

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
DEPARTMENT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCE

6.375 Complex Digital Systems
Spring 2006 - Quiz - March 24, 2006
80 Minutes

NAME: _____ SCORE: _____

Please write your name on every page of the quiz.

Not all questions are of equal difficulty, so look over the entire quiz and budget your time carefully.

Please carefully state any assumptions you make.

Enter your answers in the spaces provided below. If you need extra room for an answer or for scratch work, you may use the back of each page but please *clearly indicate where your answer is located*.

A list of useful equations is printed at the end of this quiz. You can detach this sheet for reference and do not have to hand this in. *We will not grade anything written on the equation sheet.*


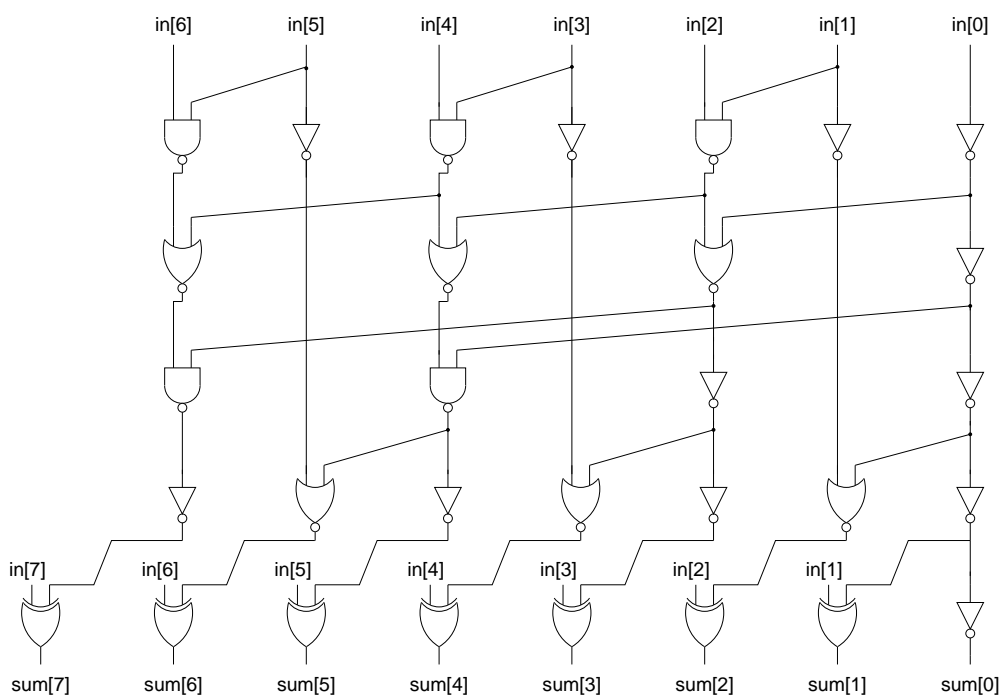
You will also receive a separate handout containing a copy of the relevant Bluespec lecture slides and code. *We will not grade anything written on the Bluespec slides.*

You must not discuss the quiz's contents with other students who have not yet taken the quiz. If, prior to taking it, you are inadvertently exposed to material in a quiz — by whatever means — you must immediately inform the instructor or a TA.

	Points	Score
Problem 1	25	
Problem 2	25	
Problem 3	25	
Problem 4	25	

Problem 1 : Logical Effort for Incrementer Carry Chain (25 total points)

The following diagram illustrates two different incrementer architectures. For all parts of this question you should assume that the delay unit (τ) for this process is 10 ps and that the parasitic delay of a minimum-sized inverter (P_{inv}) is 1.

**Ripple-Carry Architecture****Parallel-Prefix Architecture****Part 1.A : Critical paths for the adder architectures (5 points)**

Draw a line through the critical path for both the ripple-carry and the parallel-prefix architectures. When determining the critical path you can assume that XOR gates are slower than NAND/NOR gates which are slower than inverters.

Part 1.B : Optimal delay of the adder architectures (10 points)


Use logical effort to calculate the optimal delay of the critical path for both architectures in picoseconds. *You should ignore all gates which are not on the critical path!* Do not use branching effort. Ignore the fact that some gates have a fanout greater than one. The desired input capacitance of the isolated carry chain is 6 fF (since we are ignoring gates which are not on the critical path this is the input capacitance for a single gate). The load capacitance of every sum output is 60 fF. Show all your work.

Part 1.C : Gate sizing for the adder architectures (10 points)

Identify the optimum gate sizes for each gate in the critical path for both architectures. The gate sizes should be in femtofarads of input capacitance.

