



Lund University

Computer Architecture, EITF20 Exercise 2 Solutions

Mohammad Attari (Masoud)

December 17, 2019

Problem 1

In the following i1 and i2 are instructions immediately after the branch.

a)

Instruction	Clock cycle no.																		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
ld x1, 0(x2)	F	D	X	M	W														
addi x1, x1, 1		F	D	D	D	X	M	W											
sd x1, 0(x2)			F	F	F	D	D	D	X	M	W								
addi x2, x2, 4						F	F	F	D	X	M	W							
sub x4, x3, x2								F	D	D	D	X	M	W					
bnez x4, Loop									F	F	F	D	D	D	X	M	W		
i1												F	F	F	D				
i2															F				
ld x1, 0(x2)																F	D	X	

There are 16 cycles between loop instances, the last loop takes 18 cycles:

X3 is x2+40 so the loop takes 10 iterations. 9 loops with 16 cycles, and 1 with 18. T=9*16+18

b)

Instruction	Clock cycle no.																		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
ld x1, 0(x2)	F	D	X	M	W														
addi x1, x1, 1		F	D	D	X	M	W												
sd x1, 0(x2)			F	F	D	X	M	W											
addi x2, x2, 4					F	D	X	M	W										
sub x4, x3, x2						F	D	X	M	W									
bnez x4, Loop							F	D	D	X	M	W							
i1								F	F										
ld x1, 0(x2)										F	D	X	M	W					

There are 9 cycles between loop instances, the last loop takes 12 cycles:

X3 is x2+40 so the loop takes 10 iterations. 9 loops with 9 cycles, and 1 with 12. T=9*9+12

c)

Instruction	Clock cycle no.																		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
ld x1, 0(x2)	F	D	X	M	W														
addi x1, x1, 1		F	D	D	X	M	W												
sd x1, 0(x2)			F	F	D	X	M	W											
addi x2, x2, 4					F	D	X	M	W										
sub x4, x3, x2						F	D	X	M	W									
bnez x4, Loop							F	D	X	M	W								
i1								F	D										
ld x1, 0(x2)									F	D	X	M	W						

There are 8 cycles between loop instances, the last loop takes 12 cycles:

X3 is x2+40 so the loop takes 10 iterations. 9 loops with 8 cycles, and 1 with 12. T=9*8+11

Problem 2

a) States: F = Fetch, D = Decode, E = Execute, M = Memory, W = Writeback

Instruction	Clock cycle no.																		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
LW R5, (R3)	F	D	E	M	W														
LW R6, 8(R3)		F	D	E	M	W													
MUL R5, R5, R6		F	D	D	D	E1	E2	E3	E4	E5	M	W							
ADD R3, R3, R2			F	F	F	O	D	D	D	D	E1	E2	M	W					
SW R5, (R3)				-	-	F	F	F	F	F	D	D	D	E	M	W			

b) States: I = Issue, E = Execute, M = Memory, W = Writeback

Instruction	Clock cycle no.																		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
LW R5, (R3)	I	E	M	W															
LW R6, 8(R3)	I	E	M	W															
MUL R5, R5, R6		I	-	-	E1	E2	E3	E4	E5	W*									
ADD R3, R3, R2			I	E1	E2	W*													
SW R5, (R3)				W	-	-	E**	-	-	-	M	W							

*No M stage needed for Mul and Add

**SW can start E in cycle 8 since it's only address calculation and by that cycle R3 is already read from CDB

c) States: I = Issue, R=ReadOp, E = Execute/Memory, W = Writeback

Instruction	Clock cycle no.																		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
LW R5, (R3)	I	R	E	W															
LW R6, 8(R3)	I	R	E	W															
MUL R5, R5, R6			I*	R	E1	E2	E3	E4	E5	W									
ADD R3, R3, R2				I	R	E1	E2	W											
SW R5, (R3)					I	-	-	-	-	-	R	E	W						

