

Programming Assignment #2

Tests due: 3/12/23 @ 11:59pm

Implementation due: 3/19/23 @ 11:59pm

Assignment Link: https://classroom.github.com/a/FBEYM_ry

Please read through the entire writeup before beginning the programming assignment

Objectives

1. Implement a data structure based on linked lists
 - a. Gain a better understanding of constant vs linear runtime
 - b. See an example of a tactical optimization that improves constant factors
 - c. Be able to compare access-by-reference vs access-by-index

Useful Links

1. [The Scala API](#)
 - a. [Scala Collections](#)
 - b. [scala.collections.mutable.Map](#)
2. [Scala Tour](#)
3. [Scala Resources](#)
4. [ScalaTest: Writing your first test](#)
5. [Maps in Scala](#)
6. [Scala File I/O \(Scala Cookbook Excerpt\)](#)
7. [Scala Exercises](#)

Submission Process, Late Policy and Grading

Testing due date: 3/12/23 @ 11:59PM

Implementation due date: 3/19/23 @ 11:59PM

Total points: 30 (5 for testing + 20 for implementation + 5 for runtime)

The project grade is the grade assigned to the latest (most recent) submission made to Autolab (or 0 if no submissions are made). Autolab will pull your submission from your GitHub repository, so you must make sure that any changes you want to be included in your grade have been **committed and pushed**.

- If your submission is made before the deadline, you will be awarded 100% of the points your project earns.
- If your submission is made up to 24 hours after the deadline, you will be awarded 75% of the points your project earns.
- If your submission is more than 24 hours after the deadline, but within 48 hours of the deadline, you will be awarded 50% of the points your project earns.
- If your submission is made more than 48 hours after the deadline, it will not be accepted.

You will have the ability to use three grace days throughout the semester, and at most two per assignment (since submissions are not accepted after two days). Using a grace day will negate the 25% penalty per day, but will not allow you to submit more than two days late. Please plan accordingly. You will not be able to recover a grace day if you decide to work late and your score is not sufficiently higher. Grace days are automatically applied to the first instances of late submissions, and are non-refundable. For example, if an assignment is due on a Friday and you make a submission on Saturday, you will automatically use a grace day, regardless of whether you perform better or not. Be sure to test your code before submitting, especially with late submissions in order to avoid wasting grace days.

Keep track of the time if you are working up until the deadline. Submissions become late after the set deadline. Keep in mind that submissions will close 48 hours after the original deadline and you will not be able to submit your code after that time.

Note: No late submissions will be accepted for the testing portion of the assignment, and no grace days can be used on the testing portion of the assignment.

Setup

In order to complete this project, you must have completed PA0. If you are working on a machine other than the one you used in PA0, you must at least complete steps 2 and 4 in order to get IntelliJ and GitHub working properly.

Once you have ensured your development environment is setup as in PA0, you can accept the PA2 assignment in GitHub Classroom ([here](#)), and create a new IntelliJ project from VCS with your newly created repository.

Instructions

In this PA you will be filling in the body of 11 functions to complete the implementation of a sorted linked list data structure. Additionally, you will be required to write tests which can verify the correctness of valid implementations, and detect bugs in incorrect implementations.

Hint: Yes...there are 11 functions to write, but if you plan before you code, you will realize that only around ~4 of them actually required detailed implementation, and the other functions can be largely implemented by calling those 4 functions with only slight modification.

The linked list you will be implementing in this PA differs from the vanilla Linked Lists discussed in class so far in 3 distinct ways:

- **Ordering:** The data in the linked list must be stored in sorted order. This affects how items are inserted, as well as how items are searched for.
- **Hints:** Many of the functions have versions which take a node in the list as a hint. This hint provides those functions with a (hopefully) better location to start their search from.
- **Duplicates:** Rather than having a distinct linked list node for every insertion into the list, linked list nodes now store a count in addition to the value. If there are multiple insertions of the same value into the list, the count of the associated linked list node is incremented instead of creating a new node.

After you complete your tests, make sure to commit **and push** your work to GitHub, and submit to the PA2 testing submission in Autolab. After completing your implementation, make sure to commit **and push** your work to GitHub, and submit to the PA2 implementation submission in Autolab.

Hint: It is advised that you commit and push frequently rather than waiting until you've completed everything.

Hint: Although you will get feedback from Autolab about correctness of your solutions, you should get in the habit of testing locally, and adding test cases as needed. This will be a more effective/efficient means of development, and will also give you a better understanding of the content of this programming assignment in the process.

