

EECS 151/251A Homework 10

Due Wed, April 24th, 2024

Introduction

This homework is meant to test your understanding of parallelism and HW accelerator design. There are five total questions. Please check Ed first if you have any questions.

Problem 1

We would like to evaluate the list processor architectures in the lecture slides using the FOM (figure of merits). The FOM here is defined as:

$$\frac{\frac{\text{Nodes}}{\text{Cycle}} \cdot F}{\sqrt{\text{Cost}}}$$

For this problem, ignore the delay and cost of control logic. Use 16-bit components for the sum part (15-bit is enough though).

component	delay (ns)	cost
N -bit register with CE	$t_{clk-q} = t_{setup} = 0.5$	$32N + 4$
N -bit multiplexer	1	$3N + 12$
N -bit adder	$2\log_2(N) + 2$	$8N - 1$
Memory	10 for async read	0
N -bit zero comparator	$0.5\log_2(N)$	$4N - 4$

1. Draw a circuit diagram for an 8-bit zero comparator based on the delay and cost in the table.
2. Justify the delay and cost equations for each component (i.e explain why could these be realistic based on our FOM).
3. Formulate the FOM in terms of NPC (nodes per cycle), T (clock period), and the number of components shown below.

component	number
8-bit register with CE	a
16-bit register with CE	b
8-bit multiplexer	c
16-bit multiplexer	d
8-bit adder	e
16-bit adder	f
8-bit zero comparator	g

4. Fill out the following table for the list processor architectures in the lecture slides.

Architecture	NPC	T (ns)	a	b	c	d	e	f	g	FOM
1										
2										
3										
4										

5. Would it help if nodes are aligned in memory for the first architecture? The address of the next node is always stored at an even address. Explain why or why not.

6. Design a list processor with 32×16 memory and maximize the FOM. The memory stores the value and the next address of a node at the same address and has 10 ns acync-read delay. Use 8-bit components for the address part. No need to design a control logic. What is the maximized FOM of your design?

Problem 2 - 251 only, Extra Credit for 151

For this problem, you will design HW to implement the **add** function for an 8-bit Binary Search Tree. A Binary Search Tree is an ordered data structure comprised of nodes. Each nodes has at most 2 children. For every node n , the right child of n contains data that is greater than the data at n , while the left child contains data that is smaller. **add** takes as input an 8-bit input, $data$, and traverses through the tree nodes to place the data point. At every node, if the data is greater than the data at that node, you will then visit the **right** child node of that node. If the data is than the data at that node you will travel to the **left** child node of that node. If you receive data that is equal to the current node, then you should not add the data. If the node you wish to travel to does not exist, then you will get a free address and create a new node there. In order to get a free address, you will have access to a separate data structure, **free list**, which will be a Linked List (similar to the one discussed in lecture). You should return whether or not you added the data to the tree.




Figure 1: Adding the number 43 to the Binary Search Tree. You first start at the head, and since 43 is greater than 40, you travel to the right child node which has 45. Then you travel to 42, which has no right child. You then must create a new node by getting the free address, assigning that memory location to 42's right child pointer, and then setting the data of the corresponding node to 43. You also return a 1 to signify that you added a new node.

Here are some things to note:

1. You are given a 256 x 8-bit memory with 1 8-bit read port and 1 8-bit write port.
2. Assume there is already some data in the tree.

3. Each node in the free list is 2 8-bit words. Word 0 contains the address of the next element in the free list. Word 1 contains the data (free address available for use).
4. A Binary Search Tree node has 3 words. Word 0 has the address of the right child node, while Word 1 has the address of the left child node. Word 2 contains the data in the node. You can assume a child node does not exist if the corresponding address is 0.
5. The free list points to chunks of memory where 3 words exist contiguously. You do not have to worry about blocks not being big enough to fit a full tree node.
6. Adding a tree node means that you do the following:
 - (a) You set the corresponding left/right pointer of the parent node to the address you receive from the free list.
 - (b) You create a new node at the correct address and zero out the left and right child pointers.
 - (c) You add the inputted data as the data of the node you just created.
 - (d) You update the pointer to the free list and return.