*Issue has to wait for WAW hazard on R5

problem 3

$$\text{Miss Rate} = \frac{\text{misses}}{\text{memory accesses}} = \frac{\text{misses/instruction}}{\text{memory accesses/instruction}}$$

$$\text{miss rate}_{16\text{KB inst}} = \frac{3.82/1000}{7} = .004$$

$$\text{miss rate}_{16\text{KB data}} = \frac{40.9/1000}{0.36} = .114$$

$$\text{miss rate}_{32\text{KB unified}} = \frac{43.3/1000}{1 + 0.36} \approx .0318 = 3.18\%$$

$$\% \text{ instruction references} = \frac{1}{1.36} = 74\%$$

$$\% \text{ data references} = \frac{.36}{1.36} = 26\%$$

$$\text{miss rate}_{\text{split}} = (74\% \times .004) + (26\% \times .114) = .0326 = 3.26\%$$

$$3.18\% < 3.26\%$$

\Rightarrow unified cache comes out on top by a bit
 When we use effective miss rate as a performance metric

problem 3

Average memory access time (AMAT) =

$$\frac{1}{\text{instruction references}} \times (\text{hit time} + \text{instruction miss rate} \times \text{miss penalty}) \\ + \frac{1}{\text{data references}} \times (\text{hit time} + \text{data miss rate} \times \text{miss penalty})$$

$$\text{AMAT}_{\text{split}} = 74\% \times (1 + 0.004 \times 200) + 26\% (1 + 0.114 \times 200) \\ = 7.52$$

{ alternatively:

$$\text{AMAT}_{\text{split}} = 1 + 3.26\% \times 200 = 7.52$$

$$\text{AMAT}_{\text{unified}} = 74\% (1 + 0.0318 \times 200) + 26\% (1 + 0.0318 \times 200) \\ = 7.62$$

{ alternatively:

$$100\% (1 + 0.0318 \times 200) + 26\% (1) = 7.62$$

$$7.52 < 7.62$$

split cache has a better AMAT

Problem 4

$$\text{Execution time} = IC \times \left(CPI_{\text{idle}} + \frac{\text{mem accesses}}{\text{instruction}} \times \text{miss rate} \times \text{miss penalty} \right) \times T_C$$

$$\text{speedup} = \frac{1}{1-p + p_s}$$

$$\text{Execution time}_{\text{enhanced}} = IC \times \left(\frac{CPI_{\text{idle}}}{\text{speedup}} + \dots \right) \times T_C$$

$$1) \text{ speedup}_{\text{int}} = \frac{1}{1 - .423 + \frac{.423}{1.8}}$$

$$1.3341$$

$$\begin{aligned} \text{Execution time}_{\text{int}} &= IC \times \left(2.3 \left(1 - .423 + \frac{.423}{1.8} \right) + \underbrace{\left(1 + .2375 + .0966 \right)}_{\text{instructions}} \times 0.207789 \right. \\ &\quad \left. \times 50 \right) \times T_C \end{aligned}$$

$$= 15.72$$

$$2) \text{ speedup}_{\text{data}} = \frac{1}{1 - (.2375 + .0966) + \frac{(.2375 + .0966)}{2}}$$

$$\text{Execution time}_{\text{data}} = IC \times \left(2.3 \times \left(1 - .3341 + \frac{.3341}{2} \right) + 1.3341 \times 0.207789 \times 50 \right) \times T_C$$

$$= 15.78$$

problem 4

9)

$$\text{Extraction time}_{\text{each}} = TC \times (2.3 + 1.3341 \times 0.136385 \times 50) \times 1.1 \times TC$$

$$= 12.59$$

best performance

Problem 5

a)

