

Assignment 1 Multiple Choice

50 points

You will get 2 points for a correct answer and 0 points if you do not answer the question. If the answer is wrong, you will get -0.66 points. Only one alternative is correct for each question.

a) When a cache is used in a microcontroller such as the EFM32GG:

1. It enables turning off the RAM and FLASH memories completely.
2. The cache must be fully-associative to enable any energy savings.
3. It is desirable to have many cache misses, because that is the scenario in which the cache can save the most power.
4. The cache can reduce the number of memory accesses, thereby reducing energy consumption and execution time.

Correct answer: Alternative 4

b) A pipelined CPU:

1. Enables the execution of multiple instructions in multiple cores to improve throughput.
2. Enables concurrent partial execution of instructions in different stages of the pipeline to improve throughput.
3. Executes computations with different operands of the same instruction in different stages of the pipeline to improve latency.
4. Splits the execution of instructions in multiple stages to reduce instruction latency.

Correct answer: Alternative 2

c) In the C programming language:

1. The extern keyword applied to a variable means that the variable is declared and defined at that point and can be accessed from other .c files.
2. The static keyword applied to global variables means these variables will be visible in all .c files.
3. The code in a .c file can only reference symbols declared in that same .c file.
4. 2-dimensional arrays use row-major layout and are implemented as pointers to the first element of the array.

Correct answer: Alternative 4

d) Dynamic power consumption in a digital circuit is due to:

1. That transistors take up area.
2. Switching activity and capacitance.
3. That transistors are not perfect switches and leak current.
4. That the transistor gate dielectric is not a perfect insulator.

Correct answer: Alternative 2

e) The main benefit of a modular kernel is:

1. Functionality can be added at boot time by loading modules.
2. Functionality can be added after reboot, so that a small kernel with basic functionality can be extended with modules.
3. Functionality can be added after re-compiling the kernel. For instance, new device drivers can be loaded in this way.
4. Functionality can be added at run-time. For instance, new device drivers can be loaded in this way.

Correct answer: Alternative 4

f) ARM Thumb 2 instruction set encoding contains:

1. 16 and 32 bit instructions.
2. 16 bit instructions only.
3. 16, 32 and 64 bit instructions.
4. 32 bit instructions only.

Correct answer: Alternative 1

g) Compared to using interrupts to handle devices, polling:

1. Enables more energy efficient access to devices since the CPU can be put in sleep mode.
2. Makes it easier to write programs that handle multiple devices concurrently since the CPU is more responsive.
3. Avoids the use of supervisor mode, resulting in fewer OS context switches and therefore lower power consumption.
4. Typically requires more CPU involvement, therefore wasting energy and computing power.

Correct answer: Alternative 4

h) Regarding bi-directional buses:

1. Both masters and slaves may act as transmitters, depending on the phase of the protocol.
2. A wire can be safely driven to both 0 and 1 at the same time by different transmitters.
3. Asynchronous bi-directional buses include a clock line to enable synchronization between endpoints in the bus.
4. An address bus is never necessary thanks to the bi-directional nature of the bus.

Correct answer: Alternative 1

Listing 1: Example Code

```
a = b + c ;  
e = a + d ;  
f = e + d ;
```

i) The C code in Listing 1 need to be compiled for a RISC processor with a load-store architecture. In this architecture, an instruction *cannot* use the same register as both a source and a destination operand. Assume that a variable is loaded into a register the first time it is used in the code and written back to memory after the last time it is used in the code. What is the minimum number of registers?

1. The code can be executed using at least two different registers.
2. The code can be executed using at least three different registers.
3. The code can be executed using at least four different registers.
4. The code can be executed using at least five different registers.

Correct answer: Alternative 2

j) When maximizing the time a system can operate without recharging the battery, the metric that we should optimize for is:

1. Real Time Performance
2. Energy Delay Product
3. Energy
4. Performance

Correct answer: Alternative 3

k) Operating system platform code is:

1. The part of the kernel code specific to a given function (e.g., scheduling, networking).
2. A part of the userspace code specific to a given CPU.
3. Not a part of the kernel code.
4. The part of the kernel code specific to each platform.

Correct answer: Alternative 4

l) In memory mapped I/O, device registers are mapped into memory addresses such that:

1. Special I/O-instructions are always used to access the devices.
2. Regular load and store instructions can be used to communicate with the devices.
3. A specific range of memory addresses is always reserved to represent pins in the device, and this range is unused when the device has no pins.
4. Interrupts are always triggered when the contents of that memory range changes.

Correct answer: Alternative 2

m) In the I2C bus:

1. A low to high transition on the SDA line while SCL is high defines the start of a transaction.
2. Every slave has a 32 bit address according to the I2C-standard.
3. A high to low transition on the SDA line while SCL is high defines the start of a transaction.
4. Only three wires are required.

