

Data Structures and Algorithms

RAJASTHAN TECHNICAL UNIVERSITY, KOTA
Syllabus II Year-III Semester: B.Tech. Computer Science and Engineering
3CS4-05: Data Structures and Algorithms

Credit-3 Max.

Marks: 150 (IA:30, ETE:120)

1 Stacks: Basic Stack Operations, Representation of a Stack using Static Array and Dynamic Array, Multiple stack implementation using single array, Stack Applications: Reversing list, Factorial Calculation, Infix to postfix Transformation, Evaluating Arithmetic Expressions and Towers of Hanoi.

2 Queues: Basic Queue Operations, Representation of a Queue using array, Implementation of Queue Operations using Stack, Applications of Queues- Round Robin Algorithm. Circular Queues, DeQueue Priority Queues. Linked Lists: Introduction, single linked list, representation of a linked list in memory, Different Operations on a Single linked list, Reversing a single linked list, Advantages and disadvantages of single linked list, circular linked list, double linked list and Header linked list.

3 Searching Techniques: Sequential and binary search. Sorting Techniques: Basic concepts, Sorting by: bubble sort, Insertion sort, selection sort, quick sort, heap sort, merge sort, radix sort and counting sorting algorithms.

4 Trees: Definition of tree, Properties of tree, Binary Tree, Representation of Binary trees using arrays and linked lists, Operations on a Binary Tree, Binary Tree Traversals (recursive), Binary search tree, B-tree, B+ tree, AVL tree, Threaded binary tree.

5 Graphs: Basic concepts, Different representations of Graphs, Graph Traversals (BFS & DFS), Minimum Spanning Tree(Prims &Kruskal), Dijkstra's shortest path algorithms. Hashing: Hash function, Address calculation techniques, Common hashing functions, Collision resolution: Linear and Quadratic probing, Double hashing.

1: Introduction

What Are Data Structures and Algorithms?

If this book is about data structures and algorithms, then perhaps we should start by defining these terms.. We begin with a definition for “algorithm..”

Algorithm: A finite sequence of steps for accomplishing some computational task.. An algorithm must

- Have steps that are simple and definite enough to be done by a computer, and
- Terminate after finitely many steps..

This definition of an algorithm is similar to others you may have seen in prior computer science courses.. Notice that an algorithm is a sequence of steps, not a program.. You might use the same algorithm in different programs, or express the same algorithm in different languages, because an algorithm is an entity that is abstracted from implementation details.. Part of the point of this course is to introduce you to algorithms that you can use no matter what language you program in.. We will write programs in a particular language, but what we are really studying is the algorithms, not their implementations..

The definition of a data structure is a bit more involved.. We begin with the notion of an abstract data type..

Abstract data type (ADT): A set of values (the **carrier set**), and operations on those values..

Here are some examples of ADTs:

Boolean—The carrier set of the Boolean ADT is the set { true, false }.. The operations on these values are negation, conjunction, disjunction, conditional, is equal to, and perhaps some others..

Integer—The carrier set of the Integer ADT is the set {, -2, -1, 0, 1, 2, }, and the operations on these values are addition, subtraction, multiplication, division, remainder, is equal to, is less than, is greater than, and so on.. Note that although some of these operations yield other Integer values, some yield values from other ADTs (like true and false), but all have at least one Integer value argument..

String—The carrier set of the String ADT is the set of all finite sequences of characters from some alphabet, including the empty sequence (the empty string).. Operations on string values include concatenation, length of, substring, index of, and so forth..

Bit String—The carrier set of the Bit String ADT is the set of all finite sequences of bits, including the empty strings of bits, which we denote λ : { λ , 0, 1, 00, 01, 10, 11, 000, }.. Operations on bit strings include complement (which reverses all the bits), shifts (which rotates a bit string left or right), conjunction and disjunction (which combine bits at corresponding locations in the strings, and concatenation and truncation..

The thing that makes an abstract data type *abstract* is that its carrier set and its operations are mathematical entities, like numbers or geometric objects; all details of implementation on a computer are ignored.. This makes it easier to reason about them and to understand what they are.. For example, we can decide how `div` and `mod` should work for negative numbers in the Integer ADT without having to worry about how to make this work on real computers.. Then we can deal with implementation of our decisions as a separate problem.. Once an abstract data type is implemented on a computer, we call it a data type..

Data type: An implementation of an abstract data type on a computer..