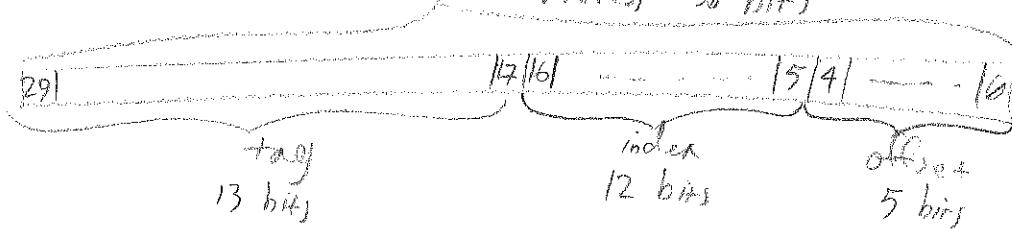
16B memory \Rightarrow 16B = 2^{30} \Rightarrow 30-bit address

$$\text{No of sets} = 2^{\text{index}} = \frac{\text{cache size}}{\text{block size} * \text{blocks/sets}}$$

$$= \frac{512 \text{ kB}}{32 \text{ B} * 4} = \frac{2^{19}}{2^5 * 2^2}$$

$$2^{\text{index}} = 2^{12} \Rightarrow \text{index} = 12 \text{ bits}$$

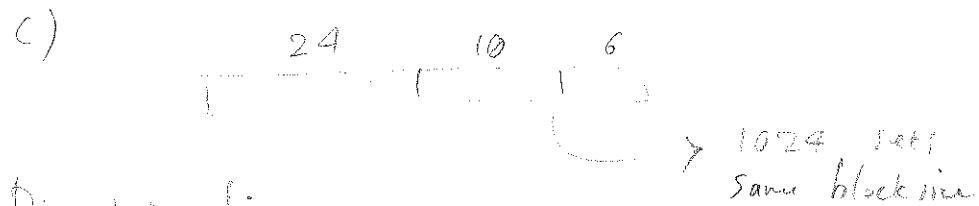
$$\text{block size} = 32 \text{ B} = 2^5 \text{ B} \Rightarrow \text{offset} = 5 \text{ bits}$$



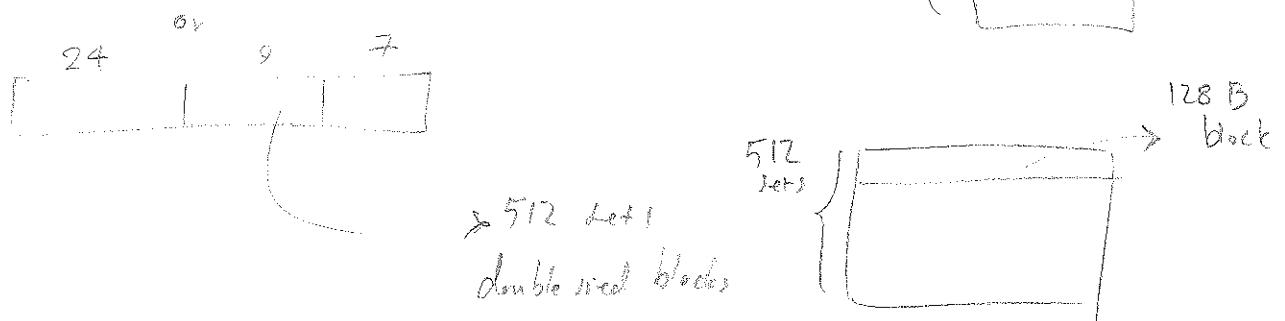
$$\text{tag} = 30 - 12 - 5 = 13$$

Problem 5

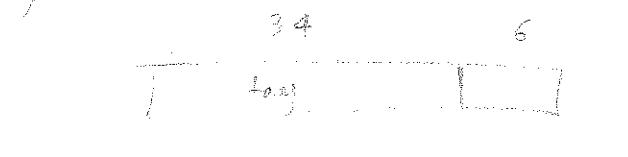
- b)
- ① break down the address into fields
 - ② select the set based on index
 - ③ read the tag (2, since it's 2-way association) in the cache
→ compare with the address tag (valid bit must be set).
 - ④ if it has a hit, send proper portion to the CPU,
if it was a miss, get the block from low-level memory



