0907731 Advanced Computer Architecture (Fall 2009)
Midterm Exam

الاسم: ..... رقم التسجيل: ..... رقم الشعبة: 1

==========

Instructions: Time **75** min. Closed books & notes. No calculators or mobile phones. Show your work clearly and limit your answer to the space provided. Every question has 5 points.

==========

Q1. List 5 reasons why, over time, latency improvements lag bandwidth improvements in the technologies of memory, disk, processor, and network.

1) Moore's Law helps BW more than latency
2) Distance limits latency
3) Bandwidth easier to sell ("bigger=better")
4) Latency helps BW, but not vice versa
5) Bandwidth hurts latency
6) Operating System overhead hurts Latency more than Bandwidth

Q2. Refer to the dynamic and static power consumption formulae of microprocessors. What are the major techniques to reduce these consumptions?

$$Power_{dynamic} = 1/2 \times CapacitiveLoad \times Voltage^2 \times FrequencySwitched$$

Dynamic power can be lowered by:
1) Fewer transistors
2) Lower power supply voltage
3) Lower frequency
4) Turning off clock of inactive modules

$$Power_{static} = Current_{static} \times Voltage$$

Static power can be lowered by:
1) Fewer transistors
2) Larger transistors
3) Lower power supply voltage
4) Turning the power supply off inactive modules

Q3. You have two ISA design alternatives: (i) include 3 basic addressing modes, and (ii) include 7 basic and advanced addressing modes. Discuss the advantages and disadvantages of each alternative.

Including only 3 basic addressing modes has the following advantages:

1) Keeps the processor simple and "RISC"

2) Gives short instructions

Including several addressing modes has the following advantages:

1) Gives more flexibility to the programmer

2) Needs fewer instructions to write a program

Q4. Assume that you have a typical 5-stage pipelined processor that uses forwarding and stalls to solve data hazards, resolves branches in the decode stage, and has one branch delay slot. After you do pipeline scheduling, loop unrolling (total two replicas), and code modification to remove inter-loop dependencies for the following code sequence, the minimum number of cycles needed to execute this code sequence for the first iteration is: _____ (*don't count execution cycles overlapped with the second iteration, and show your work clearly*)

```
L1:   lw      R2, 0(R1)
      lw      R3, 1000(R1)
      add     R3, R3, R2
      addi    R1, R1, #-4
      bne     R1, R5, L1
      nop
```

```
                          1 2 3 4 5 6 7 8 9 0 1 2 3 4
L1:   lw      R2, 0(R1)   F D E M W
      lw      R3, 1000(R1)  F D E M W
      lw      R6, -4(R1)      F D E M W
      add     R3, R3, R2        F D E M W
      lw      R7, 996(R1)         F D E M W
      addi    R1, R1, #-8           F D E M W
      add     R7, R7, R6              F D E M W
      bne     R1, R5, L1                F D E M W
      add     R3, R3, R7                  F D E M W
L1:   lw      R2, 0(R1)                     F D E M W
```

9 Cycles

Q5. In a speculative superscalar processor, describe what the processor needs to do when committing a miss predicted branch instruction.

1) Flush all later instructions in the ROB
2) Flush all instructions in the fetch, issue, and execution queues
3) Restart fetching instructions at the correct branch direction

Q6. The Pentium III and Pentium 4 (NetBurst) processors use different implementations of register renaming. Outline the main differences between the two implementations.

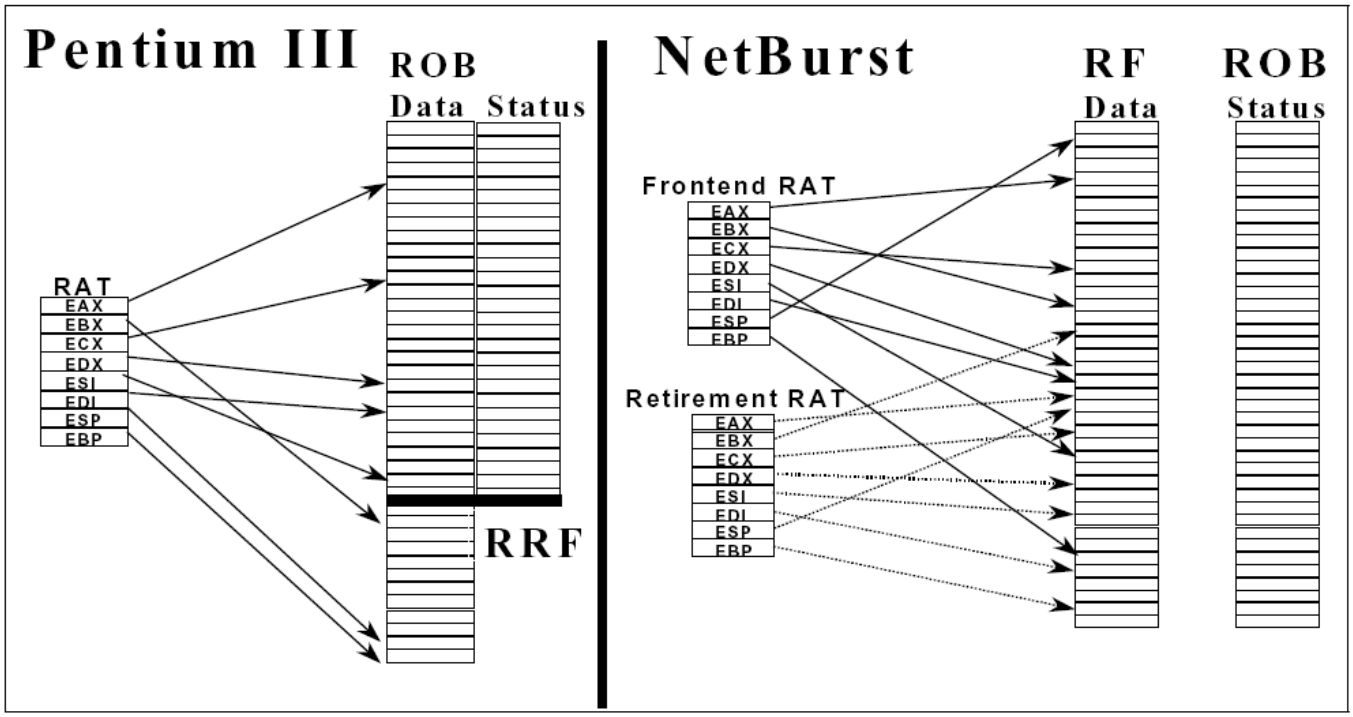

The Pentium III (P6) microarchitecture **allocates the data result registers and the ROB entries as a single, wide entity with a data and a status field**. The ROB data field is used to store the data result value of the uop, and the ROB status field is used to track the status of the uop as it is executing in the machine. These ROB entries are allocated and deallocated sequentially and are pointed to by a sequence number that indicates the relative age of these entries. Upon retirement, the result data is physically **copied from the ROB data result field into the separate Retirement Register File (RRF)**. The RAT points to the current version of each of the architectural registers such as EAX. This current register could be in the ROB or in the RRF.

The NetBurst microarchitecture **allocates the ROB entries and the result data Register File (RF) entries separately**. The ROB entries, which track uop status, consist only of the status field and are allocated and deallocated sequentially. A sequence number assigned to each uop indicates its relative age. The sequence number points to the uop's entry in the ROB array, which is similar to the P6 microarchitecture. The Register File entry is allocated from a list of available registers in the 128-entry RF–not sequentially like the ROB entries. Upon retirement, **no result data values are actually moved from one physical structure to another**.

<Good Luck>