Thus, for example, the Boolean ADT is implemented as the `boolean` type in Java, and the `bool` type in C++; the Integer ADT is realized as the `int` and `long` types in Java, and the `Integer` class in Ruby; the String ADT is implemented as the `String` class in Java and Ruby..

Abstract data types are very useful for helping us understand the mathematical objects that we use in our computations, but, of course, we cannot use them directly in our programs.. To use ADTs in programming, we must figure out how to implement them on a computer.. Implementing an ADT requires two things:

- Representing the values in the carrier set of the ADT by data stored in computer memory, and
- Realizing computational mechanisms for the operations of the ADT..

Finding ways to represent carrier set values in a computer's memory requires that we determine how to arrange data (ultimately bits) in memory locations so that each value of the carrier set has a unique representation.. Such things are data structures..

Data structure: An arrangement of data in memory locations to represent values of the carrier set of an abstract data type..

Realizing computational mechanisms for performing operations of the type really means finding algorithms that use the data structures for the carrier set to implement the operations of the ADT.. And now it should be clear why we study data structures and algorithms together: to implement an ADT, we must find data structures to represent the values of its carrier set and algorithms to work with these data structures to implement its operations..

A course in data structures and algorithms is thus a course in implementing abstract data types.. It may seem that we are paying a lot of attention to a minor topic, but abstract data types are really the foundation of everything we do in computing.. Our computations work on data.. This data must represent things and be manipulated according to rules.. These things and the rules for their manipulation amount to abstract data types..

Usually there are many ways to implement an ADT.. A large part of the study of data structures and algorithms is learning about alternative ways to implement an ADT and evaluating the alternatives to determine their advantages and disadvantages.. Typically some alternatives will be better for certain applications and other alternatives will be better for

other applications.. Knowing how to do such evaluations to make good design decisions is an essential part of becoming an expert programmer..

{ **DATA STRUCTURE:** -Structural representation of data items in primary memory to do storage & retrieval operations efficiently.}

--FILE STRUCTURE: Representation of items in secondary memory.

While designing data structure following perspectives to be looked after.

- i. Application(user) level: Way of modeling real-life data in specific context.
- ii. Abstract(logical) level: Abstract collection of elements & operations.
- iii. Implementation level: Representation of structure in programming language.

Data structures are needed to solve real-world problems. But while choosing implementations for it, its necessary to recognize the efficiency in terms of TIME and SPACE.

TYPES:

- i. Simple: built from primitive data types like int, char & Boolean.
eg: Array & Structure
- ii. Compound: Combined in various ways to form complex structures.
1:Linear: Elements share adjacency relationship& form a sequence.
Eg: Stack, Queue , Linked List

2: Non-Linear: Are multi-level data structure. eg: Tree, Graph.

ABSTRACT DATA TYPE :

Specifies the logical properties of data type or data structure.
Refers to the mathematical concept that governs them.

They are not concerned with the implementation details like space and time efficiency.

They are defined by 3 components called Triple =(D,F,A)

D=Set of domain

F=Set of function

A=Set of axioms / rules

Review Questions

1. What would be the carrier set and some operations of the Character ADT?
2. How might the Bit String ADT carrier set be represented on a computer in some high level language?
3. How might the concatenation operation of the Bit String ADT be realized using the carrier set representation you devised for question two above?

1: Introduction

4. What do your answers to questions two and three above have to do with data structures and algorithms?

Exercises

1. Describe the carrier sets and some operations for the following ADTs:
 - (a) The Real numbers
 - (b) The Rational numbers
 - (c) The Complex numbers
 - (d) Ordered pairs of Integers
 - (e) Sets of Characters
 - (f) Grades (the letters A, B, C, D, and F)
2. For each of the ADTs in exercise one, either indicate how the ADT is realized in some programming language, or describe how the values in the carrier set might be realized using the facilities of some programming language, and sketch how the operations of the ADT might be implemented..

