

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.

ANS: Here's the code we wrote to check this. The first to\_lower() ("mine") was coded by student Dmitri.

```

1 #include <iostream>
2 #include <string>
3 #include <chrono>
4 #include <ctime>
5 #include <cstdio>
6
7 using namespace std;
8 using namespace std::chrono;
9
10 void to_lower_mine(char *p) {
11     while(*p) {
12         if(*p >= 'A' && *p <= 'Z') {
13             *p -= ('A' - 'a');
14         }
15         p++;
16     }
17 }
18
19 /// solution #1 , using subscripting :
20 void to_lower_1 (char* p) {
21     /// assume that lower and upper case characters have consecutive integer
22     values
23     for(int i = 0; p[i]; ++i) {
24         if( 'A' <= p[i] && p[i] <= 'Z') // is upper case ( only characters are
25         considered
26             /// upper case     e . g . not !@#)
27             p[i] -= ('A' - 'a'); // make lower case
28     }
29 }
30
31 /// solution #1.1 , using subscripting and a local variable :
32 void to_lower_2(char* p) {
33     /// assume that lower and upper case characters have consecutive integer
34     values
35     for(int i = 0; char c = p[i]; ++i) {
36         if('A' <= c && c <= 'Z') /// is upper case ( only characters are
37         considered upper
38             /// case     e . g . not !@#)

```

```

33         p[i] -= ('A' - 'a'); // make lower case
34     }
35 }
36 void to_lower_3(char* p) {
37     // assume lower/upper case chars are consecutive integer
38     if(p == 0) return ; // just in case some user tries to lower (0)
39     for (; *p; ++p) { // p is already initialized
40         if('A' <= *p && *p <= 'Z') // is upper case
41             *p -= ('A' - 'a'); // make lower case
42     }
43 }
44 void test(const string& ss) { // rely on visual inspection
45     string s = ss ; // take a copy to avoid modifying original
46     cout << s << '\n';
47     char* p = &s[0]; // the characters are stored with a terminating 0
48     to_lower_3(p);
49     cout << p << '\n';
50 }
51 int main() {
52     string s = "Hello, World!"; // read examples into std :: string rather than
53     // C style string
54     // to avoid the possibility of overflow
55     char *p = &s[0];
56     chrono::time_point<chrono::system_clock> t_1, t_2;
57     chrono::duration<double> elapsed_seconds;
58     t_1 = chrono::system_clock::now();
59     for(int i = 0; i < 100000000; i++) {
60         to_lower_mine(p);
61         s = "Hello, World";
62     }
63     t_2 = chrono::system_clock::now();
64     elapsed_seconds = t_2 - t_1;
65     cout << "\nto_lower_mine() elapsed time: "
66         << elapsed_seconds.count();
67     t_1 = chrono::system_clock::now();
68     for(int i = 0; i < 100000000; i++) {
69         to_lower_1(p);
70         s = "Hello, World";
71     }
72     t_2 = chrono::system_clock::now();
73     elapsed_seconds = t_2 - t_1;
74     cout << "\nto_lower_1() elapsed time: " << elapsed_seconds.count();
75     t_1 = chrono::system_clock::now();
76     for(int i = 0; i < 100000000; i++) {
77         to_lower_2(p);
78         s = "Hello, World";
79     }
80     t_2 = chrono::system_clock::now();
81     elapsed_seconds = t_2 - t_1;
82     cout << "\nto_lower_2() elapsed time: " << elapsed_seconds.count();
83     t_1 = chrono::system_clock::now();
84     for(int i = 0; i < 100000000; i++) {
85         to_lower_3(p);
86         s = "Hello, World";
87     }
88     t_2 = chrono::system_clock::now();
89     elapsed_seconds = t_2 - t_1;
90     cout << "\nto_lower_3() elapsed time: " << elapsed_seconds.count();
91     return 0;
92 }

```

...and here's the (very much machine/compiler dependent) output I'm getting now:

```
to_lower_mine() elapsed time: 12.964
to_lower_1() elapsed time: 13.059
to_lower_2() elapsed time: 13.192
to_lower_3() elapsed time: 13.073
```

The differences are in how the loops are formed, what the conditional for the loop is, how the "to\_upper()" action is executed and how the update is performed. Dmitri's to\_lower() uses a while loop conditioned on the NUL character

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"
  #include <chrono>
3 #include <ctime>

5 int main()
  {
7     std::chrono::time_point<std::chrono::system_clock> start, end;

9     int initializer [7] = {1,2,3,4,5,6,7};
    int* first = initializer;
11    int* last = initializer+7;

13    //To compare, we'll do exactly the same with a vector:
    start = std::chrono::system_clock::now();