Problem 2 : RC Modeling of Register File Write Bitline (25 total points)

In this problem we will be revisiting the register file write bitline you analyzed in Lab 2. Remember that the write bitline must drive the D input port of 32 flip-flops. The combined gate capacitance of these flip-flops can be a significant load on the write bitline. The load on the write bitline is further increased by wire capacitance, since flip-flops are usually large and thus often spread apart. The following figure illustrates the write bitline including a reasonable final stage of the bitline driver. For this problem we will only consider this final stage even though the real driver might include many stages. As you determined in the lab assignment, each bitcell is $15.68\ \mu\text{m}$ wide and the input capacitance of the bitcell's D port is $3\ \text{fF}$. The following figure illustrates the register file write bitline. The bitline is routed on Metal 2. You can ignore any via resistance or capacitance. Remember that the driver PMOS/NMOS sizes are in units of minimum NMOS transistor width ($0.36\ \mu\text{m}$). For example, the NMOS for the last stage of the bitline driver is $0.36\ \mu\text{m} \times 32 = 11.52\ \mu\text{m}$.



The following table lists various parameters for a $0.18\ \mu\text{m}$ technology which you may find useful when solving this problem. Remember that there is a list of equations at the end of this quiz.

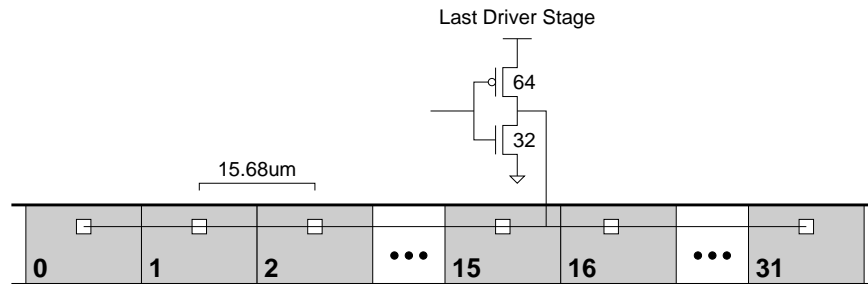
Transistor Process Parameters	Value
Desired ratio of PMOS/NMOS widths	2
PMOS gate capacitance per μm of transistor width	$1.5\ \text{fF}/\mu\text{m}$
NMOS gate capacitance per μm of transistor width	$1.5\ \text{fF}/\mu\text{m}$
PMOS drain capacitance per μm of transistor width	$0.3\ \text{fF}/\mu\text{m}$
NMOS drain capacitance per μm of transistor width	$0.3\ \text{fF}/\mu\text{m}$
PMOS effective on resistance	$6.6\ \text{k}\Omega\mu\text{m}$
NMOS effective on resistance	$3.3\ \text{k}\Omega\mu\text{m}$
Parameters for Metal 2 Wire	Value
Wire resistance per unit length	$0.4\ \Omega/\mu\text{m}$
Wire capacitance per unit length	$0.2\ \text{fF}/\mu\text{m}$

Part 2.A : Delay calculation with end-of-line driver (10 points)

Draw a simple RC model for the register file write bitline. Only include the final stage of the driver. Use a lumped π wire model. Use the RC model to determine the delay of the write bitline. Express your answer in RC time constants. This part is very similar to the question asked in Lab 2.

Part 2.B : Delay calculation with middle-line driver (15 points)

There is no reason we have to position the write bitline driver at one end of the bitline. In this part we will evaluate moving the driver to the middle of the bitline. The following figure illustrates the new design.



Draw a new RC model for the register file write bitline. Use the RC model to determine the delay of the write bitline. Express your answer in RC time constants. How does this new design compare to the baseline design evaluated in Part 2.A? Does this approach help mitigate wire resistance, wire capacitance, or both?

Problem 3 : Bluespec Synthesis (25 total points)

Consider the algorithm for binary multiplication presented in Lecture 7 (Introduction to Bluespec):

```

      1001      // d = 4'd9
x 0101      // r = 4'd5
-----
      1001      // d << 0 (since r[0] == 1)
      0000      // 0 << 1 (since r[1] == 0)
      1001      // d << 2 (since r[2] == 1)
      0000      // 0 << 3 (since r[3] == 0)
-----
0101101      // product (sum of above) = 45

```

This algorithm is actually quite similar to the software multiplication algorithm you implemented for SMIPS in Lab 1. For this problem we will explore implementing this as a hardware module in Bluespec.

