



Department of Electrical Engineering and Computer Science

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

6.5660 Spring 2023

Quiz I

You have 80 minutes to answer the questions in this quiz. In order to receive credit you must answer the question as precisely as possible.

Some questions are harder than others, and some questions earn more points than others. You may want to skim them all through first, and attack them in the order that allows you to make the most progress.

If you find a question ambiguous, be sure to write down any assumptions you make. Be neat and legible. If we can't understand your answer, we can't give you credit!

Write your name and submission website email address on this cover sheet.

**This is an open book, open notes, open laptop exam.
NO INTERNET ACCESS OR OTHER COMMUNICATION.**

This quiz is printed double-sided.

Please do not write in the boxes below.

I (xx/12)	II (xx/10)	III (xx/10)	IV (xx/14)	V (xx/10)	VI (xx/10)	VII (xx/10)	VIII (xx/4)	Total (xx/80)

Name:

Submission website email address:

You can answer the feedback questions on the back of the quiz before the official start time.

This page intentionally left blank.

I Paper reading questions

1. [4 points]:

Which of the following statements are true about the OKWS system (as described in the assigned reading)?

(Circle True or False for each choice.)

- A. **True / False** The database proxy requires services to authenticate to the proxy server before being able to issue queries.
- B. **True / False** When each service starts, it must register with the okd dispatcher to receive HTTP requests.
- C. **True / False** The okld launcher can access the state of every service.
- D. **True / False** Each service is responsible for constructing the exact HTTP response that will be sent back on a client's TCP connection.

2. [4 points]:

Which of the following statements are true about iOS devices, according to the assigned reading?

(Circle True or False for each choice.)

- A. **True / False** The Secure Enclave encrypts and authenticates the data it stores in off-chip DRAM because it does not trust the CPU/operating system.
- B. **True / False** In order to implement secure boot, the boot ROM will contain the secret key of Apple SK_{APPLE} .
- C. **True / False** If an adversary learns the Exclusive Chip Identification (ECID) of a device, they can perform a downgrade attack on that device.
- D. **True / False** The secret UID of the device is entangled with the passcode to produce a file class key that is deleted once the device is locked.

This page intentionally left blank.

3. [4 points]:

Which of the following statements are true about the EXE symbolic execution system (as described in the assigned reading)?

(Circle True or False for each choice.)

- A. True / False** Given enough time, EXE will find every possible execution of the application that leads to a bug.

- B. True / False** EXE can reason about integer overflow in C code.

- C. True / False** EXE's compiler translates a C program into an SMT query to STP.

- D. True / False** EXE can reason about calls to a symbolic function pointer value.

This page intentionally left blank.

II Timing attacks

Consider the following C function:

```
int compareTwoStrings(char b[])
{
    char *s = "?????";
    int flag = 0;
    int i = 0;
    while (s[i] != 0 && b[i] != 0) {
        if (s[i] != b[i]) {
            flag = 1;
            break;
        }
        i++;
    }
    if (s[i] != 0 || b[i] != 0)
        return 1;
    if (flag == 0)
        return 0;
    else
        return 1;
}
```

The above code is running on a CPU and you know everything about it except the characters in the string s , but you do know the length of s beforehand, e.g., $n = 5$, $n = 7$. You are allowed to call the function `compareTwoStrings()` repeatedly with different arguments that you can control. You can also assume that you can time each function call precisely.

4. [10 points]: Describe a timing attack that allows you to discover the characters in s . Describe all necessary calls to `compareTwoStrings()` with appropriate arguments, and explain how you can infer characters in s through timing measurements. You can only make a number of calls that is polynomial in the length of s , i.e., n . Specify how many total calls are made.

This page intentionally left blank.