15
17    auto beginning = std::chrono::high_resolution_clock::now();
    for(int i = 0; i < 100000; ++i)
    {
19        vector<int> v(first, last);
        vector<int>::iterator p = v.begin(); // take a vector
21        ++p; ++p; ++p; // point to its 3rd element
        vector<int>::iterator q = p;
23        ++q; // point to it's 4th element

25        p = v.insert(p,99); // p points at the inserted element
        p = v.erase(p); // p points at the element after the erased
one
    }
27    auto ending = std::chrono::high_resolution_clock::now();
29    end = std::chrono::system_clock::now();
    std::chrono::duration<double> elapsed_seconds = end-start;
31    std::time_t end_time = std::chrono::system_clock::to_time_t(end);

33    std::cout << "finished computation at " << std::ctime(&end_time)
                << "elapsed time: " << elapsed_seconds.count() << "s\n";
35    cout << "My poorly-written function took: "
           << std::chrono::duration_cast<std::chrono::nanoseconds>(ending-beginning).
count()
37           << " nanoseconds\n";
    //To compare, we'll do exactly the same with a list (provide timing code):
39    {
        list<int> v(first, last);
41        list<int>::iterator p = v.begin(); // take a list
        ++p; ++p; ++p; // point to its 3rd element
```

```
43     list<int>::iterator q = p;
44     ++q; // point to it's 4th element
45 }
}
```

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.

ANS: In class we developed code for this purpose:

```

1 #include <iostream>
  #include<cstdlib>
3 #include<ctime>
  using namespace std;
5 using namespace std::chrono;

7 int main() {
  srand(time(0));
9  chrono::time_point<chrono::system_clock> t1, t2;
  chrono::duration<double> elapsed_seconds;
11  ///putIn();
  int randNo{0};
13  vector<int> vint;
  vector<int>::iterator vit = vint.begin();
15  list<int> lint;
  list<int>::iterator lit = lint.begin();
17  int N = 100;
  int n = 50000; /// repeat do_something() n times
19  t1 = system_clock::now(); /// begin time
  vint.insert(vit, rand()%N);
21  //cout << "\nvint[0] = " << vint[0];
  for (int i = 0; i<n; i++) { /// timing loop
23    randNo = rand()%N;
    vit = vint.begin();
25    //cout << "\n*vit = " << *vit;
    while(*vit<randNo && !(vit==vint.end())) ++vit;
27    vint.insert(vit, randNo);
    ///for(j:vint) cout << j << " ";
29  }
  t2 = system_clock::now(); // end time
31  elapsed_seconds = t2-t1;
  cout << "\nInsterting into sorted vector " << n << " times took "
33    << elapsed_seconds.count() << " seconds\n";
  ///<< duration_cast<seconds>(t2-t1).count() << "seconds\n";
35  t1 = system_clock::now(); /// begin time
  //vint.insert(vit, rand()%N);
  //cout << "\nvint[0] = " << vint[0];
  //cin.get();
37  for (int i = 0; i<n; i++) { /// timing loop
    randNo = rand()%N;
41    lit = lint.begin();
    //cout << "\n*lit = " << *lit;
43    while(*lit<randNo && !(lit==lint.end())) ++lit;
    lint.insert(lit, randNo);
45    ///for(j:vint) cout << j << " ";
  }
47  t2 = system_clock::now(); // end time
  elapsed_seconds = t2-t1;
49  cout << "\nInsterting into sorted list " << n << " times took "
    << elapsed_seconds.count() << " seconds\n";
51 }

```

Perhaps it's strange, but, while inserting into `vector` is slower for smaller  $N$ , for larger  $N$ , inserting into `list` is slower! I put the thing in a loop with  $N =$  powers of 10:

$N = 10$ :

Inserting into sorted vector 50000 times took 7.402 seconds

Inserting into sorted list 50000 times took 6.492 seconds

$N = 100$ :

Inserting into sorted vector 50000 times took 30.364 seconds

Inserting into sorted list 50000 times took 35.798 seconds

$N = 1000$ :

Inserting into sorted vector 50000 times took 69.311 seconds

Inserting into sorted list 50000 times took 72.806 seconds

$N = 10000$ :

Inserting into sorted vector 50000 times took 98.779 seconds

Inserting into sorted list 50000 times took 112.134 seconds

For an explanation and bigger picture on this, see `cpp-benchmark-vector-list-deque` “<https://baptistewicht.com/posts/2012/12/cpp-benchmark-vector-list-deque.html>”

“List is clearly slower than the other two data structures that exhibit the same performance. This comes from the very slow linear search. Even if the two other data structures have to move a lot of data, the copy is cheap for small data types.”

4. Do the drill from Chapter 20.