(1) [3] What is one of the simplest (and oldest) ways to exploit parallelism among instructions?

(2) [3] Clock rates have grown by a factor of 1000 while power consumed has only grown by a factor of 30. How was this accomplished? Why was this so effective?

(3) [3] Which metric should the architect worry about, power or energy? Why?

(4) [4] As minimum feature sizes decrease, a) what happens to power density? b) what happens to wire resistance, and why?

(5) [4] What is "Yield"? Is it important? Why or why not?

(6) [4] It is difficult for the internal processing elements on a chip to communicate with the outside world (things that are on other chips, for example). Explain why that is.

(7) [6] There are two different types of power consumption, static and dynamic. Which one contributed the most to overall power consumption 20 years ago? What is the situation now, and why has it changed?

(8) [4] What is a benchmark program? What is the **perfect** benchmark?

(9) [4] Do benchmarks need to change over time? Why or why not?

(10) [4] Machines today use registers - often as many as they can. Give 2 advantages and 2 disadvantages to using registers.

(11) [4] Register allocation is a very effective and useful compiler optimization. However, it cannot be used in certain situations for correctness reasons - explain why.

(12) [4] What are the 3 pipeline hazards? Which one can be solved by throwing money at it?

(13) [3] What is an "unbalanced" pipeline?

(14) [4] What is forwarding, and why is it useful?

(15) [3] The Intel 80x86 has a CISC instruction set, and yet performs as well as RISC machines. How is that accomplished?

(16) [6] In an MOS device, there is a gate, drain, and source. Briefly explain how this device works. Drawing a picture is recommended.

(17) [4] What type of addressing mode does a control flow changing instruction use? Why?

(18) [4] What are the two primary goals of a compiler (in order)?

(19) [4] There are a variety of addressing modes that can be used. Why is Displacement mode a good one to choose?

(20) [4] In the 5-stage pipeline described in class, is it possible to have a WAW hazard? Why or why not?

(21) [5] Processor A requires 600 instructions to execute a given program, uses 1 cycles per instruction, and has a cycle time of 5 ns. Processor B requires 3 cycle per instruction, but only requires 200 instructions to do the same program. What must the cycle time of Processor B be in order to give the same CPU time as Processor A? (Show your work)

(22) [6] You are responsible for designing a new processor, and for a variety of reasons you must use a fixed 12 bit instruction size. You would like to support 16 different operations, use a 3-operand instruction format, and have 8 registers. If it is possible to do this, draw what an instruction would look like. If it is not possible, explain why, and show what you would do to fix

(23) [5] An important program spends 60% of its time doing Floating Point operations, and 40% of its time doing integer arithmetic. By redesigning the hardware you can either make the Floating Point unit 50% faster (take 50% as long), or the integer unit 80% faster (take 20% as long). Which should you do, and why?

(24) [5] The standard MIPS has a 5-stage pipeline, and uses a load delay slot. If the machine is redesigned to be a 7-stage pipeline, with the following stages:

F D RR E1 M1 M2 WB (where RR stands for Register Read)

Assuming the machine has bypass/forwarding logic, how many load delay slots will this new design require if the memory returns the value at the end of M2? At the end of M1?