Direct mapping:



Fully association:

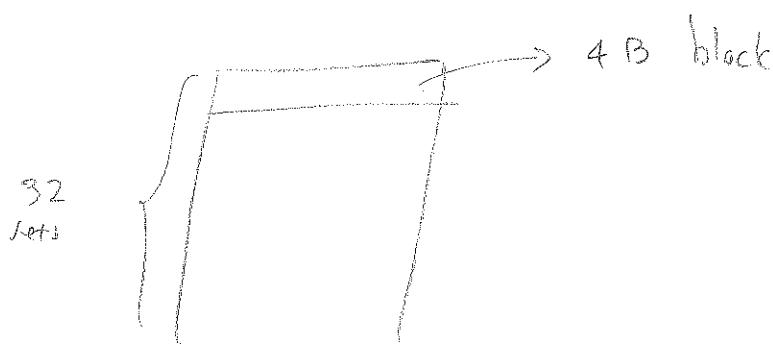


Same cache size means 1024 blocks in 7 set!



Problem 5

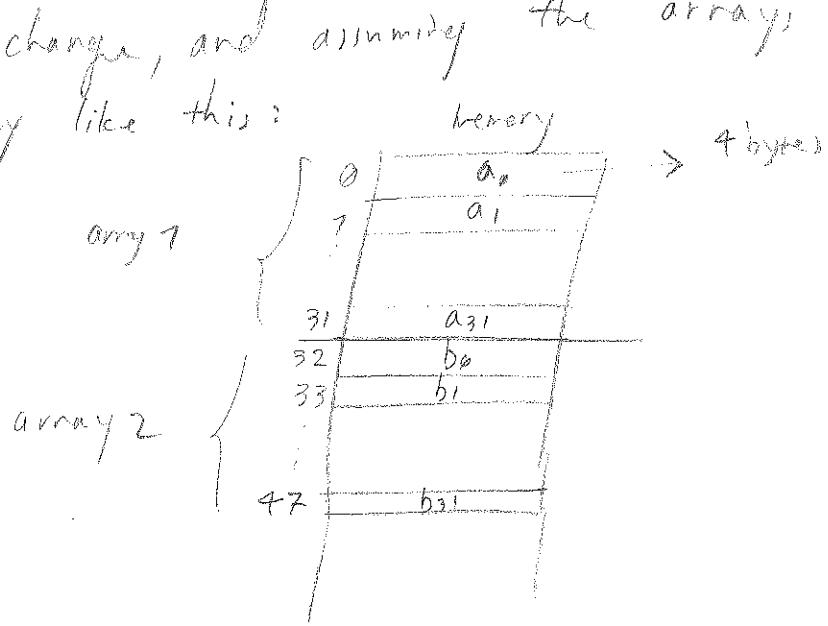
d) the cache looks like this:



the loop iterates over 2 arrays (pointed to by R5 and R6)

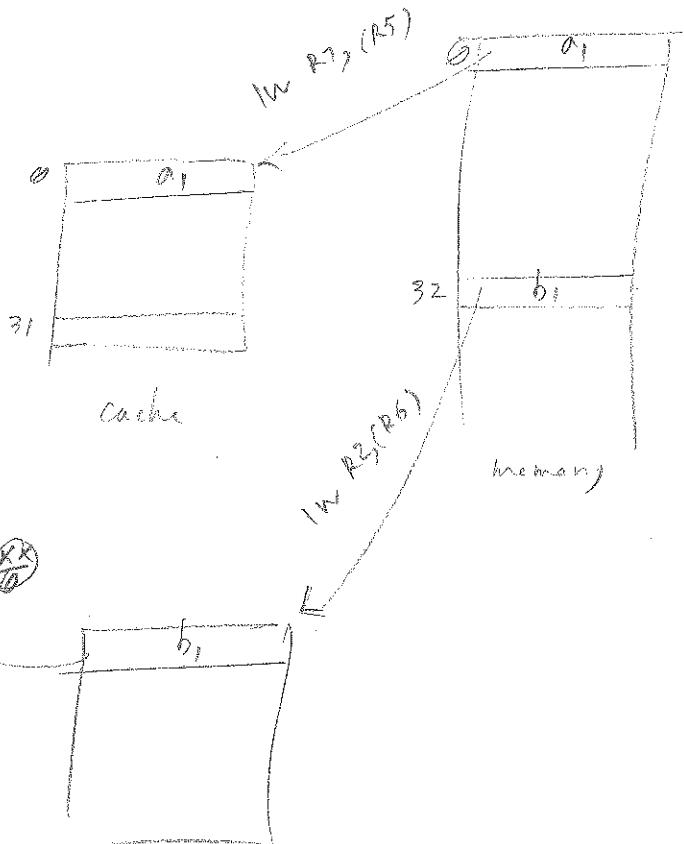
for 32 iterations, ^{one at} load the arrays' elements in each iteration, add them up, store the sum in R4, and repeats for the next element in the arrays.