0. Testing Phase (due 3/12/23)

[5/30 points]

For this part, modify the file `SortedListTests.scala` to include new test cases.

Your test cases will first be run against a correct implementation of `SortedList`. If your tests fail the correct implementation you will receive 0 points for the testing phase.

Your test cases will then be run against several broken implementations of `SortedList`. You will get points for each broken implementation that at least one of your tests fail.

1. Define the search functions

[5/30 points]

Implement the following functions related to searching for a particular value in our linked list:

```
def findRefBefore(elem: T): Option[SortedListNode[T]]
```

If the list contains `elem`, then this function should return a reference to that list node.

If the list does not contain `elem`, then this function should return a reference to the node that would immediately precede a node with value `elem` if it were to be inserted into the list.

If `elem` is smaller than every element in the list (ie the `elem` would be inserted at the head of the list) then this function should return `None`.

For a list of length n , this function should run in $O(n)$.

```
def findRefBefore(elem:T, hint: SortedListNode[T]): Option[SortedListNode[T]]
```

The return value of this function should be the same as the return value of `findRefBefore(elem: T)`.

The second argument to this function, `hint`, is a reference to a node in the list. The search for the node containing `elem` should begin at `hint` rather than the head of the list.

If `hint` is at position i and the node with value `elem` is at position j , then this function should run in $O(|i - j|)$.

```
def findRef(elem: T): Option[SortedListNode[T]]
```

If the list contains `elem`, then this function should return a reference to that list node.

Otherwise this function should return `None`.

For a list of length n , this function should run in $O(n)$.

```
def findRef(elem: T, hint: SortedListNode[T]): Option[SortedListNode[T]]
```

The return value of this function should be the same as the return value of `findRef(elem: T)`.

The second argument to this function, `hint`, is a reference to a node in the list. The search for the node containing `elem` should begin at `hint` rather than the head of the list.

If `hint` is at position i and the node with value `elem` is at position j , then this function should run in $O(|i - j|)$.

Hint: Only one of these functions really needs a detailed implementation if you then use that function as the basis for defining the other 3.

2. Define the indexing functions [5/30 points]

Implement the following functions related to getting an element at a specific position in the list:

```
def getRef(idx: Int): SortedListNode[T]
```

Return a reference to the node containing the element at the provided index. Note that the index provided is with respect to the elements that have been inserted, not the actual nodes of the list since there may be fewer nodes than elements.

Throw an `IndexOutOfBoundsException` if the provided index is invalid (i.e., `idx < 0` or `idx >= length`).

For a list of length n , this function should run in $O(n)$.

```
def apply(idx: Int): T
```

Return the value at the provided index. Note that the index provided is with respect to the elements that have been inserted, not the actual nodes of the list since there may be fewer nodes than elements.

Throw an `IndexOutOfBoundsException` if the provided index is invalid (i.e., `idx < 0` or `idx >= length`).

For a list of length n , this function should run in $O(n)$.

Hint: Only one of these functions really needs a detailed implementation if you then use that function as the basis for defining the other one.

3. Define the insertion functions [5/30 points]

Implement the following functions related to inserting an element into the list:

```
def insert(elem: T): SortedListNode[T]
```

Insert `elem` into the linked list and return the linked list node where it was inserted.

The element should be placed so that the list remains in sorted order. After the insertion, the returned node's `next` method should return a reference to the next greatest element, and the `prev` method should return a reference to the next least element.

If `elem` is already in the list, no new node should be created. Instead, increment the count of the existing node.

For a list of length n , this function should run in $O(n)$.

```
def insert(elem: T, hint): SortedListNode[T]
```

Insert `elem` into the linked list and return the linked list node where it was inserted.

This function should behave exactly like `insert(elem: T)` but with the search for the insertion point beginning at `hint` rather than the head of the list.

If `hint` is at position i and the node with value `elem` should be at position j , then this function should run in $O(|i - j|)$.

Hint: Much of the logic for insertion can rely on previously defined functions.

4. Define the removal functions

[5/30 points]

Implement the following functions related to removing an element from the list:

```
def remove(ref: SortedListNode[T]): T
```

Remove a single instance of the element held by the node referenced by `ref` from the list, and return its value.

If this is the last element of its value, then the node referenced by `ref` should also be removed from the linked list.

