

HW2: Three Dictionaries

Due: Thursday, October 18, at 11:59 PM, via GSC

You may work on your own or with one (1) partner.

The *dictionary* is a common abstract data type for representing a table of key-value pairs. For example, suppose you wanted a data structure for keeping track of your friends' birthdays:

Key	Value
Anne	Aug. 27, 1997
Bob	Mar. 5, 1997
Carol	Dec. 29, 1996

To represent such a table, you could use a dictionary with names as keys and birthdays as values. In this assignment, we consider dictionaries as bank ledgers mapping account numbers to account information.

There are many possible concrete implementations of the dictionary ADT, and in this assignment you will implement operations for three:

- a linked list,
- a binary search tree, and
- a sorted vector.

In `dictionaries.rkt`¹ I've supplied headers for the functions that you'll need to write along with some tests.

Your task

The Account struct

In Homework 1, we used an `Account` class, but to keep things simple in Homework 2, we use a plain `Account` struct:

```
struct Account:  
  let id: nat?  
  let owner: str?  
  let balance: num?
```

The annotations after the colons are contracts—predicates, in this case, that must hold for each field. That is, the `id` field must be a natural number, the `owner` field must be a string, and the `balance` field can be any kind of number.

¹<http://goo.gl/ZUuKwp>

The starter code also defines six example accounts, a function for cloning accounts (so that you don't modify the example accounts in testing), and an `account_transfer` function (used later).

The LEDGER interface

A ledger is, for our purposes, a collection of `Accounts` that supports efficient lookup by account ID. That is, it's a kind of read-only dictionary where the keys are account IDs and the values are the `Accounts` themselves.

We are going to represent a ledger with three different concrete data structures, so we will express the ledger abstract data type (ADT) as an interface:

```
interface LEDGER:
    # How many accounts are in this ledger?
    def len(self) -> nat?
    # Looks up an account by ID, returning False if not found.
    def lookup(self, id: nat?) -> OrC(Account?, False)
```

You will implement operations for three classes that satisfy the `LEDGER` interface.

The linked-list representation

The linked list data definition is as follows. First, the list itself is a singly-linked list of nodes:

```
# A BareList is one of:
# - nil()
# - node(Account, BareList)
# where the `account.id` values are unique.
struct nil:
    pass

struct node:
    let element: Account?
    let link: OrC(nil?, node?)
```

As with the `Account` struct, the `node` struct uses contract to express constraints on its fields: The `element` field must be an `Account` struct, and the `link` field can be either a `nil` (the end of the list) or another `node`.

We encapsulate the singly-linked list in a `ListLedger` class, whose one field, `contents`, is the `BareList` as described above.

```
class ListLedger (LEDGER):
    let contents: OrC(nil?, node?)

    def __init__(self):
```

```
self.contents = nil()
```

```
...
```

In addition to the constructor for `ListLedger`, the starter code contains a definition of a `push_back` method, which adds an account to the end of the list. Your job is to implement four additional methods:

1. `ListLedger.len(self) -> nat?` returns the number of accounts in the list, or equivalently, the number of nodes.
2. `ListLedger.push_front(self, account: Account?) -> VoidC` adds an account to the front of the list.
3. `ListLedger.pop_front(self) -> Account?` removes the first node from the list and returns its account, or errors on empty.
4. `ListLedger.lookup(self, id: nat?) -> OrC(Account?, False)` looks up an account by ID, returning `False` if no `Account` with the given ID can be found.

Note that the accounts are stored in no particular order.

A small number of tests are provided, but you should write more.

The binary search tree representation

```
# A BareBst is one of
# - leaf()
# - branch(BareBst, Account, BareBst)
# where for a branch branch(l, acct, r), all the account.ids
# l are less than acct.id, and all the account.ids in r are greater
# than acct.id. (This is the binary search tree property.)

struct leaf:
  pass

struct branch:
  let left: OrC(leaf?, branch?)
  let element: Account?
  let right: OrC(leaf?, branch?)
```

That is, a `BareBst` is either the empty tree, or a tree node containing one account and two subtrees with (potentially) more accounts. The accounts are stored in order of increasing account number à la the binary search tree property. This means that lookups needs only proceed down a single path in the tree rather than searching everywhere, for time complexity $\mathcal{O}(\log n)$ (when the tree is balanced).

As with `ListLedger` above, the tree nodes are encapsulated in a class, `BstLedger`, which has one field, which contains the root of the tree:

```
class BstLedger (LEDGER):
    let contents: OrC(leaf?, branch?)

    def __init__(self):
        self.contents = leaf()

    ...
```

In addition to the constructor for `BstLedger`, the starter code contains a definition of an `insert` method, which adds an account to bottom of the tree. The `insert` method does not necessarily produce a balanced tree—we’ll learn that later—but it does produce a predictable tree, which can be useful for testing.

Your job is to implement two additional methods:

5. `BstLedger.len(self) -> nat?` returns the number of accounts in the tree, or equivalently, the number of nodes.
6. `BstLedger.lookup(self, id: nat?) -> OrC(Account?, False)` looks up an account by ID, returning `False` if no `Account` with the given ID can be found. This method must use the BST property to find the element rather than searching the whole tree.

A small number of tests are provided, but you should write more.

The sorted vector representation

The third class is a sorted vector. Unlike the linked list and binary search tree above, there are no auxiliary structs to define. The `VecLedger` class has one field, which is a vector of accounts, sorted by ascending ID:

```
class VecLedger (LEDGER):
    let contents: VecC[Account?]

    def __init__(self, vec: VecC[Account?]):
        self._assert_sorted(vec)
        self.contents = vec

    ...
```

Rather than provide a method for adding accounts to the vector, the constructor for `VecLedger` takes an already-sorted vector of accounts. Thus, the caller is responsible for supplying such a thing. Before initializing the `contents` field, the constructor calls a private helper method, `VecLedger._assert_sorted`, to check that the vector meets the class’s invariant.

In addition to the constructor for `VecLedger`, the starter code contains a definition of a `len` method. Your job is to implement two additional methods:

7. `VecLedger._assert_sorted(self, vec)` scans vector `vec`, checking that the account IDs are strictly increasing; it calls `error` if it discovers that they are not.
8. `VecLedger.lookup(self, id: nat?) -> OrC(Account?, False)` looks up an account by ID, returning `False` if no `Account` with the given ID can be found. This method must take advantage of the invariant that the `contents` vector is sorted by account ID. This means that lookups proceed by binary search, not by a linear scan of the vector, for time complexity $\mathcal{O}(\log n)$.

A small number of tests are provided, but you should write more.

Working with the LEDGER ADT

Because all three ledger classes implement the `LEDGER` interface, we can write code that works with instances of all three of them. The starter code provides a function header:

```
def ledger_transfer(amount: num?,
                   from_id: nat?,
                   to_id: nat?,
                   ledger: LEDGER!):
```

This function should look up the accounts specified by `from_id` and `to_id` in `ledger`, and transfer `amount` from `from_id` and to `to_id`. The contract annotation `LEDGER!` limits `ledger` to be used only according to the `LEDGER` interface, which is all that `ledger_transfer` needs in order to do its work. It should look up both accounts, calling `error` if either does not exist. Then it should perform the transfer with the provided `account_transfer` function.

9. Complete the definition of the `ledger_transfer` function.

Deliverables

The provided file `dictionaries.rkt`, containing

- definitions of the eight methods and one function described above, and
- sufficient tests to be confident of your code's correctness.