

CSCI 2021: Practice Final Exam SOLUTION

Spring 2023

University of Minnesota

Exam period: 20 minutes

Points available: 40

Background: Nearby are several C files along with two attempts to compile them on the left. Study these and answer the questions that follow.

```

1 > gcc vf_weak_var.c vf_strong_func.c vf_main.c # COMPILER 1
2 /usr/bin/ld: warning: size of symbol 'foo' changed from 4 to 14
3 /usr/bin/ld: warning: type of symbol 'foo' changed from 1 to 2
4 > file a.out
5 a.out: ELF 64-bit LSB pie executable, x86-64, version
6 > ./a.out
7 -573193927
8 4
9 > rm a.out
10
11 > gcc vf_strong_var.c vf_strong_func.c vf_main.c # COMPILER 2
12 /usr/bin/ld: multiple definition of 'foo';
13 collect2: error: ld returned 1 exit status
14 > file a.out
15 a.out: cannot open 'a.out' (No such file or directory)

```

```

1 // FILE: vf_main.c
2 #include <stdio.h>
3 int foo(int x);
4 int main(){
5     printf("%d\n",foo);
6     printf("%d\n",foo(2));
7     return 0;
8 }
9
10 // FILE: vf_strong_func.c
11 int foo(int x){
12     return 2*x;
13 }
14
15 // FILE: vf_strong_var.c
16 int foo = 0;
17
18 // FILE: vf_weak_var.c
19 int foo;

```

Problem 1 (10 pts): Why does COMPILER 1 succeed while COMPILER 2 fails? Mention pertinent properties of ELF files in your answer.

SOLUTION: COMPILER 1 succeeds because the integer int foo is uninitialized and therefore weak. It is overridden by the strong symbol int foo(int x) so the resulting ELF file has only the function version. COMPILER 2 fails as the C file initializes int foo=0; making both definitions strong. Two strong symbols with the same name cannot exist in an ELF file causing linking to fail.

Problem 2 (10 pts): Nearby is the output of pmap showing page table virtual memory mapping information for a running program called memory_parts. Answer the following questions about this output.

(A) The mapped memory references something called libc-2.26.so. Describe this entity and what kind of information you would expect to find at the mapped locations.

SOLUTION: This is the C standard library. It is a shared object with the .so extension and is likely to contain binary assembly instructions standard C functions like printf() and malloc().

(B) Why does pmap only show a limited number of virtual addresses? What would happen if the program attempted to access an address not listed in the output? Example: address 0x00 is not in the listing.

SOLUTION: The page table only contains mapped pages for program. These mapped addresses are what is shown. The large number of other addresses are unmapped. Attempting to access these unmapped addresses will result in errors such as segmentation faults; this usually causes the program to be immediately terminated.

```

> pmap 7986
7986: ./memory_parts
00005579a4abd000      4K r-x-- memory_parts
00005579a4cbd000      4K r---- memory_parts
00005579a4cbe000      4K rw--- memory_parts
00005579a4cbf000      4K rw--- [ anon ]
00005579a53aa000     132K rw--- [ heap ]
00007f441f2e1000    1720K r-x-- libc-2.26.so
00007f441f48f000    2044K ----- libc-2.26.so
00007f441f68e000     16K r---- libc-2.26.so
00007f441f692000     8K rw--- libc-2.26.so
00007f441f694000     16K rw--- [ anon ]
00007f441f698000    148K r-x-- ld-2.26.so
00007f441f88f000     8K rw--- [ anon ]
00007f441f8bb000     4K r---- gettysburg.txt
00007f441f8bc000     4K r---- ld-2.26.so
00007f441f8bd000     4K rw--- ld-2.26.so
00007f441f8be000     4K rw--- [ anon ]
00007fff96ae1000    132K rw--- [ stack ]
00007fff96b48000     12K r---- [ anon ]
00007fff96b4b000     8K r-x-- [ anon ]
total                4276K

```

Problem 3 (10 pts): Below is an initial memory/cache configuration along with several memory load operations. Indicate whether these load operations result in cache hits or misses and show the state of the cache after these loads complete.

----- SOLUTION -----

MAIN MEMORY				DIRECT-MAPPED Cache, 8-byte lines 4 Sets, 8-bit Address = 3-bit tag			
Addr	Addr Bits	Value		Set	V	Tag	Blocks/Line
				0-3			0-3 4-7
10	000 10 000	10					
14	000 10 100	11					
18	000 11 000	12					
1C	000 11 100	13					
20	001 00 000	20		00	1	010	200 201
24	001 00 100	21		01	1	001	22 23
28	001 01 000	22		10	1	000	10 11
2C	001 01 100	23		11	0	-	-
30	001 10 000	100					
34	001 10 100	101					
38	001 11 000	102					
3C	001 11 100	103					
40	010 00 000	200					
44	010 00 100	201					
48	010 01 000	202					
4C	010 01 100	203					

HITS OR MISSES?	
OPERATION	HIT/MISS?
1. Load 0x48	Miss
2. Load 0x4C	Hit
3. Load 0x24	Miss

FINAL CACHE STATE						
Set	V	Tag	Blocks/Line	0-3	8-7	Changed?
00	1	001	20	21		***
01	1	010	202	203		***
10	1	000	10	11		
11	0	-	-	-		

Problem 4 (10 pts): Nearby is the definition for `base_scalvec()` which scales a vector by multiplying each element by a number. Write an optimized version of this function in the space provided. Mention in comments why you performed certain transformations.

```

1 int vget(vector_t vec, int idx){
2     return vec.data[idx];
3 }
4 void vset(vector_t vec, int idx, int x){
5     vec.data[idx] = x;
6 }
7 void base_scalevec(vector_t *vec, int *scale){
8     for(int i=0; i < vec->len; i++){
9         int cur = vget(*vec,i);
10        int new = cur * (*scale);
11        vset(*vec,i,new);
12    }
13 }

```

```

1 ////////////////////////////////////////////////// SOLUTION ///////////////////////////////////
2 void opt_scalevec(vector_t *vec, int *scale){
3     // locals to avoid memory access
4     int *data = vec->data, len = vec->len;
5     int scal = (*scale), i;
6     // unroll x2 with duplicate vars to
7     // enable pipelining
8     for(i=0; i < len-2; i+=2){
9         // no function calls - inline bodies
10        // to improve register use
11        int cur0 = data[i+0];
12        int new0 = cur0 * scal;
13        data[i+0] = new0;
14        int cur1 = data[i+1];
15        int new1 = cur1 * scal;
16        data[i+1] = new1;
17    }
18    // cleanup loop
19    for(; i<len; i++){
20        int cur0 = data[i+0];
21        int new0 = cur0 * scal;
22        data[i+0] = new0;
23    }
24 }

```