The following module implements this algorithm using two shifters to form an iterative multiplier:

```

interface I_mult;
  method Action start( Bit#(16) x, Bit#(16) y );
  method Bit#(32) result();
endinterface

module mkMult ( I_mult );
  Reg#(Bit#(32)) product <- mkReg(0);
  Reg#(Bit#(32)) d      <- mkReg(0);
  Reg#(Bit#(16)) r      <- mkReg(0);

  rule cycle ( r != 0 );
    if ( r[0] == 1 )
      product <= product + d;
    d <= d << 1;
    r <= r >> 1;
  endrule

  method Action start( Bit#(16) x, Bit#(16) y ) if ( r == 0 );
    d <= zeroExtend(x);
    r <= y;
    product <= 0;
  endmethod

  method Bit#(32) result() if ( r == 0 );
    return product;
  endmethod
endmodule

```

Diagram the hardware that the Bluespec compiler should produce for this module, including interface ports. Clearly circle and label which parts correspond to the rule, the scheduler, the `start` and `result` methods. Label which wire or wires correspond to `CAN_FIRE_cycle` and `WILL_FIRE_cycle`, as well as all ports corresponding to method ready and enable signals.



Problem 4 : Rule Scheduling in Bluespec (25 total points)

In this problem we will explore the behavior of the pipeline used in Lab 3 and presented in class. The reference code has been included in a separate handout.

In order to gain fine-grained control over the scheduling, it is often desirable to split large rules with case statements into multiple rules. Consider the `execute` rule. It only interacts with the `dataReqQ` on a memory operation, so one natural partitioning is to create an `execMem` rule which handles Load and Store operations.

Similarly the `execute` rule only interacts with the `pc` when the current instruction is a branch. Therefore one design choice might be to separate the handling of branch instructions into a separate rule. However this choice is actually too restrictive. In point of fact, the `execute` stage only sets `pc` on a *taken* branch. Consider the design where `execute` is split into four rules, `execALU`, `execMem`, `execBr_NotTaken`, and `execBr_Taken`.

For reference, here is the code for the `execBr_NotTaken` and the `execBr_Taken` function.

```
function Bool isBranch( Instr i );
  // Returns True if i is a Branch
endfunction

function Bool branchTaken( Instr i );
  // If given a branch instruction, returns True if the branch is taken,
  // otherwise returns False.
  // Note that in some cases this involves reading the RegFile.
endfunction

rule execBr_NotTaken ( instRespQ.first() matches tagged LoadResp .ld
  &&& ld.tag == epoch
  &&& unpack(ld.data) matches .inst
  &&& !stallfunc(inst)
  &&& isBranch(inst)
  &&& !branch_taken(inst) );

  pcQ.deq();
  instRespQ.deq();

endrule
```

```
rule execBr_Taken ( instRespQ.first() matches tagged LoadResp .ld
    &&& ld.tag == epoch
    &&& unpack(ld.data) matches .inst
    &&& !stallfunc(inst)
    &&& isBranch(inst)
    &&& branch_taken(inst) );

Addr next_pc;

case (inst) matches
  tagged J    .it :
    next_pc = { pcQ.first()[31:28], it.target, 2'b0 };
  tagged JR   .it :
    next_pc = rf.rd1(it.rsrc);
  tagged JAL  .it :
    begin
      wbQ.enq( WB_ALU {dest: 31, data: pcQ.first()} );
      next_pc = { pcQ.first()[31:28], it.target, 2'b0 };
    end
  tagged JALR .it :
    begin
      wbQ.enq( WB_ALU {dest: it.rdst, data: pcQ.first()} );
      next_pc = rf.rd1(it.rsrc);
    end
  //BLEZ, BGTZ, BTZ, BGEZ, BEQ, BNE
  default:
    next_pc = pcQ.first() + (sext(it.offset) << 2);
endcase

pc <= next_pc;
epoch <= epoch + 1;
pcQ.deq();
instRespQ.deq();

endrule
```

After splitting this system the rules have the following resource usage. (Note that the FIFO clear methods are unused.)

pcGen	discard	execALU	execMem
pc.read	epoch.read	epoch.read	epoch.read
epoch.read	pcQ.deq	instRespQ.first	instRespQ.first
pc.write	instRespQ.deq	instRespQ.deq	instRespQ.deq
pcQ.enq		pcQ.deq	pcQ.first
instReqQ.enq		wbQ.enq	pcQ.deq
		wbQ.find1,2	wbQ.enq
		rf.rd1,2	wbQ.find1,2
			rf.rd1,2
			dataReqQ.enq