with no block in cache, and assuming the arrays are stored in main memory like this:

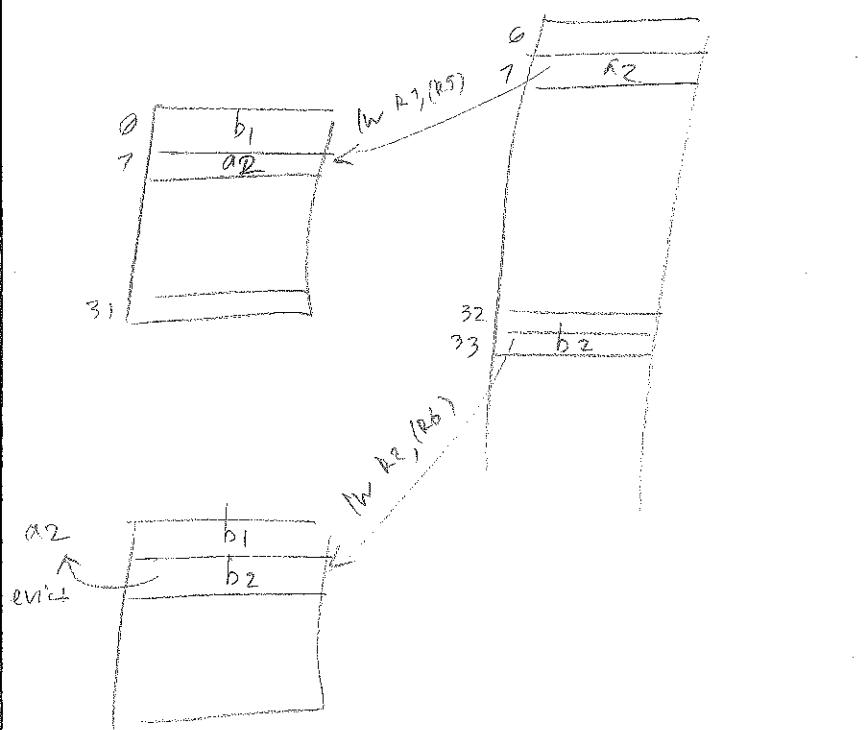


each LW brings 4B (a block) into the cache.

so in the first iteration of the loop, we would have:



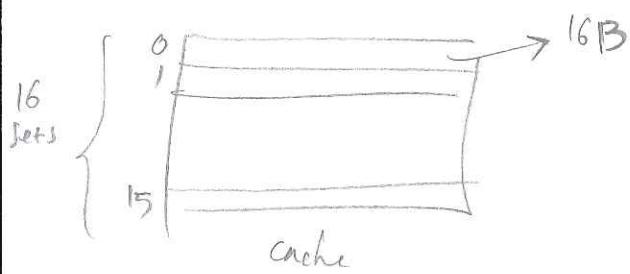
2nd iteration :



and so on!