Please do the following:

1. Write the **add** function as described using RT notation. You can use assume the current address of the free list pointer is stored in a register *FREE*. The head of the Binary Search Tree is stored in a register *HEAD*. Assume that you start working when the START signal is high, as we did in lecture. You should have one loop, designated by a repeat block. Make sure to maintain the correctness of the free list.
2. What is the minimal number of states you need for the Finite State Machine for the Controller (keep in mind what all you can do in a single cycle!)? Draw a high level Finite State Machine (In each state, write what operations are happening, explain what outcomes lead to what transitions). Make sure you draw the diagram clearly.
3. What operations from the critical path could be moved to a different state? If you were to move those operations, what extra HW do you need?
4. Given the following sequence of adds to the tree in Figure 1:
add(20), add(20), add(20),add(25), add(35)
 Which of the following architectural changes would most significantly improve the performance in terms of the average **latency** (in terms of cycles) of an an **add** request?
 - (a) Adding 1 read port
 - (b) Adding 1 write port
 - (c) Modifying the read and write ports to be 24 bits wide (3 words)

Problem 3:

You will be designing a 2-way set associative cache using a single read and write port RAM. If both ways are full, the cache always kicks out the data from Way 0. If the data is dirty (has been written to by the CPU), then the data is written to memory. The cache has 16 lines (8 lines per set) which are 128 bits wide. Both memory address and data words are 32 bits wide. The interface between the CPU and cache is shown below.


Between CPU and cache:

- Read
 - CPU waits til `cache_ready` is high.
 - CPU sets `ren` to 1 and `addr` to the source address.
 - Cache reads those signals immediately. The cache then checks both ways and set `valid` to 1 and `dout` to the data in case of a cache hit.
 - On a miss, the caches sets `cache_ready` to low and checks if both ways are valid. If both are valid, then the cache must perform a memory write with the data from way 0 if it is dirty. If both are invalid, you may read into either way. If one is invalid, you must read into that way.
 - The cache then performs a memory read into the appropriate way.
 - When memory returns the data (`mem_valid` is 1), cache sets `valid` to 1 and `dout` to the read data. The cache additionally set `cache_ready` high again.
 - Cache stores the tag and data at the next positive clock edge.
 - Write
 - CPU waits til `cache_ready` is high.
 - CPU sets `wen` to 1, `addr` to the destination address, and `din` to the data to write.
 - On a hit, cache stores the tag and the data. It also sets the dirty bit to 1.
 - On a miss, the caches sets `cache_ready` to low and checks if both ways are valid. If both are valid, then the cache must perform a memory write with the data from way 0 if it is dirty. If both are invalid, you may read into either way. If one is invalid, you must read into that way.
 - The cache then performs a memory read into the appropriate way (must read the the rest of the block) and then writes the data. It must set the corresponding dirty bit high.
 - When the cache is done writing the data, it sets `cache_ready` high again.
1. Draw the structure of the cache. Make sure to show the two ways, the tags for each line, as well as any additional flag bits needed. How big (in bits) is the full cache (include the size of the tag, flag bits, and data)?
 2. Draw the FSM for the cache controller. Assume the memory responds with a `mem_ready` signal on a memory read. Caches do not need to wait on a memory write. Assume to read from memory you set the `mem_rd` signal high (and then set it low on receive) and to write to memory you set `mem_write` high (and then you can set it low on the next clock signal).

3. Draw a full block diagram with the complete controller logic. You may use logic gates, multiplexers, arithmetic blocks and flip-flops.

Problem 4

Given the following FSM:



1. Draw the circuit diagram for the FSM using binary encoded states.
2. Draw the circuit diagram for the FSM if we receive 3 bits per cycle (you only need to output the final output). Make sure use loop unrolling.

Problem 5

```
char* strstr(const char *main_str, const char *substr);
```

`strstr` is a C function that takes 2 strings, `main_str` and `substr` and outputs a 1 if `substr` is in `main_str`.

For this problem, assume that you are designing an accelerator connected to a 1-byte wide memory block. Assume the accelerator is given the address of the two strings. The strings are formatted as traditional ASCII characters followed by a null character. Perform the following:

1. Write a RT Notation for a simple version of this accelerator.
2. Draw the datapath and the state transfer diagram of the controller.
3. Write a RT Notation for a performance-optimized design.
4. Describe the revised datapath for performance optimization.