Exam solutions (DAT093)

Lars Svensson

Dec 21, 2016

What follows are brief suggested answers/solutions. Other solutions than those listed here may be acceptable.

1. (a) *The abbreviation "PVT" is used to describe design margins. For each of the three letters, briefly describe why a design margin may be needed to guarantee that performance targets are met.*

   P stands for **P**rocess. Nominally identical circuits may have different performance due to manufacturing tolerances. A supply voltage that allows the slowest manufactured circuit to meet the requirements will be larger than what's needed for a circuit with nominal performance.

   V stands for **V**oltage. Even when a certain supply voltage is provided, the actual voltage at the active circuits may be lower due to resistive and inductive voltage drops in the supply wires. These effects may be compensated for by increasing the provided supply voltage somewhat.

   T stands for **T**emperature. CMOS switching speed is reduced when temperature rises (of course, the circuit's own power dissipation serves to increase temperature). A circuit designed to fulfill speed requirements also at the highest allowed temperature will be unnecessarily fast at nominal temperatures.

   (b) *The "cost" of design margins is frequently an increased power dissipation. Explain why and/or how.*

   Supply voltage is increased in order to be sure that the requirements are filled, also in the case of variations. A higher supply voltage means higher voltage swing, which means higher dynamic power dissipation.

2. *The driving trend of electronic system development over the past few decades is known as* **Moore's Law**.

   (a) *What is the common definition of the "law"?*

   Originally, Gordon Moore stated that the most economical system size for silicon integration would increase by a factor of 2 every year. In a subsequent publication, Moore suggested that the rate would slow down somewhat, to a doubling every 2 years. However, the important characteristic is that the transistor count has grown *exponentially* for *decades*; the exact value of the exponent is less important.

   (b) *Discuss how the "law" has enabled the design and manufacturing of increasingly able, high-performance systems.*

Larger numbers of devices always enable more complex designs, for example by adding a graphics unit and a video decoder to the main CPU in a smartphone. Larger numbers of devices make it possible to increase the capacity of on-chip caches, which is a huge performance benefit. The scaling associated with the "law" also increases the switching speed of individual transistors.

(c) *The previous task focused on the* benefits *of the "law". Next, select and discuss* **two challenges** *that must be addressed in order to benefit from the progression of the "law".*

The increased number of devices on a chip requires continuously smaller devices with new process nodes. Scaling (such as according to Dennard) has required repeated breakthroughs in every aspect of manufacturing, with lithography being one example. At this point, however, it is no longer clear how to scale devices further, as the smallest dimensions are now limited by atom-scale distances.

Increasing device counts also bring ever higher complexity of the designs themselves. The amount of design data to be managed grows at least as fast as the number of devices (since scaling also requires more data to be kept per device). A billion-transistor design cannot be completed with tools built to cope with a mere million-transistor design.

3. (a) *On the topic of microprocessor design, briefly describe the expected trends for clock frequency, number of cores per processor, and power budget per processor.*

Clock frequency is not expected to grow significantly in the future.

Number of cores appears to increase exponentially just as the number of transistors per chip[1].

Power budget for an entire *processor* is essentially fixed by the affordable packaging technology. (Budget *per core* is therefore falling as the number of cores rise.)

(b) *Power-directed microprocessor design will often result in other design choices than purely performance-directed design. Briefly discuss one aspect which is affected by the power perspective, and outline how the power-aware design might be different.*

Many performance-enhancing technologies involve carrying out certain activity even though it is not known that the results will be needed. Examples include branch prediction and speculation. Also, additional datapaths are increasingly difficult to keep busy, so improved performance may become more and more power-expensive as issue width grows.

4. (a) *Describe some possible benefits and drawbacks of basing a design around a Zync chip compared with using a pure FPGA solution.*

A software component facilitates re-use of the same hardware (such as an adder) for many operations (such as different additions in an algorithm). The hard processor core has better performance per power than a corresponding synthesized processor would provide. Finally, field upgrade of software is still somewhat easier than for an FPGA configuration.

[1] Google for "kilocore" ...

Drawbacks include a higher development complexity (VHDL + Java, or equivalent) and a higher price.

(b) *Describe some possible benefits and drawbacks of basing a design around a Zync chip compared with using a pure software-on-microprocessor solution.*

The configurable hardware makes it possible to design hardware accelerators which may provide improved performance compared to a pure software solution (the performance improvement may be exchanged for power reduction by reducing the supply voltage). For this benefit to be realized, some of the computations would typically have to be a "bad fit" to the fixed processor (in terms of wordlength or logic operations), or be easily parallelized across many more datapaths than those provided in the processor.

As for the previous case, the drawbacks include a higher complexity and a higher price.

(c) *What might characterize a product where the Zync solution would be superior to either of these two alternatives?*

Both sets of benefits listed above would need to be exemplified by the same product or application. Look for complex algorithms with many disparate characteristics, but that still contain identifiable computational "hot-spots" that might benefit from specialized hardware implementations.

5. (a) *Surface mounting of components is the default method for new-system development. Outline surface-mounting benefits and drawbacks, and give some examples of when it would* **not** *be a good choice.*

Benefits (compared with through-hole mounting) include small physical size, two-sided mounting, and undisturbed inner PCB layers. On the other hand, manual modification or other re-work is much more difficult. Due to mechanical stresses, surface mounting may also be a bad choice for large, heavy, or high-power (hot) components.

(b) *A Multi-Chip Module (MCM) may be an alternative to a "conventional" packaging strategy for an electronic system. Outline some benefits and drawbacks that can be expected when going for an MCM solution. In what kind of systems may MCMs be attractive?*

MCMs offer higher-density packaging and therefore potentially higher performance. Reliabilty may also be improved due to fewer electrical connections. Except in high-volume applications, the cost may be higher than for a conventionally packaged system, and component selection may be limited by bare-die availability.

High performance requirements and less cost sensitivity may for example be found in military and aerospace domains.

6. *Emilia, a designer of embedded electronic systems, has just received an emergency call: a product built by her employer includes a power supply unit which has proved to be unreliable at the power level at which it currently operates. All units fed by this supply must therefore reduce their power drain if at all possible.*

Emilia is assigned to review an FPGA design that includes a soft processor core imple-

mented in VHDL. It is not possible to replace the FPGA part (that would require a new PCB) but she does have spare room both in the FPGA itself and in the processor program memory (which is external to the FPGA). Also, the supply voltage for the FPGA can be adjusted via software that runs on the processor.

Suggest at least three possible approaches for Emilia, and discuss why they might help her reach her goal as quickly as possible!

Here are some:

- Emilia's first action should be to ascertain that the supply voltage is the lowest needed to meet the specifications. If not, a reduction might be a very quick way to a significant improvement.
- If her application has different "phases" with different performance requirements, she may be able to introduce dynamic voltage regulation to reduce average dissipation.
- She might be able to introduce application-specific instructions in the processor and thereby reduce the number of instructions executed; this in turn could allow the clock frequency and the number of memory accesses to be reduced. Each of these changes will improve dynamic dissipation.
- A lower clock frequency (as per the previous bullet) may let Emilia reduce the supply voltage even more for a further improvement of the dissipation.