This function should run in $O(1)$.

```
def removeN(ref: SortedListNode[T], n: Int): T
```

Remove `n` instances of the element held by the node referenced by `ref` from the list, and return its value.

If after removing `n` elements, there are no more elements of the referenced value, then the node referenced by `ref` should also be removed from the linked list.

Throw an `IllegalArgumentException` if `n` is larger than the number of elements held by `ref`.

This function should run in $O(1)$.

```
def removeAll(ref: SortedListNode[T]): T
```

Remove all instances of the element held by the node referenced by `ref` from the list, return its value, and remove the node from the linked list.

This function should run in $O(1)$.

Hint: Only one of these functions really needs a detailed implementation if you then use that function as the basis for defining the other 2.

5. Runtime complexity

[5/30 points]

The final 5 points for this assignment will come after the submission deadline from separate testing that ensures your functions have the desired runtime complexity.

Additional Notes

Doubly Linked Lists

A doubly linked list is a collection data structure that stores each of its elements wrapped in a linked list node (`SortedLinkedListNode` for this project).

A doubly linked list contains a `headNode` reference to the first node in the list and a `lastNode` reference to the last node in the list. Each node contains the value it stores, and references to the following and preceding nodes (if they exist...see the section on `Option` types in Scala). In this assignment, each node also contains a count that defines how many times the associated value has been inserted into the list.

To insert an element into a doubly-linked list, all of the relevant references need to be updated.

Make sure you handle all of the corner cases!

1. `newNode.next` and `newNode.prev` need to be set appropriately
2. `newNode.prev.next` needs to be updated if it exists
3. `newNode.next.prev` needs to be updated if it exists
4. `headNode` needs to be updated if `newNode` is inserted at the front
5. `lastNode` needs to be updated if `newNode` is inserted at the end
6. `SortedLinkedList.length` and `SortedLinkedListNode.count` should be updated

Ordering in Scala

`SortedLinkedList` uses a scala provided class called `Ordering` to define a sort order. This class defines a method: `compare(a, b)` that behaves as follows:

- `compare(a, b) < 0` if `a < b`
- `compare(a, b) == 0` if `a = b`
- `compare(a, b) > 0` if `a > b`

There is a little scala magic involved in how you call `Ordering.compare`, so

`SortedLinkedList` defines a helper function, `compare`, to call it for you. In general, the following conditions should hold for all nodes of the linked list, excluding the first/final one:

```
compare(node.prev.get.value, node.value) <= 0
compare(node.value, node.next.get.value) <= 0
```

Operations that modify the contents of the linked list (`insert` and `update`) should ensure that the inserted/modified node is moved to the correct location.

Option Types in Scala

Because in general there is no guarantee that a `SortedLinkedListNode` will have a following or preceding node, we cannot store references to `SortedLinkedListNodes` directly. Similarly, for many of the `SortedLinkedList` functions, there is no guarantee that a given node will be found, so we cannot return a reference to a `SortedLinkedListNode` directly.

In both cases, we use the [Option](#) class in Scala. The `Option[T]` class has two subclasses, `Some[T]` and `None`. `Some[T]` is used in a situation where we actually have something of type `T`, and `None` is used when we do not.

If for example, we have a function that may or may not return an `Int`, then in Scala we would declare the return type of that function to be `Option[Int]`. To return an integer, say 2 for example, our function would return `Some(2)`. To return nothing our function would return `None`.

Other Tips

1. It is important to understand that the mental model we have of our collection, and its actual implementation are different. For example, if the linked list in this assignment has values 1, 1, 1, 2, 3 inserted into it, the sequence it represents is 1, 1, 1, 2, 3. There are 5 elements in that order. The actual implementation of our linked list however would only have 3 actual linked list nodes.
2. The linked list implementation already has the update and iterator methods defined for you. Feel free to reference them as needed.
3. When writing tests, it is useful to start with tests that test one function at a time, before moving on to tests that test the interaction between multiple functions. If you test multiple functions in the same test without being sure they are individually correct, then if the test fails it is hard to know where to begin. For example, if you are trying to test the remove function, but you create the linked list that you will test it on by calling insert, then if the test fails you don't actually know if the remove call was incorrect, or the insert call was incorrect. Building your example lists manually (by explicitly creating the nodes and setting the next/prev refs by hand) you can have more focused testing.
4. For testing, try to think about what corner cases or oddities would trip you up when implementing the list yourself. Think about details which could be easily overlooked. Then write tests which check for these corner cases and details. In order to do this, you must have a good understanding of the problem you are trying to solve. **Make sure you understand exactly how the linked list in this assignment is supposed to behave before writing tests.** Draw out examples on paper, and ask clarifying questions as needed.