Review Question Answers

1. We must first choose a character set; suppose we use the ASCII characters.. Then the carrier set of the Character ADT is the set of the ASCII characters.. Some operations of this ADT might be those to change character case from lower to upper and the reverse, classification operations to determine whether a character is a letter, a digit, whitespace, punctuation, a printable character, and so forth, and operations to convert between integers and characters..
2. Bit String ADT values could be represented in many ways.. For example, bit strings might be represented in character strings of “0”s and “1”s.. They might be represented by arrays or lists of characters, Booleans, or integers..
3. If bit strings are represented as characters strings, then the bit string concatenation operation is realized by the character strings concatenation operation.. If bit strings are represented by arrays or lists, then the concatenation of two bit strings is a new array or list whose size is the sum of the sizes of the argument data structures consisting of the bits from the first bit string copied into the initial portion of the result array or list, followed by the bits from the second bit string copied into the remaining portion..
4. The carrier set representations described in the answer to question two are data structures, and the implementations of the concatenation operation described in the answer to question three are (sketches of) algorithms..

LINKED LIST:

A dynamic data structure.

Linear collection of data items.

Direction is associated with it.


Logical link exists b/w items. Pointers acts as the logical link.

Consists of nodes that has two fields.

- Data field : info of the element.
- Next field: next pointer containing the address of next node.

TYPES OF LINKED LIST:

i. Singly or chain: Single link b/w items.



ii. Doubly: There are two links, forward and backward link.



iii. Circular: The last node is again linked to the first node. These can be singly circular & doubly circular list.

ADVANTAGES:

Linked list use dynamic memory allocation thus allocating memory when program is initialised. List can grow and shrink as needed. Arrays follow static memory allocation .Hence there is wastage of space when less elements are declared. There is possibility of overflow too bcoz of fixed amount of storage.

Nodes are stored incontinuously thus insertion and deletion operations are easily implemented.

Linear data structures like stack and queues are easily implemented using linked list.

DISADVANTAGES:

Wastage of memory as pointers require extra storage.

Nodes are incontinuously stored thereby increasing time required to access individual elements. To access nth item arrays need a single operation while linked list need to pass through (n-1) items.

Nodes must be read in order from beginning as they have inherent sequential access.

Reverse traversing is difficult especially in singly linked list. Memory is wasted for allocating space for back pointers in doubly linked list.

DEFINING LINKED LIST:

```
struct node {  
  
int info;  
  
    struct node *next; \\next field. An eg of self referencetial structure.(#)  
  
} *ptr;
```

(#)Self Referencetial structure: A structure that is referencing to another structure of same type. Here "next" is pointing to structure of type "node".

-ptr is a pointer of type node. To access info n next the syntax is: ptr->info; ptr->next;

OPERATIONS ON SINGLY LINKED LIST:

- i. Searching
- ii. Insertion
- iii. Deletion
- iv. Traversal
- v. Reversal
- vi. Splitting
- vii. Concatenation

Some operations:

a: Insertion :

```
void push(struct node** headref, int data)----- (1)  
  
{  
  
    struct node* newnode = malloc(sizeof(struct node));  
  
    newnode->data= data;  
  
    newnode->next= *headref;  
  
    *headref = newnode;  
  
}
```


(1) : headref is a pointer to a pointer of type struct node. Such passing of pointer to pointer is called Reference pointer. Such declarations are similar to declarations of call by reference. When pointers are passed to functions ,the function works with the original copy of the variable.

i. Insertion at head:

```
struct node* head=NULL;
for(int i=1; i<6;i++)
```

```
{ push(&head,i); \\ push is called here. Data pushed is 1.'&' used coz references are
passed in function arguments.
```

```
    }
return(head);
}
```

```
# :o\p: 5 4 3 2 1
```

```
# :Items appear in reverse order.(demerit)
```

ii. Insertion at tail:

```
struct node* add1()
{ struct node* head=NULL;
struct node* tail;
push(&head,1);
tail = head;
for(int i=2 ;i<6; i++)
{ push(&(tail->next),i);
tail= tail->next;
}
return(head);
}
```

```
# : o\p: 1 2 3 4 5
```

b. Traversal:

```
int count( struct node* p)
```

```
{int count =0;
```

```
struct node* q;
```

```
current = q;
```

```
while(q->next != NULL)
```

```
{ q=q->next;
```

```
count++; }
```

```
return(count);
```

```
}
```

c. Searching:

```
struct node* search( struct node* list, int x)
```

```
{ struct node* p;
```

```
for(p= list; p ->next != NULL; p= p->next )
```

```
{
```

```
    if(p->data==x)
```

```
        return(p);
```

```
    return(NULL);
```

```
}}
```

IMPLEMENTATION OF LISTS:

i : Array implementation:

```
#define NUMNODES 100
```

```
structnodetype
```

```
{ int info ,next ;
```

```
};
```

```
structnodetype node[NUMNODES];
```

:100 nodes are declared as an array node. Pointer to a node is represented by an array index. Thus pointer is an integer b/w 0 & NUMNODES-1 . NULL pointer is represented by -1. node[p] is used to reference node(p) , info(p) is referenced by node[p].info & next by node[p].next.

ii : Dynamic Implementation :

This is the same as codes written under defining of linked lists. Using `malloc()` and `free()` there is the capability of dynamically allocating & freeing variable. It is identical to array implementation except that the next field is a pointer rather than an integer.

