EECS 151/251A Discussion 8

TA: Christopher Yarp
Mar 22, 2019

Problem 1: CPU Overclocking

As we have learned in class, you can typically decrease the propagation delay of gates by increasing V_{DD}. This can reduce the critical path delay allowing the clock to run faster. This is the general principle behind CPU overclocking.

Let us take a hypothetical CPU which nominally operates with a 3 GHz clock at a V_{DD} of 0.9 V. Let us assume that the V_{DD} required to support a particular clock frequency can be approximated as the linear function $V_{DD} = 0.0002f + 0.3$ where f is the clock frequency in MHz and V_{DD} is in V.

(a) Approximately what V_{DD} is required to run the CPU at 4 GHz?

(b) What is the ratio of active power dissipation between an overclocked CPU running at 4 GHz compared to the same CPU running at 3 GHz?

Version: 2 - 2019-03-22 11:29:48-07:00

Problem 2: Parallelism for Lower Power

You have developed an ASIC chip which transcodes user uploaded videos for a popular video sharing website. The parent company, *Digits*, wants to reduce the cost of transcoding these videos by reducing the dynamic power consumption of your chip. However, they do not want this power savings to come at the expense of fewer videos transcoded per hour. They also want the time to transcode a single video to be close to the existing solution.

Fortunately, the transcoding algorithm can be parallelized by working on separate segments of video at the same time. However, parallelizing the transcoding requires some additional logic to handle the splitting an re-assembly of the video. If the number of parallel transcoders is N and the critical path of the unparallelized chip is τ_{orig}, the overhead for parallelization is $0.1 \cdot N \cdot \tau_{orig}$ additional delay added to the critical path and $0.2 \cdot N \cdot A_{transcoder}$ additional chip area (in addition to the N replications of the encoder).

For this problem, we will assume that the clock rate for version A of this chip was set at the maximum rate possible and was 3 GHz.

For version B of the chip, the decision was made to implement 4 parallel transcoder blocks along with the required additional logic.

(a) What is the new critical path delay of version B (with version A running at the same V_{DD} as vesion A)?

(b) If version B was run at the full clock rate using all 4 of the transcoder blocks, how much time would it take to encode 1 video as compared to version A? (*Assume the number of cycles required to encode the video can be split evenly across the decoders*)

(c) Using all 4 transcoders, what clock rate is required to transcode a video in the same amount of time as version A of the chip?

(d) What is the chip size of version 2 as compared to version 1?

(e) How much power is consumed in version B of the chip vs. version A of the chip when a single video is transcoded in the same amount of time (max rate of version A)? (*Assume capacitance of chips is directly proportional to its area. Assume the activity factor is equivalent for version A and B. Also assume that the V_{DD} required to support a particular clock frequency in versions A and B of the chip can be approximated as the linear function $V_{DD} = 0.0002f + offset$ where f is the clock frequency in MHz and V_{DD} is in V. Version A and B have the same slope but different offsets.*)