Academic Integrity

As a gentle reminder, please re-read the academic integrity policy of the course. I will continue to remind you throughout the semester and hope to avoid any incidents.

What Constitutes a Violation of Academic Integrity?

These bullets should be obvious things not to do (but commonly occur):

- Turning in your friend's code/write-up (obvious).
- Turning in solutions you found on Google with all the variable names changed (should be obvious). This is a copyright violation, in addition to an AI violation.
- Turning in solutions you found on Google with all the variable names changed and 2 lines added (should be obvious). This is also a copyright violation.
- Paying someone to do your work. You may as well not submit the work since you will fail the exams and the course.
- Posting to forums asking someone to solve the problem.

Note: Aggregating every [stack overflow answer|result from google|other source] because you "understand it" will likely result in full credit on assignments (if you aren't caught) and then failure on every exam. Exams don't test if you know how to use Google, but rather test your understanding (i.e., can you understand the problems to arrive at a solution on your own). Also, other students are likely doing the same thing and then you will be wondering why 10 people that you don't know have your solution.

Other violations that may not be as obvious:

- Working with a tutor who solves the assignment with you. If you have a tutor, please contact me so that I may discuss with them what help is allowed.
- Sending your code to a friend to help them. If another student uses/submits your code, you are also liable and will be punished.
- Joining a chatroom for the course where someone posts their code once they finish, with the honor code that everyone needs to change it in order to use it.
- Reading your friend's code the night before it is due because you just need one more line to get everything working. It will most likely influence you directly or subconsciously to solve the problem identically, and your friend will also end up in trouble.

What Collaboration is Allowed?

Assignments in this course should be solved individually with only assistance from course staff and allowed resources. You may discuss and help one another with technical issues, such as how to get your compiler running, etc.

There is a gray area when it comes to discussing the problems with your peers and I do encourage you to work with one another to solve problems. That is the best way to learn and overcome obstacles. At the same time you need to be sure you do not overstep and not plagiarize. Talking out how you eventually reached the solution from a high level is okay:

"I used a stack to store the data and then looked for the value to return."

but explaining every step in detail/pseudocode is not okay:

"I copied the file tutorial into my code at the start of the function, then created a stack and pushed all of the data onto the stack, and finished by popping the elements until the value is found and use a return statement."

The first example is OK but the second is basically a summary of your code and is not acceptable, and remember that you shouldn't be showing any code at all for how to do any of it. Regardless of where you are working, you must always follow this rule: Never come away from discussions with your peers with any written work, either typed or photographed, and especially do not share or allow viewing of your written code.

What Resources are Allowed?

With all of this said, please feel free to use any [files|examples|tutorials] that we provide directly in your code (with proper attribution). Feel free to directly use anything from lectures or recitations. You will never be penalized for doing so, but should always provide attribution/citation for where you retrieved code from. Just remember, if you are citing an algorithm that is not provided by us, then you are probably overstepping.

More explicitly, you may use any of the following resources (with proper citation/attribution):

- Any example files posted on the course webpage (from lecture or recitation).
- Any code that the instructor provides.
- Any code that the TAs provide.
- Any code from the tour of Scala (<https://docs.scala-lang.org/tour/tour-of-scala.html>)
- Any code from Scala Collections (<https://docs.scala-lang.org/overviews/collections-2.13/introduction.html>)
- Any code from Scala API (<https://www.scala-lang.org/api/2.13.0/>)

Omitting citation/attribution will result in an AI violation (and lawsuits later in life at your job). This is true even if you are using resources provided.

Amnesty Policy

We understand that students are under a lot of pressure and people make mistakes. If you have concerns that you may have violated academic integrity on a particular assignment, and would like to withdraw the assignment, you may do so by sending me an email BEFORE THE VIOLATION IS DISCOVERED BY ME. The email should take the following format:

Dear Dr. Mikida,

I wish to inform you that on assignment X, the work I submitted was not entirely my own. I would like to withdraw my submission from consideration to preserve academic integrity.

J.Q. Student
Person #12345678
UBIT: jqstuden

When we receive this email, student J would receive a 0 on assignment X, but would not receive an F for the course, and would not be reported to the office of academic integrity.