Exam solutions (DAT092, DAT093)

Lars Svensson

Jan 7, 2016

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

1. (a) *A technology implementation platform for electronic systems, in a broad sense, may be considered to comprise four parts. What are these?*

   The physical implementation technology; partial hardware designs such as cell libraries and more complex hardware IP blocks; firmware/software libraries such as device drivers needed for those hardware IP blocks; and the tools to put it all together.

   (b) *Briefly discuss how requirements on manufacturing costs, time-to-market, and power dissipation may influence a decision of FPGA vs standard-cell-ASIC platform.*

   Manufacturing costs per part is typically much lower for a purpose-built ASIC than for an FPGA. Of course, the total cost of the ASIC includes the higher NRE costs incurred during design, so volumes have to be quite high to justify the choice from a pure cost perspective.

   The time-to-market for an FPGA solution is typically much shorter than for an ASIC.

   Power dissipation may be much lower for a purpose-built chip at the same performance level.

2. *Pick* any three *of the following challenging trends for future electronic system design, and suggest possible approaches to managing or handling them.*

   - *Growing complexity*
   - *Increasing performance requirements*
   - *Shrinking development time*
   - *Stricter power requirements*
   - *Reduced manufacturing cost*

   Here are some approaches mentioned during the course:

   - *Growing complexity:* Larger development teams, CAD support at higher levels of abstraction
   - *Increasing performance requirements:* Parallelism in microprocessors and other computing elements

- *Shrinking development time:* Platform design and other forms of design re-use
- *Stricter power requirements:* Specialized, highly efficient hardware blocks for different classes of computations
- *Reduced manufacturing cost:* Reconfigurability in order to meet several needs with one piece of hardware

3. (a) *In digital hardware, what component of power dissipation is targeted by clock gating? Mention at least two design concerns when clock gating is used.*

   Clock gating targets the dynamic dissipation only.

   When introducing the clock gating, you need to make sure that the transformation is "safe", that is, that the clock is turned off only when no behavioral changes would occur. This is often simple when the gating signal is derived locally, and often very difficult when an external signal is used. Also, the gate in the clock distribution path will add delay, which must be taken into account when balancing clock trees or otherwise when calculating setup- and hold-times for the logic blocks.

   (b) *Power-directed microprocessor design will often result in other design choices than mere performance-directed design. Briefly discuss at least two design aspects which are affected by the power perspective, and outline how the power-aware design might be different.*

   Branch prediction and speculation improve performance, but will sometimes involve wasteful calculations whose results are disregarded when the results of a later computation becomes known. Such calculations contribute to dissipation without contributing to the result.

   Also, cache memory size and operation can be selected for high performance and for low power. For example, speculatively fetching data into a cache may improve performance but will have a power cost.

4. (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 in order to guarantee that performance targets are met.*

   P stands for Process variations. Even if devices are designed to be identical, their actual performance may differ due to predictable or unpredictable variations in manufacturing. The margins are needed to make sure that even the slowest devices meet the speed constraints.

   V stands for Voltage variations. Supply voltage is ideally fixed and constant but in practice varies due to supply-line impedances and variable currents. The margin is needed to ensure that the chip works as advertised even at low actual voltages.

   T stands for Temperature variations. MOS devices work better at low temperatures but self-heat due to static and dynamic power dissipation, leading to slower operation and increased leakage. The margin is needed to ensure that performance targets are met at the highest end of the intended temperature range.

   (b) *Why is manufacturing yield an important issue in electronics production? Give at least two examples of what can be done to improve yield for a silicon chip.*

   The yield is defined as the percentage of successfully manufactured parts, which in

turn determines the profit margin for each part (it costs as much to build a faulty part as a working one, but you only have revenue for the latter).

Some yield-improvement techniques discussed in class are speed binning, design margining, and adaptive techniques such as RAZOR.

5. (a) *Software benchmark suites do not always usefully predict the performance of a certain hardware/software combination. Give at least two possible reasons for such insufficiency.*

   Here are some:

   - The instruction mix of the benchmark suite is too disparate from that of the target software. As an example, a suite that is heavy on floating-point computations may say little about the performance of an integer-dominated target software.
   - The application memory footprint may be significantly different from that of the target code, even if the instruction mix is identical.
   - Compilers and other tools may be optimized to give good results on well-known benchmark suites[1], in ways unlikely to benefit other codes.

   (b) *If the target hardware for a system is not yet available, it may be necessary to use a simulator to verify that the performance requirements are met. Briefly discuss some issues with this practice, such as the cost/accuracy tradeoff and how the simulation results can be used to steer higher-level design decisions.*

   In general, higher simulation accuracy will increase cost, since finer aspects of system behavior must be simulated. Simulation-time differences of two to three orders of magnitude may be encountered.

   In, for example, an incremental benchmark-driven processor development, small modifications to the processor (maybe an added instruction) are carried out, and the effect of each modification is evaluated through simulation. Only the modifications that contribute to the performance goals are kept.

6. *Modern design process models may be exemplified by "agile" and "lean" approaches. Pick* any two *of the following aspects of design and discuss how the modern approaches differ from their predecessors.*

   (a) Documentation

   (b) Customer involvement

   (c) Testing and verification

   In modern process models, documentation is less emphasized than in traditional models; customer involvement is more important and is used throughout development; and testing and verification are more tightly integrated with design and development. (There is obviously *much* more to say about this.)

[1] Compare with car mileage: a standard sequence of speeds and stops defines "mixed driving"; so manufacturers tune gearboxes etc for this case. It is very difficult to achieve comparable milage in practice.