Correct answer: Alternative 3

n) Several loop transformations can be used to reduce the execution time of loops. Which of the following statements about them is *true*:

1. Loop unrolling increases code size and results in more branch instructions being executed, but efficient branch prediction and better scheduling ultimately result in improved performance.
2. Loop fusion is used to merge dependent loops with different iteration spaces into a single loop, resulting in fewer branches.
3. Loop fission consists of splitting a loop into two (or more) loops in order to improve cache usage by avoiding conflict misses.
4. Loop tiling is used to remove loop nesting and improve the memory access pattern of the loop.

Correct answer: Alternative 3

- o)** When optimizing the energy use of the EFM32GG, which of the statements is *false*:
1. Analog bias currents should be turned off completely to reduce the energy consumption of the analog circuits that are in use.
 2. Unconnected GPIO pins should be disabled to avoid energy consumption caused by their Schmitt triggers acting on floating inputs.
 3. Unused RAM blocks should be turned off to save energy.
 4. Prescaling should be done as early as possible in the clock tree.

Correct answer: Alternative 1

- p)** The scheduler and filesystem code in the Linux kernel:
1. Is platform specific code initialized at boot time.
 2. Is code that needs to know specific details of the hardware it is running on.
 3. Is generalized driver code loaded as modules at boot time.
 4. Is system independent code.

Correct answer: Alternative 4

- q)** Dependencies and hazards are related concepts. Which of the following statements about them is *true*:
1. A dependency always generates a hazard in a CPU pipeline.
 2. Branch prediction can help avoid dependencies.
 3. Dependencies occur in the instruction stream, while hazards occur in the processor pipeline.
 4. RAW, WAW and WAR dependencies can turn into name hazards.

Correct answer: Alternative 3

- r)** The C tool flow has the following stages:
1. Preprocessor, compiler, prelinker, assembler, and linker, in this order.
 2. Preprocessor, compiler, assembler, and linker, in this order.
 3. Preprocessor, compiler, prelinker, preassembler, assembler, and linker, in this order.
 4. Preprocessor, preassembler, compiler, assembler, and linker, in this order.

Correct answer: Alternative 2

- s) A microcontroller has been sleeping in its deepest sleep mode and a reset cycle is required to start it up. The energy used by the microcontroller is being harvested from the environment and accumulated in an energy bank. The start-up process is too long to be carried out on a single charge of the energy bank. This means:
1. The system is designed incorrectly and the microcontroller can never be initialized.
 2. The system should have never been put into such a deep sleep mode to avoid this problem in the first place.
 3. Start-up must be split into several phases and timed sleeps must be carried out in-between.
 4. The energy bank has too low capacity and must be replaced.

Correct answer: Alternative 3

- t) Regarding the stages of the boot process of an OS such as uClinux on an embedded system:
1. There are three stages: bootloader, kernel initialization, and userspace initialization, in this order.
 2. There are two stages: kernel load and userspace initialization, in this order.
 3. Bootloader, kernel initialization and userspace initialization can be executed in any order.
 4. Kernel load and userspace initialization can be executed in any order.

Correct answer: Alternative 1

- u) Das U-Boot is a commonly used bootloader for embedded systems. Which of the statements regarding Das U-Boot is *false*:
1. It supports command line interaction through RS-232.
 2. It supports scripting of the boot process.
 3. It is easy to port to new hardware.
 4. It only supports reading the OS kernel from flash.

Correct answer: Alternative 4

- v) Flattened Device Trees (FDTs) enable:
1. Recompilation of the kernel for different platforms.
 2. Decompression of the kernel at boot time.
 3. Execution of the init process with parallel dependency resolution.
 4. The kernel to initialize correctly for the available hardware.

Correct answer: Alternative 4

- w) Which statement regarding the fork and execve functions are *false*:
1. A call to execve involves the loader.
 2. A call to execve starts a (possibly) different program.
 3. A call to fork creates a replica of the process, except for the returned value by fork (which is 0 in the parent process and the PID of the child in the child process)
 4. A call to fork creates a replica of the process, except for the returned value by fork (which is the PID of the child in the parent process and 0 in the child process)

Correct answer: Alternative 3

- x) In order to reduce the energy consumption of a microcontroller, we decide to apply techniques that reduce the power consumed by its caches. To keep good performance, the content of the caches must be kept alive at all times. In that case, we should:
1. Apply clock gating to the caches.
 2. Apply power gating to the caches.
 3. Apply both clock gating and power gating to the caches.
 4. Not apply clock gating nor power gating to the caches.

Correct answer: Alternative 1

- y) Consider the conditional IT instruction and all its variants (ITT, ITE...) compared to regular branch instructions such as BNE and others, and choose the correct answer:
1. The IT instructions are especially useful for long sequences of instructions executed conditionally.
 2. The IT instructions are more useful for short sequences of instructions executed conditionally.
 3. Regular branch instructions are always more efficient regarding performance, while the IT instructions are more energy friendly.
 4. Performance is always the same regardless of what type of instruction is used.

Correct answer: Alternative 2