

EEE4120F 2021

Quiz0 0 - Some Impressions

Please try to do this quiz (Quiz0) by Wednesday 17 Mar 23h55. This is a first quiz and does not count for marks, it is just to give you an idea of quiz structures used in this course. All the quizzes except this one are to be done as individual assignments, not as teamwork. However, this particular quiz (Quiz0) you may discuss these questions with other students before, while, or after completing the quiz (but note this is not going to be the case for subsequent quizzes).

Part 1

Q1. What is Amdahl's law, in regards to computing? Select the best option below that explains it...



- A. **It is a formula for estimating speedup of a workload, in comparison to a baseline sequential processor working on that workload**, that is achievable in response to the part of that workload being distributed among computing resources and worked on in parallel, where each parallel processing resource performs processing at a speed equivalent to the baseline.
- B. It is a formula for comparing the cost between two implementations: the baseline that works through items of a workload sequentially, to an enhanced, parallel system, that can distribute a workload among multiple processing resources, where each processor on its own would take an equivalent time to work through the workload as does the baseline.
- C. Amdahl's law is a rather simpler concept that which is described in the other options. It is just used for calculating the speed between two implementations. i.e. $S_{\text{par}} / S_{\text{seq}}$, which is usually just expressed as a number (usually to around 2 decimal places) without units.
- D. Amdahl's law is entirely irrelevant to most multi-processor computing platforms nowadays. Whether you have a fairly late-model Intel Celeron or an i7 (since it's not like the one has a clock speed double the other), your PC is going to be taking similarly long to run a similar task regardless of the processor you're using.

Q2 In computer engineering terminology, what does 'co-design' mean? Select the best option below.

A. It most often refers to situations where an individual developer is working on both the hardware and software part of the system. Coincidentally, this is typically what computer engineers do. The "co" in "co-design" (the activity) or in "co-designer" (the individual) can be assumed to mean "computer".

B. The concept of co-design has its foundations around simulation. It refers to a development process whereby each item that is done for the design of the physical product is equivalently applied to the design for the (nowadays usually virtual) simulation of that product.



C. It is generally considered an approach where the hardware design is worked on simultaneously with the software define. Typically, this refers to a development team with team members are working simultaneously on these different parts and following an agreed-upon method for coordinating and collaborating on design aspects.

D. The process of co-design, or more formally "collective design" is a formal term used to refer to process of designing an embedded system that involves a design team, i.e. multiple team members, rather than a design process that involves only one designer. The process of "co-design" thus typically leads to better thought out projects whereby more of the client's requirements are likely to be better achieved than having just one person do the design.

Q3. Why would the concept of an embedded system possibly be mixed with the notion of a high performance computer... leading to our use of the term High Performance Embedded System (HPES)? (Select the most logical response.)

A. It's just because so many microprocessors or microcontroller nowadays are essentially parallel processors (they have multiple processing codes). Writing parallel code is by definition HPC (high performance computing), so swapping if we just call it an embedded system (ES) instead of computer (C) we just get out HPES. And that's a Q.E.D. for this answer.



B. It's not really about microcontrollers being parallel. That's just one of many new technologies improvements that have led to HPES. It is actually more about High Performance Computers becoming more task-specific ... thus many of the technologies traditionally associated with HPC (e.g. threads, message passing, use of CUDA) coming into the embedded systems domain. But it is even more than that, HPES involves HPC plus the potential of designing task-specific performance-enhancing subsystems that can make the embedded system concerned a complicated mixture of parallel processing, interfacing, digital logic/SoC design among other techniques ... to deliver a pretty powerful, likely pretty small, and probably highly specialized solution.

C. This is actually quite a misnomer. While HPES and HPC both have High Performance in the name, they actually refer to different things. High Performance in terms of HPC is, as everyone knows, about achieving high processing speeds by bringing more processing resources to bear on a workload. High Performance in regards to HPES is less about the processing and more about the product (as typically they are a type of product, not a machine in some dedicated processing location); the high-performance typically is more about the systems physical attributes, such as enclosure quality (is it robust, looks professional) and robustness (e.g. can it withstand high temperatures).

D. None of these points is remotely correct. HPES and HPC is actually the same thing. Actually, more correctly, HPES can be considered a superset of HPC. i.e. any computer implementing HPC (e.g., in a dedicated data centre or in someone's office) is also an HPES.

Q4. What is Spatial Computing?

A. It is a parallel processing method whereby different tasks happen at different times.

B. It is a coding paradigm whereby you need to adhere to a particular style of spacing and indenting of your code.



C. It is a programming paradigm whereby computation is described as happening in different spaces instead of different times.

D. It is a methodology for thinking about concurrent programming in terms of what processors are responsible for what results.

