

Write your responses to following questions on separate paper. Use complete sentences where appropriate and write out code using proper style and syntax. You can hand write your responses or type set with a computer.

1. Consider the problem of writing a function, `void to_lower(char* s)`, that replaces all uppercase characters in the C-style string `s` with their lowercase equivalents. For example, `Hello, World!` becomes `hello, world!`. Do not use any standard library functions. A C-style string is a zero-terminated array of characters, so if you find a `char` with the value 0 you are at the end.

Below you will find a number of solutions to that problem for comparison.

```

1 #include <iostream>
  #include <string>
3 using namespace std;
  // solution #1, using subscripting:
5 void to_lower(char* p) {
    // assume that lower and upper case characters have consecutive integer values
7     for (int i = 0; p[i]; ++i) {
        if ('A' <= p[i] && p[i] <= 'Z') // is upper case (only characters are considered
9             // upper case - e.g. not !@#$)
                p[i] -= ('A' - 'a'); // make lower case
11 }
}

```

To avoid repeating `p[i]`, we could introduce a variable to hold the correct character:

```

// solution #1.1, using subscripting and a local variable:
2 void to_lower(char* p) {
    // assume that lower and upper case characters have consecutive integer values
4     for (int i = 0; char c = p[i]; ++i) {
        if ('A' <= c && c <= 'Z') // is upper case (only characters are considered
        upper
6             // case - e.g. not !@#$)
                p[i] -= ('A' - 'a'); // make lower case
8     }
}

```

- (a) Is this really an improvement? How so? If not, why not?

Here's another solution to the problem that uses pointer dereferencing:

```

1 void to_lower(char* p) {
    // assume that lower and upper case characters have consecutive integer values
3     if (p==0) return; // just in case some user tries to_lower(0)
    for (; *p; ++p) { // p is already initialized
5         if ('A' <= *p && *p <= 'Z') // is upper case
                *p -= ('A' - 'a'); // make lower case
7     }
}
9 void test(const string& ss) // rely on visual inspection

```

```

11 {
12     string s = ss; // take a copy to avoid modifying original
13     cout << s << '\n';
14     char* p = &s[0]; // the characters are stored with a terminating 0
15     to_lower(p);
16     cout << p << '\n';
17 }
18
19 int main()
20 {
21     test("Hello , World!");
22
23     string s; // read examples into std::string rather than C-style string
24             // to avoid the possibility of overflow
25     while (cin>>s && s!="quit") // take input, one word at a time, until "quit"
26         test(s);
27 }

```

(b) Is this an improvement in efficiency, as measured by time/operation? Test this using the `chrono` library, as described in the next problem.

2. Look at PPP20.7.1 and time the insert/delete operations using the `chrono` library as partially illustrated below:

```

1 #include "std_lib_facilities.h"
2 #include <chrono>
3 #include <ctime>
4
5 int main()
6 {
7     std::chrono::time_point<std::chrono::system_clock> start, end;
8
9     int initializer [7] = {1,2,3,4,5,6,7};
10    int* first = initializer;
11    int* last  = initializer+7;
12
13    //To compare, we'll do exactly the same with a vector:
14    start = std::chrono::system_clock::now();
15
16    auto beginning = std::chrono::high_resolution_clock::now();
17    for(int i = 0; i < 100000; ++i)
18    {
19        vector<int> v(first, last);
20        vector<int>::iterator p = v.begin(); // take a vector
21        ++p; ++p; ++p; // point to its 3rd element
22        vector<int>::iterator q = p;
23        ++q; // point to it's 4th element
24
25        p = v.insert(p,99); // p points at the inserted element
26        p = v.erase(p); // p points at the element after the erased
27    one
28    }
29    auto ending = std::chrono::high_resolution_clock::now();
30    end = std::chrono::system_clock::now();
31    std::chrono::duration<double> elapsed_seconds = end-start;
32    std::time_t end_time = std::chrono::system_clock::to_time_t(end);
33
34    std::cout << "finished computation at " << std::ctime(&end_time)
35              << "elapsed time: " << elapsed_seconds.count() << "s\n";

```



3. (PPP20.20) Run a small timing experiment to compare the cost of using vector and list. You can find an explanation of how to time a program in 26.6.1. Generate  $N$  random int values in the range  $[0:N)$ . As each int is generated, insert it into a vector `int[]` (which grows by one element each time). Keep the vector sorted; that is, a value is inserted after every previous value that is less than or equal to the new value and before every previous value that is larger than the new value. Now do the same experiment using a list `int[]` to hold the ints. For which  $N$  is the list faster than the vector? Try to explain your result.
4. Do the drill from Chapter 20.