NOTE : Major demerit of dynamic implementation is that it may be more time consuming to call upon the system to allocate & free storage than to manipulate a programmer- managed list. Major advantage is that a set of nodes is not reserved in advance for use.

Arrays

Introduction

A structured type of fundamental importance in almost every procedural programming language is the array..

Array: A fixed length, ordered collection of values of the same type stored in contiguous memory locations; the collection may be ordered in several dimensions..

The values stored in an array are called **elements**.. Elements are accessed by *indexing* into the array: an integer value is used to indicate the ordinal value of the element.. For example, if a is an array with 20 elements, then $a[6]$ is the element of a with ordinal value 6.. Indexing may start at any number, but generally it starts at 0.. In the example above $a[6]$ is the seventh value in a when indexing start at 0..

Arrays are important because they allow many values to be stored in a single data structure while providing very fast access to each value.. This is made possible by the fact that (a) all values in an array are the same type, and hence require the same amount of memory to store, and that (b) elements are stored in contiguous memory locations.. Accessing element $a[i]$ requires finding the location where the element is stored.. This is done by computing $b + (i \times m)$, where m is the size of an array element, and b is the base location of the array a .. This computation is obviously very fast.. Furthermore, access to all the elements of the array can be done by starting a counter at b and incrementing it by m , thus yielding the location of each element in turn, which is also very fast..

Arrays are not abstract data types because their arrangement in the physical memory of a computer is an essential feature of their definition, and abstract data types abstract from all details of implementation on a computer.. Nonetheless, we can discuss arrays in a “semi-abstract” fashion that abstracts some implementation details.. The definition above abstracts away the details about how elements are stored in contiguous locations (which indeed does vary somewhat among languages).. Also, arrays are typically types in procedural programming languages, so they are treated like realizations of abstract data types even though they are really not..

In this book, we will treat arrays as implementation mechanisms and not as

ADTs.. **Varieties of Arrays**

In some languages, the size of an array must be established once and for all at program design time and cannot change during execution.. Such arrays are called **static arrays**.. A

chunk of memory big enough to hold all the values in the array is allocated when the array is created, and thereafter elements are accessed using the fixed base location of the array.. Static arrays are the fundamental array type in most older procedural languages, such as Fortran, Basic, and C, and in many newer object-oriented languages as well, such as Java.. Some languages provide arrays whose sizes are established at run-time and can change during execution.. These **dynamic arrays** have an initial size used as the basis for allocating a segment of memory for element storage.. Thereafter the array may shrink or grow.. If the array shrinks during execution, then only an initial portion of allocated memory is used.. But if the array grows beyond the space allocated for it, a more complex *reallocation procedure* must occur, as follows:

1. A new segment of memory large enough to store the elements of the expanded array is allocated..
2. All elements of the original (unexpanded) array are copied into the new memory segment..
3. The memory used initially to store array values is freed and the newly allocated memory is associated with the array variable or reference..

This reallocation procedure is computationally expensive, so systems are usually designed to minimize its frequency of use.. For example, when an array expands beyond its memory allocation, its memory allocation might be doubled even if space for only a single additional element is needed.. The hope is that providing a lot of extra space will avoid many expensive reallocation procedures if the array expands slowly over time..

Dynamic arrays are convenient for programmers because they can never be too small—whenever more space is needed in a dynamic array, it can simply be expanded.. One drawback of dynamic arrays is that implementing language support for them is more work for the compiler or interpreter writer.. A potentially more serious drawback is that the expansion procedure is expensive, so there are circumstances when using a dynamic array can be dangerous.. For example, if an application must respond in real time to events in its environment, and a dynamic array must be expanded when the application is in the midst of a response, then the response may be delayed too long, causing problems..