Q5. What is "Xilinx Vivado" you might have used it before, you might not have. Choose the best description of it below:

A. It is an application for simulating HDL code and for outputting waveforms describing how the code would perform when run e.g. on an FPGA.



B. It is an application to develop HDL code, do simulations and program (or generate bitfiles to program) FPGAs

C. It is an application that you can use to program code in Python that will be able to run on a broad range of Xilinx-compatible processors.

D. It is an IDE for developing parallel (or sequential) programs for a range of microprocessors. It supports Python, Java and C# and provides powerful debugging facilities by which you can test GUIs developed using these languages.

Q6. A quick Verilog HDL check... Which of these code segments will do the following:

Define an input pin called reset, and a 16-bit output register called count. When there is a positive edge on reset it will set count to 0.

A.

```
input reset, count;

reg [15:0] count;

always@(posedge(reset))
    if (reset)
        begin
            count = 0x10'b0;
        end
```



B.

```
input reset;

output reg [15:0] count;

always@(posedge(reset))
    begin
        count = 16'b0;
    end
```

C.

```
input reset;

output reg [15:0] count;

always@(posedge(reset))
    count <= 16'x0; you thought this looks like X marks the spot? Tricked you! 😊
```

D.

```
output reg [16:0] count; ← too big! even though it can go over the line to the var name.

input reset;

always@(posedge(reset)) ← just to avoid giving away the sneaky trick in C.
    begin
        count <= 15'd0;
    end
```

Q7.

This is a short essay question. **Yes, in quizzes you might get the occasional essay or drawing type question** ... or pretending to be in discussion with a client as in this case.

Consider the following scenario that involves a (possibly uninformed) client asking you to help with preparing a solution, and you need to think about how you would give a **short but informed reply**. (i.e. **keep it to elevator length, not long and wordy**)

Scenario:

Your client, a bean and pea counter, wants you to develop a data logging device that has two buttons, one he can press to count a bean, and the other to press to count a pea. The device is to generate a single CSV file of the following format (that can be downloaded from the device, e.g. via USB):

Time, Bean, Pea

When he presses the bean button he says it must add a line with a 1 in the Bean column like so:

timevalue, 1, 0

and it must set its 7-seg LED display to show the running total of beans counted so far for the day.

Similarly, when he presses the pea button it likewise adds a row, but with a 1 in the Pea column:

timevalue, 0, 1

and it must show on its LED display the running total of peas counted so far for the day.

He expects not to count to over 10000 beans or peas a day (but still, that's an impressive upper limit, about 20 button presses a minute for an 8-hour day, without stopping for any rest breaks).

The client wants to know:

- How many processors are needed to build this device?
- How would the processing be split up between the processors?
- What do you think of the proposed means of recording the bean and pea presses in the csv file? Is it good? Would you prefer to suggest something more efficient?

what you're to do now: Pretend you are talking to this client and want to make a short and clear response to each question above. Take a few minutes to compose yourself and formulate an elegant and informative response to these questions.

Suggested approach...

Well, I was interested to see what you guys would say.

But here are thinking points that should have crossed your mind...

Would you say, directly what you might be feeling, as in:

“Hah you dufus. What do you want to parallelize your stupid bean counter system for? It’s like simple! My calculator could add up something that simple faster than you could count to two.”

Or would you say something more useful, ethical and to the point when you are dealing with a client. For example:

“Dear sir. Certainly, that sounds to be a job that needs much accuracy... as well as patience. However, a parallel solution to this problem appears potentially extraneous, from what I have understood of this system. Nowadays, a single core computer, even on an embedded device, is like to run at many megahertz. It would be able to add thousands of numbers, whether stored in a CSV file or in binary form, in probably less than a second, if that is fast enough for your needs. Thus, only one processor is likely needed. The processing would thus not need to be split between processors, which will simplify the design. The structure of the suggested CSV file could be optimized. While there are advantages to storing a timestamp with each entry, this could be optimized by having a relative value between time stamps instead of recording these longer absolute values. Also, unless it is necessary to record a press of a bean or pea button in separate columns, this could instead be placed into a single column as, using a number to indicate the button. Or a binary value, to record if it was ... a pea or not a pea. Hopefully, that answers the question.”

Obviously the first alternative has a whole lot of problems. But it’s a good length. The second alternative, well honestly it’s bit overdone perhaps. But it’s more suited as you obviously want to be polite, professional and clear to the client. You can see I’ve been a bit restrained on the computing terms of acronyms. But basically, of course, assume you can use technical terms etc for quizzes. But for a question like this, where one is explicitly asked to report on how you would converse with a client or someone that might not be an expert, you may want to aim for appropriate wording to boost your marks for that particular type of question. But don’t be too fussed, that sort of question doesn’t come up much.