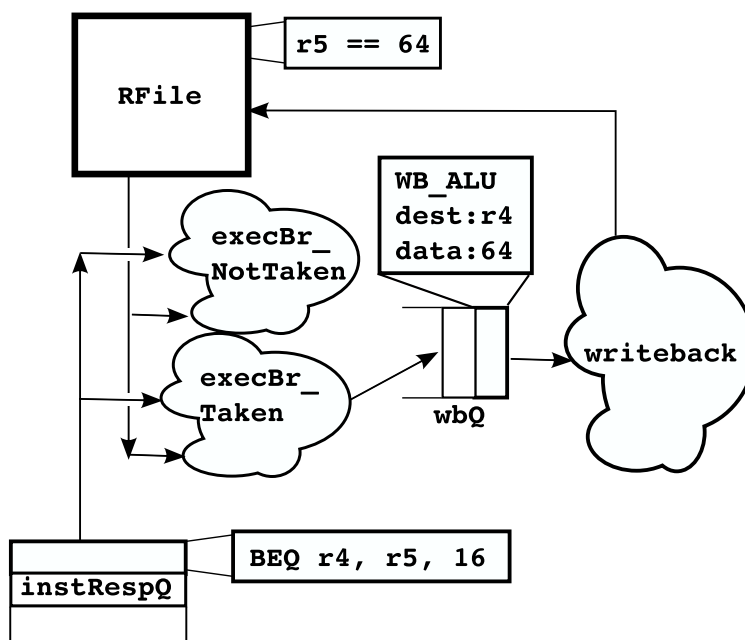
execBr_Taken	execBr_NotTaken	writeback	
epoch.read	epoch.read	wbQ.first	
instRespQ.first	instRespQ.first	wbQ.deq	
instRespQ.deq	instRespQ.deq	dataRespQ.first	
pcQ.first	pcQ.deq	dataRespQ.deq	
pcQ.deq	wbQ.find1,2	rf.wr	
wbQ.enq			
wbQ.find1,2			
rf.rd1,2			
dataReqQ.enq			
epoch.write			
pc.write			

Part 4.C : Dynamic Behavior (12 points)

Consider the following three variants of a processor:

- Behaves as if: $pcGen < execBr_NotTaken, execBr_Taken < writeback$
- Behaves as if: $writeback < execBr_NotTaken, execBr_Taken < pcGen$
- Behaves as if: $writeback < execBr_NotTaken < pcGen < execBr_Taken$

While running a program these processors reach the following state:



For each variant, answer the following. A) What rules (of those shown) will the scheduler choose to fire and why. B) What is the longest combinational path in the system (including the parts not shown)?

`pcGen < execBr_NotTaken, execBr_Taken < writeback`

A)

B)

`writeback < execBr_NotTaken, execBr_Taken < pcGen`

A)

B)

`writeback < execBr_NotTaken < pcGen < execBr_Taken`

A)

B)

Equation Sheet

Equation or Symbol	Description
g	Gate logical effort
$h = C_{out}/C_{in}$	Gate electrical effort
$f = gh$	Gate effort
p	Gate parasitic delay
p_{inv}	Parasitic delay of minimum-sized inverter
τ	Delay unit
$d = f + p$	Delay in units of τ
$d_{abs} = d\tau$	Absolute delay in seconds
$G = \prod g_i$	Path logical effort
$H = C_{out}/C_{in}$	Path electrical effort
$F = GH$	Path effort
$D = \sum d_i = \sum g_i h_i + \sum p_i$	Path delay
$f_{opt} = F^{1/N}$	Optimal stage effort
$D_{opt} = N f_{opt} + P$	Optimal path delay
$C_{in,opt,i} = C_{out,i} \times g_i / f_{opt}$	Optimal input capacitance for stage i
Delay = $\sum_{i=0}^n \left(\sum_{j=0}^{j=i} R_j \right) C_i$	Penfield-Rubenstein wire-delay model
R_d	Effective driver resistance
R_w	Total wire resistance
C_w	Total wire capacitance
Delay $\propto R_d \times C_w/2 + (R_d + R_w) \times (C_w/2 + C_{load})$	Simple lumped π model

Gate Type	Number of inputs					
	1	2	3	4	5	n
Inverter Logical Effort	1					
NAND Logical Effort		4/3	5/3	6/3	7/3	$(n+2)/3$
NOR Logical Effort		5/3	7/3	9/3	11/3	$(2n+1)/3$
XOR/XNOR Logical Effort		4	12	32		
Inverter Parasitic Delay	p_{inv}					
NAND Parasitic Delay		$2p_{inv}$	$3p_{inv}$	$4p_{inv}$	$5p_{inv}$	np_{inv}
NOR Parasitic Delay		$2p_{inv}$	$3p_{inv}$	$4p_{inv}$	$5p_{inv}$	np_{inv}
XOR/XNOR Parasitic Delay		$4p_{inv}$	$4p_{inv}$	$4p_{inv}$	$4p_{inv}$	$4p_{inv}$