Review Questions

1. If an array holds integers, each of which is four bytes long, how many bytes from the base location of the array is the location of the fifth element?
2. Is the formula for finding the location of an element in a dynamic array different from the formula for finding the location of an element in a static array?
3. When a dynamic array expands, why can't the existing elements be left in place and extra memory simply be allocated at the end of the existing memory allocation?
4. If a Ruby array a has n elements, which element is $a[n-1]$? Which is element $a[-1]$?

Exercises

1. Suppose a dynamic integer array a with indices beginning at 0 has 1000 elements and the line of code $a[1000] = a[5]$ is executed.. How many array values must be moved from one memory location to another to complete this assignment statement?
2. Memory could be freed when a dynamic array shrinks.. What advantages or disadvantages might this have?

3. To use a static array, data must be recorded about the base location of the array, the size of the elements (for indexing), and the number of elements in the array (to check that indexing is within bounds).. What information must be recorded to use a dynamic array?
4. State a formula to determine how far the from base location of a Ruby array an element with index i is when i is a negative number..
5. Give an example of a Ruby array reference that will cause an indexing error at run time..
6. Suppose the Ruby assignment `a=(1 .. 100).to_a` is executed.. What are the values of the following Ruby expressions? Hint: You can check your answers with the Ruby interpreter..
 - (a) `a[5..10]`
 - (b) `a[5...10]`
 - (c) `a[5, 4]`
 - (d) `a[-5, 4]`
 - (e) `a[100..105]`
 - (f) `a[5..-5]`
 - (g) `a[0, 3] + a[-3, 3]`
7. Suppose that the following Ruby statement are executed in order.. What is the value of array `a` after each statement? Hint: You can check your answers with the Ruby interpreter..
 - (a) `a = Array.new(5, 0)`
 - (b) `a[1..2] = []`
 - (c) `a[10] = 10`
 - (d) `a[3, 7] = [1, 2, 3, 4, 5, 6, 7]`
 - (e) `a[0, 2] = 5`
 - (f) `a[0, 2] = 6, 7`
 - (g) `a[0..-2] = (1..3).to_a`

Review Question Answers

1. If an array holds integers, each of which is four bytes long, then the fifth element is 20 bytes past the base location of the array..
2. The formula for finding the location of an element in a dynamic array is the same as the formula for finding the location of an element in a static array.. The only difference is what happens when a location is beyond the end of the array.. For a static array, trying to access or assign to an element beyond the end of the array is an indexing error.. For a dynamic array, it may mean that the array needs to be expanded, depending on the language.. In Ruby, for example, accessing the value of `a[i]` when $i \geq a.size$ produces `nil`, while assigning a value to `a[i]` when $i \geq a.size$ causes the array `a` to expand to `size i+1`..
3. The memory allocated for an array almost always has memory allocated for other data structures after it, so it cannot simply be increased without clobbering other data structures.. A new chunk of memory must be allocated from the free memory store sufficient to hold the expanded array, and the old memory returned to the free memory store so it can be used later for other (smaller) data structures..
4. If a Ruby array `a` has n elements, then element `a[n-1]` and element `a[-1]` are both the last element in the array.. In general, `a[n-i]` and `a[-i]` are the same elements..

Stacks

Introduction

Stacks have many physical metaphors: shirts in a drawer, plates in a plate holder, box-cars in a dead-end siding, and so forth.. The essential features of a stack are that it is ordered and that access to it is restricted to one end..

Stack: A dispenser holding a sequence of elements that can be accessed, inserted, or removed at only one end, called the **top**..

Stacks are also called last-in-first-out (LIFO) lists.. Stacks are important in computing because of their applications in recursive processing, such as language parsing, expression evaluation, runtime function call management, and so forth..

The Stack ADT

Stacks are containers, and as such they hold values of some type.. We must therefore speak of the ADT *stack of T*, where *T* is the type of the elements held in the stack.. The carrier set of this type is the set of all stacks holding elements of type *T*.. The carrier set thus includes the empty stack, the stacks with one element of type *T*, the stacks with two elements of type *T*, and so forth.. The essential operations of the type are the following (where *s* is a stack of *T*, and *e* is a *T* value)..

push(s,e)—Return a new stack just like *s* except that *e* has been added at the top of *s*..

pop(s)—Remove the top element of *s* and return the resulting (shorter) stack..

Attempting to pop an empty stack gives an undefined result.. Thus, a precondition of the *pop()* operation is that the stack *s* is not empty..

empty?(s)—Return the Boolean value true just in case *s* is empty..

top(s)—Return the top element of *s* without removing it.. Like *pop()*, this operation has the precondition that the stack *s* is not empty..

