consider a micro-controller, which consumes $P_{active} = 1.2mW$ in the ACTIVE/RUN mode or $P_{sleep} = 0.1mW$ in the SLEEP mode. Regular activations of executions occur at times $t = iT$, $i \in \{0, 1, 2, \dots\}$ to interrupt and notify the processor about the arrival of new tasks. Each task needs to run t_{task} amount of time. The transition from the SLEEP mode to the ACTIVE/RUN mode lasts t_2 ; the transition from the ACTIVE/RUN mode to the SLEEP mode lasts t_1 . We assume that the micro-controller remains idle during these transitions. Also, for simplicity, we assume that the power changes continuously and linearly during the transitions. For the energy supply of the system, an energy source (battery) with 27 kJ is deployed.

source: Thiele, at ETHZ.



Exercise - Dynamic Power Management (DPM) cont.

- ① Initially, neglect the transition times ($t_1 = t_2 = 0$). Assume that after each interrupt, the processor gets into the ACTIVE mode and executes a task for $t_{task} = 2$ sec. After the execution of the task, the processor returns into the SLEEP mode. Which condition must period T satisfy so that the system can execute at least 10^7 calculations (tasks)? How does this condition affect the maximum life time of the processor?

Solutions:

$$\begin{aligned}(2\text{sec} \times 1.2\text{mW} + (T - 2) \text{ sec} \times 0.1\text{mW}) \times 10^7 &\leq 27\text{kJ} \\ (T - 2) \times 10^3 \text{J} &\leq 3 \times 10^3 \text{J} \\ T &\leq 5.\end{aligned}$$

The smaller T the shorter the lifetime of the processor is.


Exercise - Dynamic Power Management (DPM) cont.

- ② Next, assume that $t_1 = 25$ ms, $t_2 = 75$ ms, $t_{task} = 100$ ms, and $T = t_1 + t_2 + t_{task} = 200$ ms. At time $t = 0$ the processor is in SLEEP mode. Schedule **S1**: Transition to the ACTIVE mode follows directly after an interrupt. After the task execution, i.e., after t_{task} , the processor returns immediately to the SLEEP mode. Draw the schedule and analyze the energy consumption in the first 800 ms.

Solutions: For every 200 ms:

$$\frac{0.1 + 1.2}{2} mW \times 75 msec + 1.2 mW \times 100 msec + \frac{0.1 + 1.2}{2} mW \times 25 msec = 185 \mu J.$$

For every 800 ms: $185 \mu J \times 4 = 740 \mu J.$




Exercise - Dynamic Power Management (DPM) cont.

- 3 Next, assume that $t_1 = 25$ ms, $t_2 = 75$ ms, $t_{task} = 100$ ms, and $T = t_1 + t_2 + t_{task} = 200$ ms. At time $t = 0$ the processor is in SLEEP mode. Schedule S2: If the processor is in the SLEEP mode when an interrupt occurs, then transition to the ACTIVE mode happens immediately. After the task execution, the processor decides whether to return to the SLEEP mode or to remain in the ACTIVE mode in order to execute the next task without delay when the next interrupt occurs. The processor makes this decision aiming at minimizing the energy consumption. Draw the schedule and analyze the energy consumption in the first 800 ms. Compute the energy difference, which can be saved on average per period T (200 ms) when Schedule S2 is used instead of Schedule S1.

Solutions: For every 400 ms:

$$\frac{0.1 + 1.2}{2} \text{ mW} \times 100 \text{ msec} + 1.2 \text{ mW} \times 225 \text{ msec} + 0.1 \text{ mW} \times 75 \text{ msec} \\ = 342.5 \mu\text{J}.$$

For every 200 ms: $342.5 \mu\text{J} / 2 = 171.25 \mu\text{J}$.



Exercise - Dynamic Power Management (DPM) cont.

- 4 Next, assume that $t_1 = 25$ ms, $t_2 = 75$ ms, $t_{task} = 100$ ms, and $T = t_1 + t_2 + t_{task} = 200$ ms. Can we have better schedule (in terms of energy saving) than **S2** for which the following condition must hold: the task denoted by the i -th interrupt must have finished by the time the $(i + 1)$ -th interrupt occurs?

Solutions: For every 400 ms:

$$\frac{0.1 + 1.2}{2} \text{ mW} \times 100 \text{ msec} + 1.2 \text{ mW} \times 200 \text{ msec} + 0.1 \text{ mW} \times 100 \text{ msec} \\ = 315 \mu\text{J}.$$

For every 200 ms: $315 \mu\text{J} / 2 = 157.5 \mu\text{J}$.

