Exam solutions (DAT091, DAT092, DAT093)

Lars Svensson

Oct 28, 2014

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

1. *Consider the following VHDL code snippet:*

```
ENTITY and_or IS
  PORT(a:STD_LOGIC;
       b:STD_LOGIC;
       c:IN STD_LOGIC;
       y1:OUT STD_LOGIC;
       y2:OUT STD_LOGIC);
END and_or;

ARCHITECTURE arch_and_or OF and_or IS
  SIGNAL x1:STD_LOGIC;
  SIGNAL x2:STD_LOGIC;
BEGIN
  x1 <= a AND b;
  y1 <= x1 OR c;
 proc1:
  PROCESS(a,b,c)
  BEGIN
    x2 <= a AND b;
    y2 <= x2 OR c;
  END PROCESS proc1;
END arch_and_or;
```

   (a) *When the process has triggered,* `a` *and* `c` *are 1 and* `b` *is 0. What are the resulting values of* `y1` *and* `y2`*?*

   `y2` is defined by the value of `c`, `OR`ed with the previous value of `x2`. As `c` is 1, `y2` also gets the value of 1, regardless of the old value of `x2`. [Note: a previously posted version was in error, in that it assumed an `AND` rather than an `OR` in the last signal assignment expression. The typo in the exam task actually made it easier to solve.]

   `y1` is not defined inside a process, and therefore any changes take effect immediately. `y1` thus takes the value 1, as given by the input values and the logic expressions.

(b) *In VHDL, what is the difference between the operators* `:=` *and* `<=`*?*
`:=` denotes variable assignment, whereas `<=` denotes signal assignment.

2. (a) *Why do you need to keep track of characteristic impedances in PCB design?*
A signal will be reflected when it reaches the end of a PCB trace, unless the load impedance connected across the output matches the characteristic impedance (C.I.) of the trace. The same phenomenon occurs when two traces with different C.I. are connected, that is, at any C.I. discontinuity along the PCB trace.

(b) *How can you control the characteristic impedances in PCB design?*
To some degree by selecting the materials; but in practical design situations, mostly by changing the width of the PCB trace, and when possible its distance from the ground plane or return wire.

(c) *List at least three important criteria to consider when choosing a packaging option for a semiconductor chip.*
Here are some:

- The thermal resistance of the package must be small enough to ensure that the chip temperature stays within specified limits.
- Small physical dimensions matter for compact systems.
- The characteristic impedances of the pins and connectors may matter when the signal wavelengths are comparable to the package dimensions.
- The cost of the package itself (and any ancillary costs for assembly, testing, etc) will often be decisive.

3. (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) *The issue of device variability is of increasing importance in integrated-circuit design. Briefly describe three possible ways to manage the effects of such variability.*

- The simplest way is to assign design margins large enough for any combination of device data to result in a chip that fulfills the specifications. The example given in the last-week lecture was of supply voltage margin to reach a given performance level, but other examples could be given too.
- The RAZOR technique, also described in the last-week lecture, relies on the ability to detect when a signal arrives too late to a sequential element, and then stop the processor pipeline and re-run the instruction at a higher supply voltage. It is an example of a family of techniques which involve a control loop to set the supply voltage just high enough to fulfill the requirements.
- Layout and logic style restrictions, in particular to more "regular" designs, may reduce the amount of variability somewhat by eliminating particularly sensitive designs.

(Also refer to suggested answers for question 5b of the 111022 exam.)

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

   - 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) *The current trends for microprocessor design include a continued growth of the number of processor cores per chip. Describe the trends for clock frequency scaling and for power dissipation* per core.

   Clock frequency is not expected to continue to quickly improve[2].

   The total power limit per package is limited by what can be economically handled by cooling. This limit seems to remain around 150W. The number of processor cores appears to increase exponentially with the transistors per chip, so dissipation per core is expected to decrease over time.

5. *The lecture on power dissipation brought up several ways to manage and/or reduce dissipation at different points in the design process. Briefly describe up to five of these. You'll get two points per separate and well-enough described approach.*

   The lecture brought up the following approaches (there are, of course, many more):

   - Choose a power-frugal implementation technology (such as ASIC over FPGA)
   - Use logic-synthesis tools to select frugal combinatorial and sequential implementations
   - Turn off the clock for logic whose results will be discarded
   - Select the lowest workable supply voltage which fulfills the requirements, and optionally dynamically tune this voltage to the speed requirements
   - Minimize the number of cache misses in processors, and forsake speculation as a method of performance improvement, since it entails throwing away some results
   - Design software to emphasize locality, again in order to reduce the number of cache misses
   - Implement systems that allow efficient modular power handling, where parts of the system may be quickly powered off without affecting other parts

[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.

[2] In fact, the rate of improvement has already slowed down considerably: the highest-clocked Intel procecssor available today appears to be the i7-4790K, at nominally 4 GHz ; this processor was introduced in Q2 2014. In February 2005, the Pentium 4 Processor Extreme Edition was already rated at 3.73 GHz—an improvement of a measly seven percent over nine years.

- Consider power requirements already at the specification level, such as to allow early decisions for power-off based on incomplete data
- Plan to use power-prediction tools from the outset

6. (a) *The lecture on real-time systems (Risat Pathan) introduced the concept of* formal proof *of system real-time properties. Despite the reassuring term, some sources of uncertainty still remain. List two of those.*

   Here are three:

   - The worst-case execution time of a task may not be deterministic, even when the task is considered in isolation.
   - Additionally, the wall-clock execution time for a task may depend on other tasks executed on the same processor in a pesudo-parallel fashion.
   - Tasks may compete for shared resources, such as memory or I/O channels, which may affect the execution time of each of them.

   (b) *Intuitively, data conversion is perfect only when it is* lossless*—when an A-to-D-to-A conversion is an identity operation. No practical implementation fulfills such a condition, and approaching the ideal gets progressively more expensive. Describe at least three design parameters which influence the error/loss in data conversion; in each case, suggest why exaggerated parameter values may be expensive.*

   - There is no reason to increase the sample rate beyond what is necessary to fulfill the Nyquist criterion, but the sample clock must also be highly stable with little jitter. A high-frequency, low-jitter clock costs power to generate.
   - A high-resolution quantization generates more signals to be switched, which will cost more power. (Other resolution costs will be investigated in DAT116.)
   - The signal of interest often needs to be separated from other signals using frequency-selective filters. Very steep filters require more components and more power.