Besides the precondition assertions mentioned in the explanation of the stack ADT operations above, we can also state some axioms to help us understand the stack ADT.. For example, consider the following axioms..

For any stack *s* and element *e*, $pop(push(s,e)) = s$..

For any stack *s* and element *e*, $top(push(s,e)) = e$..

For any stack *s*, and element *e*, $empty?(push(s,e)) = false$..

The first axiom tells us that the *pop()* operation undoes what the *push()* operation achieves.. The second tells us that when an element is pushed on a stack, it becomes the top element on the stack.. The third tells us that pushing an element on an empty stack can never make it empty.. With the right axioms, we can completely characterize the behavior of ADT operations without having to describe them informally in English (the problem is to know when we have the right axioms).. Generally, we will not pursue such an axiomatic

Stacks

specification of ADTs, though we may use some axioms from time to time to explain some ADT operations..

The Stack Interface

The stack ADT operations map stacks into one another, and this is done by having stacks as operation arguments and results.. This is what one would expect from mathematical functions that map values from the carrier set of an ADT to one another.. However, when implementing stacks in an object-oriented language that allows us to create a stack class with stack instances, there is no need to pass or return stack values—the stack instances hold values that are transformed into other values by the stack operations.. Consequently, when we specify an object-oriented implementation of the stack ADT, all stack arguments and return values vanish.. The same thing occurs as we study other ADTs and their implementations throughout this book: ADT carrier set arguments and return values will vanish from operation signatures in ADT implementation classes, and instead instances will be transformed from one carrier set value to another as operations are executed..

A Stack interface is a sub-interface of Dispenser, which is a sub-interface of Container, so it already contains an empty?() operation that it has inherited from Container.. The Stack interface need only add operations for pushing elements, popping elements, and peeking at the top element of the stack.. The diagram in Figure 1 shows the Stack interface..




Figure 1: The Stack Interface

Note that a generic or type parameter is used to generalize the interface for any element type, shown in UML as a dashed box in the upper right-hand corner of the class icon.. Note also that the preconditions of operations that need them are shown as UML properties enclosed in curly brackets to the right of the operation signatures.. Finally, because pop() no longer needs to return a stack, it can return nothing, or as a convenience, it can return the element that is popped from the stack, which is what the operation in the Stack interface is designed to do..

Using Stacks—An Example

When sending a document to a printer, one common option is to collate the output, in other words, to print the pages so that they come out in the proper order.. Generally, this means printing the last page first, the next to last next, and so forth.. Suppose a program sends several pages to a print spooler (a program that manages the input to a printer) with the instruction that they are to be collated.. Assuming that the first page arrives at the print spooler first, the second next, and so forth, the print spooler must keep the pages in a container until they all arrive, so that it can send them to the printer in reverse order.. One

way to do this is with a Stack.. Consider the pseudocode in Figure 2 describing the activities of the print spooler..

```
def printCollated(j : Job)
  Stack stack = Stack.new
  for each Page p in j do stack.push(p)
  while !stack.empty?
    print(stack.pop)
  end
end
```

Figure 2: Using A Stack to Collate Pages for Printing

A Stack is the perfect container for this job because it naturally reverses the order of the data placed into it..

Contiguous Implementation of the Stack ADT

There are two approaches to implementing the carrier set for the stack ADT: a contiguous implementation using arrays, and a linked implementation using singly linked lists; we consider each in turn..

Implementing stacks of elements of type T using arrays requires a T array to hold the of the stack and a marker to keep track of the top of the stack.. The marker can record the location of the top element, or the location where the top element would go on the next push() operation; it does not matter as long as the programmer is clear what the marker denotes and writes code accordingly..

If a static (fixed-size) array is used, then the stack can become full; if a dynamic (resizable) array is used, then the stack is essentially unbounded.. Usually, resizing an array is an expensive operation because new space must be allocated, the of the old array copied to the new, and the old space deallocated, so this flexibility is acquired at a cost..