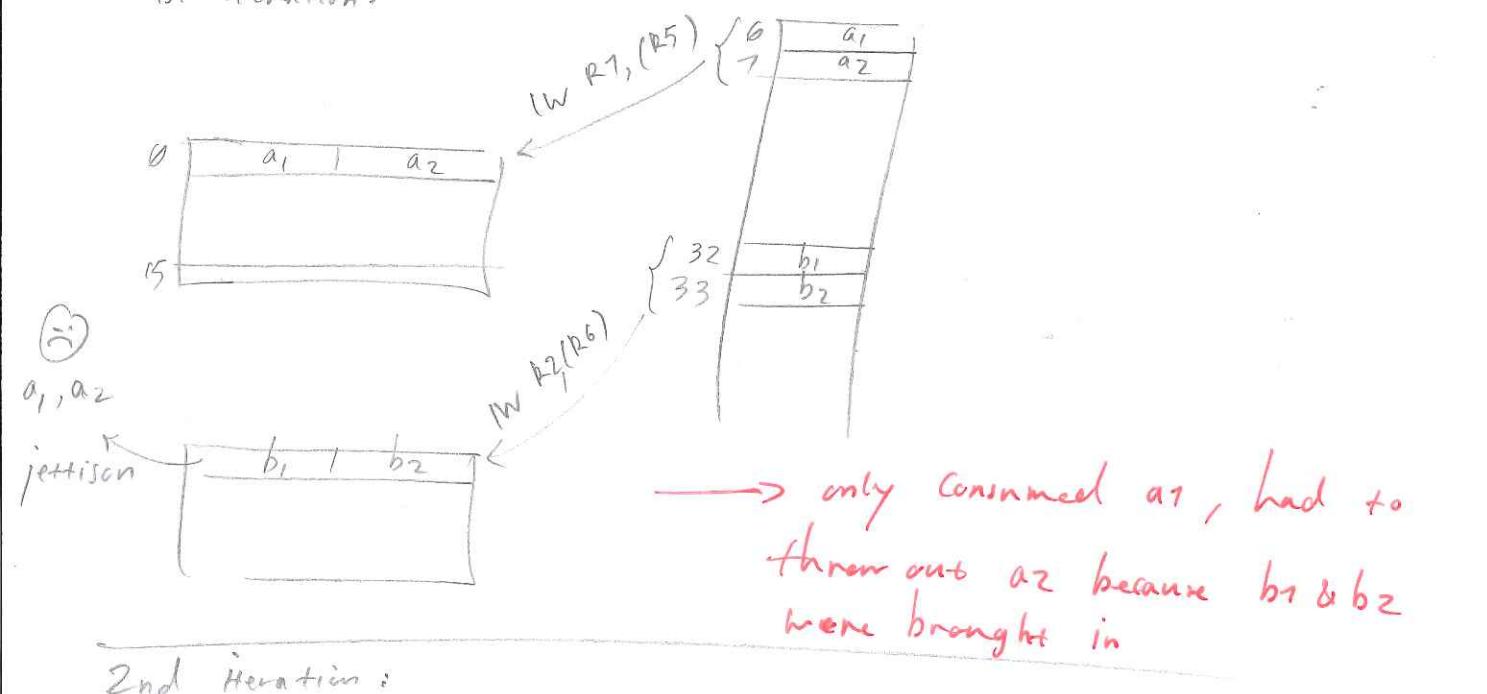
problem 5

now by increasing block size to 8B & reducing sets to 16 we have:

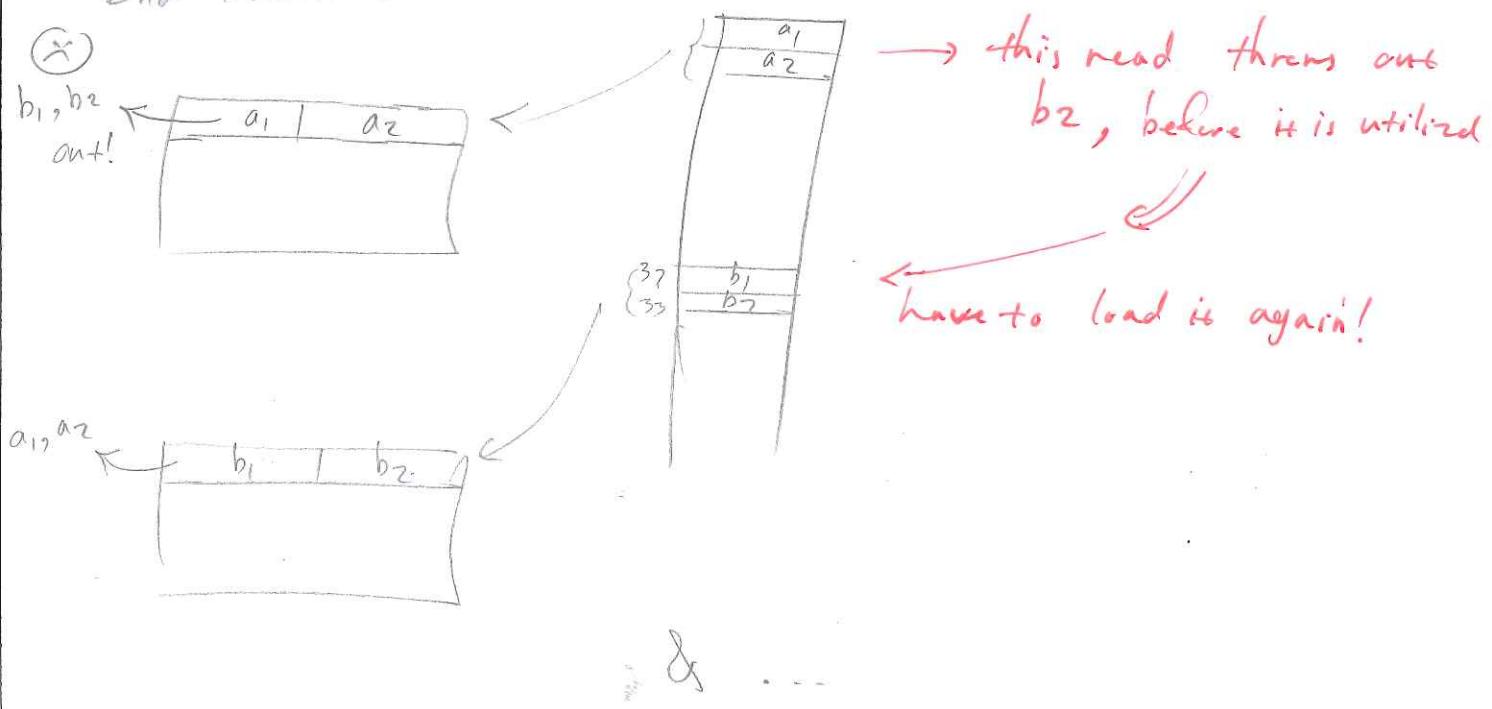


now the loop if arrays stay the same in memory

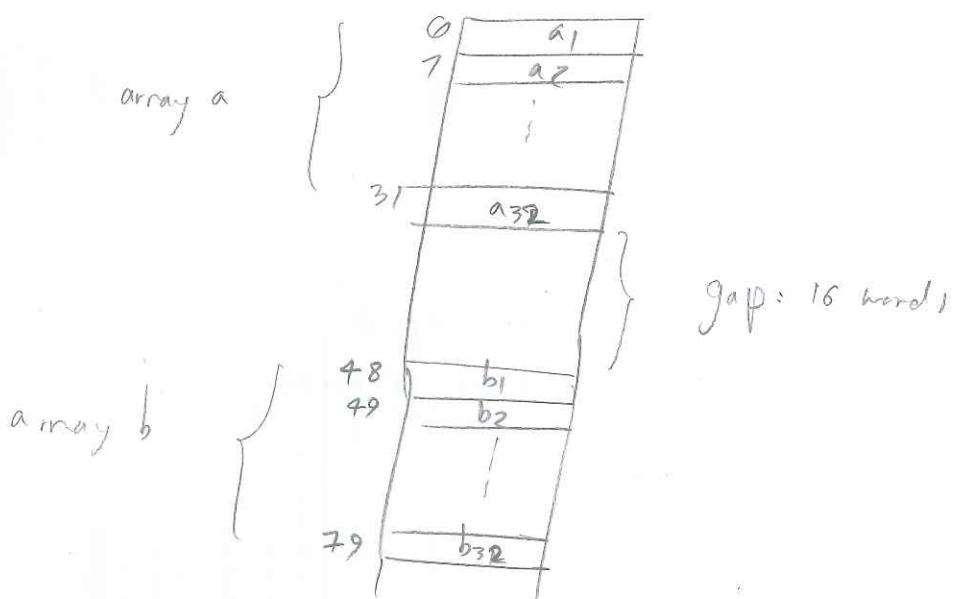
1st iteration:



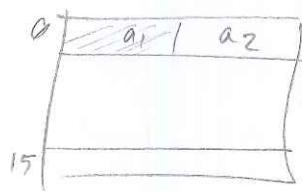
2nd iteration:



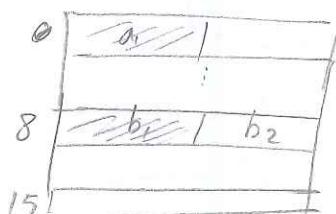
now if
the arrays are placed like this in memory:



Cache after 1st load, 1st iteration:

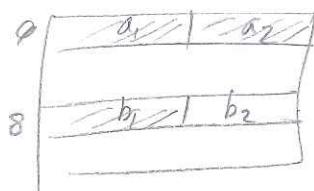


Cache after 2nd load, 1st iteration:

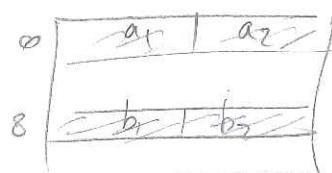


Cache after 1st load, 2nd iteration

→ a2 already in the cache. yay baby! ☺



Cache after 2nd load, 2nd iteration



→ b2 also in the cache. yay!

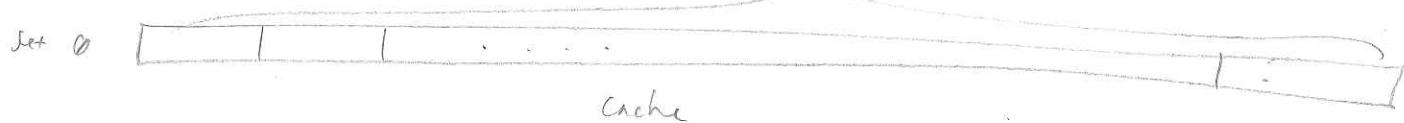
cache after 1st load, 3rd iteration

0	a ₀₀	a ₀₁
1	a ₀₂	a ₀₃
2	a ₁₀	a ₁₁
3	a ₁₂	a ₁₃
4	b ₀₀	b ₀₁
5	b ₀₂	b ₀₃
6	b ₁₀	b ₁₁
7	b ₁₂	b ₁₃

& ..

in the extreme case, we have a block size of 128B, just 1 set!
that is:

128B block



1st load, 1st iteration:

a ₀₀	a ₀₁	a ₀₂	a ₀₃	- - - - -	a ₃₀	a ₃₁
-----------------	-----------------	-----------------	-----------------	-----------	-----------------	-----------------

brings 32 elements, consumed only 1.

2nd load, 2nd iteration:

b ₀₀	b ₀₁	b ₀₂	b ₀₃	- - - - -	b ₃₀
-----------------	-----------------	-----------------	-----------------	-----------	-----------------

each time we bring 32 elements & consume only one & throw the rest out.