Exam solutions (DAT093)

Lars Svensson

Oct 23, 2017

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

1. (a) *A diagram similar to the one below was shown in one of the lectures. Please mark approximate curves for an FPGA technology and an ASIC technology in the diagram. Use the "bigger is better" convention.*

   Something along these lines might be good. Red is ASIC, blue is FPGA.

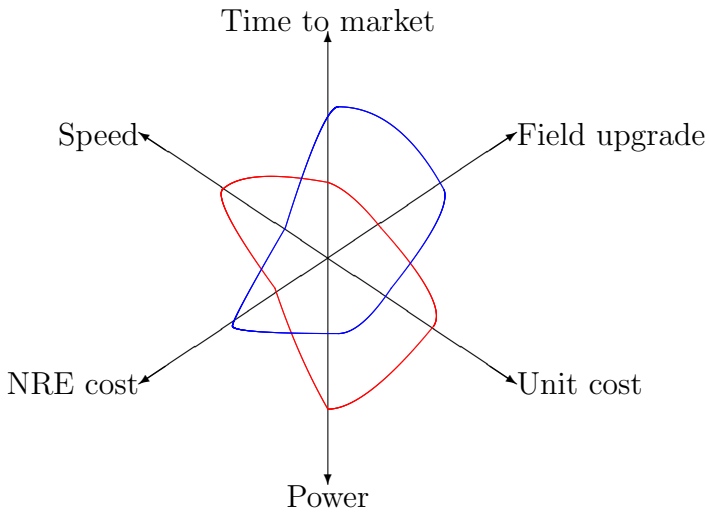

   (b) *Under what cirumstances might it be a good idea to implement a special-purpose processor in an ASIC rather than just use a processor off the shelf? Please give one cost-driven example and one performance-driven example.*

   Off-the-shelf processors are designed to perform well on a wide range of workloads. A special-purpose design could make sense for a high-performance workload with narrow characteristics. An example mentioned in the course was the multi-Tensilica processor for the Microsoft HoloLens system.

   A pure cost-driven example can be more difficult to find. A very simple, but still programmable, processor to handle some low-performance actions might make sense as a part of a larger system-on-chip. For a separate part in its own package to be cheaper than an off-the-shelf processor, it would seem to be necessary to avoid some of the resource-hungry features of standard processors (floating-point processing comes to mind) and reduce the area spent on caches; but low performance would likely make such an alternative unviable. More credibly, in a system where the mem-

ory is more expensive than the processor, it could make sense to introduce hardware and instructions to handle in-memory data compression so as to reduce memory requirements. In any case, the NRE cost for developing special-purpose tools and software would necessitate large production volumes.

2. (a) *What is the benefit of using well-established benchmark suites to evaluate processor performance? Are there also potential drawbacks? If so, specify.*

   Established suites are readily available, without large amounts of development work. Their correlation with real-world performance in the applicable application domains may be good, or at least well documented and understood. Collections of benchmark scores for existing systems are often extensive, which facilitates comparisons.

   On the other hand, vendors may have optimized their products to score highly on a well-established suite, to the detriment of the performance on slightly different workloads. Thus, established suites are best if the target software is very similar to the suite. (See the next point if not.)

   (b) *If you were to use special-purpose benchmark software to evaluate processor alternatives, what would be the important points to consider when choosing/creating the benchmark suite?*

   In order to well represent the software that will eventually run on the product, the benchmark software should be as "similar" as possible to the target software. Significant points of similarity may include instruction mix, memory footprint, and cache behavior. Additionally, cost and time of development may suggest the use of a previous version of the target software...

3. *Briefly discuss* two *(not more) important non-functional requirements, and how each of these may influence* packaging *and* PCB design.

   Here are two examples:

   Manufacturing cost is one of the most important non-functional requirements. The cost of component packages varies widely, based on performance and packing density, among other requirements. PCB cost increases with the number of copper layers.

   Physical size is important for many systems; the smallest component packages are little larger than the chip inside, and PCB technology contributes with small-pitch traces, double-sided mounting, and flexible substrates that allows fitting the PCB into awkward spaces.

4. *System power dissipation may affect overall costs in a multitude of ways, many of which have been directly and indirectly referred to in the lectures. List up to eight of these, each with a sentence that explains the link.*

   Here are some:

   - The power bill may be the most direct link, since electrical energy is not free (rule-of-thumb: SEK 1 / kWh).
   - Power-supply components (rectifiers, capacitors, connectors, wires, etc) are more expensive for higher power levels.
   - Larger-capacity batteries and power-supply transformers are heavier, therefore necessitating sturdier mechanical solutions.

- The mean time between failures (MTBF) is worse for hot equipment, and for equipment which is temperature-cycled. Failures bring warranty costs and/or intangible reputation costs.
- Component packaging gets more expensive for higher heat-transfer requirements.
- Heat removal equipment (heat sinks, heat pipes, etc) is more expensive the more heat is to be removed.
- Moving parts (such as air fans or coolant pumps) in high-capacity heat removal components wear out and fail eventually and should therefore be avoided.
- Heat generated by electronic office equipment is typically removed with air conditioning, which gets more expensive with higher capacity.

5. *. . . Your task is to match the colors with the data types.*

   These should be possible to identify even if you didn't attend this lecture!

   Black is the number of cores per processor: it remained at 1 until about 10 years ago, and has since moved to 2, 4, 8, 12 etc, but is rarely above 100.

   Red is the power: it has saturated just above 100W, since that is what can be cost-effectively removed by cooling (cf. the lecture on power).

   Yellow is the number of transistors: recent peak values are more than one billion, and its exponential increase has not yet slowed down in this diagram (which seems to end around 2015).

   Green is the clock frequency: it has flattened out at between 1 GHz and 10 GHz (your latest latop/desktop processor is probably rated between 3 GHz and 5 GHz).

   Blue therefore must be the SpecINT score (which roughly measures single-core performance; note how its rate of increase slows down when processor architects start using multicore techniques to improve performance!).

6. (a) *In discussion of design margins, what do the letters "PVT" signify?*

   Variations in **P**rocess, **V**oltage, and **T**emperature.

   (b) *How may CMOS device parameter variability increase* static *power dissipation compared to that of a hypothetical design with all-nominal parameters?*

   Assume that threshold voltages are normally distributed around the nominal values. Then, devices with low threshold voltage magnitudes will suffer from larger leakage currents, which will not be compensated for by the high-V_T specimens, as the leakage current grows exponentially with reduced threshold voltage.

   (c) *How may CMOS device parameter variability increase* dynamic *power dissipation compared to that of a hypothetical design with all-nominal parameters?*

   Threhsold voltage spread will cause a logic delay spread, as high-magnitude-threshold devices will have a lower maximum drain current. It may therefore be necessary to increase the supply voltage to meet the performance target, which will cause increased dynamic power dissipation.