III Lab 1

5. [10 points]: Ben Bitdiddle built an exploit for lab 1's exercise 5 (return-to-libc using the accidentally function) that sends the following payload req to the server from his exploit-5.py:

```
accidentally_addr = 0x555555556bbb
unlink_addr      = 0x2aaaab246ea0
grades_str_addr  = 0x7fffffffed20

req = "A" * n + \
    struct.pack("<Q", accidentally_addr) + \
    struct.pack("<Q", unlink_addr) + \
    struct.pack("<Q", grades_str_addr) + \
    b"/home/student/grades.txt"
```

For your reference, here is the disassembly of accidentally:

```
0x0000555555556bbb <accidentally+0>: endbr64
0x0000555555556bbf <accidentally+4>: push  %rbp
0x0000555555556bc0 <accidentally+5>: mov   %rsp,%rbp
0x0000555555556bc3 <accidentally+8>: mov   0x10(%rbp),%rdi
0x0000555555556bc7 <accidentally+12>: nop
0x0000555555556bc8 <accidentally+13>: pop   %rbp
0x0000555555556bc9 <accidentally+14>: ret
```

Ben wants to change the attack to jump into the middle of accidentally, bypassing the initial push %rbp; that is, he wants his attack to jump to accidentally_addr+5. Write down a new attack payload that Ben can use instead of the payload shown above, which jumps to accidentally_addr+5 instead of accidentally_addr but still succeeds. If you refer to unlink_addr, accidentally_addr, and grades_str_addr in your attack payload, your attack should work with those constants having the exact 64-bit value shown in Ben's original exploit.

This page intentionally left blank.

IV Lab 2

6. [8 points]: Suppose that you modify `/home/student/lab/zoobar/auth-server.py` and then launch the web server using `zookld.py`. At what path (if any) does the modified file exist in the bank container?

7. [6 points]: Before you fixed the security problems pointed out in lab 2 exercise 11 to protect the system library files, suppose that you ran some profile code that wrote to `/lib/hello.txt`. At what paths (if any) does this file exist in the `profile` and `dynamic` containers, respectively?

This page intentionally left blank.

V WebAssembly

Ben Bitdiddle is writing a compiler from C to WebAssembly. Ben's compiler stores all of the variables from the C program—including all of the variables on the C stack—in the WebAssembly memory, and does not use the WebAssembly stack, local variables, or global variables for storing its C variables.

8. [10 points]: Suppose there is a buffer overflow in some C program compiled with Ben's compiler. Could an adversary exploit that buffer overflow to escape from WebAssembly? Explain how or why not.

This page intentionally left blank.

VI KSplit

Alyssa P. Hacker is writing a Linux driver for her device, and the interface with the kernel is as follows (the kernel periodically runs `use_device(d)` to invoke the device driver's `bar` function):

```
struct device {
    char buf[128];
    void (*bar) ();
}

void use_device(struct device *d) {
    driver_foo(d);
    d->bar();
}
```

Ben Bitdiddle downloads Alyssa's driver, runs KSplit to generate the IDL file, and uses the generated IDL file to compile and run the resulting kernel and driver.

9. [10 points]: Suppose Alyssa's driver had the following code in its implementation of `driver_foo()`:

```
void driver_foo(struct driver *d) {
    d->bar = &machine_shutdown;
}
```

where `machine_shutdown()` is the Linux kernel function that powers off the computer.

Will Alyssa's driver cause Ben's computer to call `machine_shutdown`? Explain why or why not.

This page intentionally left blank.

VII Knox

Ben Bitdiddle has an implementation of a PIN-protected backup HSM that he verifies to be secure (i.e., match the PIN-protected backup specification from the Knox paper, as shown in Figure 2).

10. [10 points]: Ben modifies the HSM implementation to add a status light that shows when the HSM is busy processing a request. Specifically, he adds an extra output wire (intended to be wired up to the status light) that goes high when the HSM receives a request from the host, and goes low when the HSM sends the response back to the host.

Could this new implementation be proven secure with respect to the same specification? Explain in what case this could be done, or explain why it cannot be done.

VIII 6.5660

We'd like to hear your opinions about 6.5660. Any answer, except no answer, will receive full credit.

- 11. [4 points]:** Is there one paper out of the ones we have covered so far in 6.5660 that you think we should definitely remove next year? If not, feel free to say that.

End of Quiz