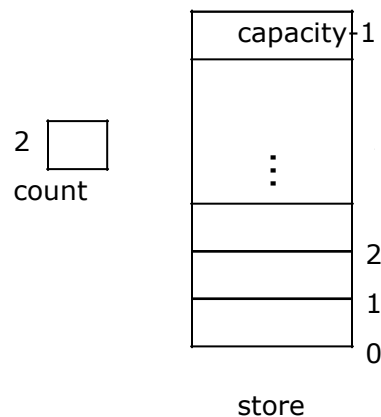


Figure 3: Implementing a Stack With an Array

We will use a dynamic array called `store` to hold stack elements, and a marker called `count` to keep track of the top of the stack.. Figure 3 below illustrates this data structure..

The marker is called `count` because it keeps track of the array location where the top element will go on the next push, and the array grows from index 0 upward, so the marker also records how many elements are in the stack..

The `store` array is allocated to hold capacity elements; this is the maximum number of items that can be placed on the stack before it must be expanded.. The diagram shows two elements in the stack, designated by the cross-hatched array locations, and the value of the `count` variable.. Note how the `count` variable also indicates the array location where the next element will be stored when it is pushed onto the stack.. The stack is empty when `count` is 0 and must be expanded when `count` equals capacity and another element is pushed on the stack..

Implementing the operations of the stack ADT using this data structure is quite straightforward.. For example, to implement the `push()` operation, a check is first made to see whether the `store` must be expanded by testing whether `count` equals capacity.. If so, the `store` is expanded.. Then the pushed value is assigned to location `store[count]` and then `count` is incremented..

A class realizing this implementation might be called `ArrayStack(T)`.. It would implement the `Stack(T)` interface and have the `store` array and the `count` variable as private attributes and the stack operations as public methods.. Its constructor would create an empty stack..

The capacity variable could be maintained by the `ArrayStack` class or be part of the implementation of dynamic arrays in the implementation language..

Linked Implementation of the Stack ADT

A linked implementation of a stack ADT uses a linked data structure to represent values of the ADT carrier set.. Lets review the basics of linked data structures..

Node: An aggregate variable composed of data and link or reference fields..

Linked (data) structure: A collection of nodes formed into a whole through its constituent node link fields..

Nodes may contain one or more data and link fields depending on the need, and the references may form a collection of nodes into linked data structures of arbitrary shapes and sizes.. Among the most important linked data structures are the following..

Singly linked list: A linked data structure whose nodes each have a single link field used to form the nodes into a sequence.. Each node link but the last contains a reference to the next node in the list; the link field of the last node contains **null** (a special reference value that does not refer to anything)..

Doubly linked list: A linked structure whose nodes each have two link fields used to form the nodes into a sequence.. Each node but the first has a predecessor link field containing a reference to the previous node in the list, and each node but the last has a successor link containing a reference to the next node in the list..

Linked tree: A linked structure whose nodes form a tree..

The linked structure needed for a stack is very simple, requiring only a singly linked list, so list nodes need only contain a value of type T and a reference to the next node.. The top element of the stack is stored at the head of the list, so the only data that the stack data structure must keep track of is the reference to the head of the list.. Figure 4 illustrates this data structure.. The reference to the head of the list is called `topNode`.. Each node has an item field (for values of type T) and a link field (for the references).. The figure shows a stack with




Figure 4: Implementing a Stack With a Singly Linked List

three elements; the top of the stack, of course, is the first node on the list.. The stack is empty when `topNode` is null; the stack never becomes full (unless memory is exhausted)..

Implementing the operations of the stack ADT using a linked data structure is quite simple.. For example, to implement the `push()` operation, a new node is created with its item field set to the new top element and its link field set to the current value of `topNode`..

The `topNode` variable is then assigned a reference to the new node..

A class realizing this implementation might be called `LinkedList(T)`.. It would implement the `Stack(T)` interface and have `topNode` as a private attribute and the stack operations as public methods.. It might also have a private inner `Node` class for node instances.. Its constructor would create an empty stack.. As new nodes are needed when values are pushed onto the stack, new `Node` instances would be instantiated and used in the list.. When values are popped off the stack, `Node` instances would be freed and their space reclaimed.. A `LinkedList` would thus use only as much space as was needed for the elements it holds..

Summary and Conclusion

Both ways of implementing stacks are simple and efficient, but the contiguous implementation either places a size restriction on the stack or uses an expensive reallocation technique if a stack grows too large.. If contiguously implemented stacks are made extra large to make sure that they don't overflow, then space may be wasted..

A linked implementation is essentially unbounded, so the stack never becomes full.. It is also very efficient in its use of space because it only allocates enough nodes to store the values actually kept in the stack.. Overall, then, the linked implementation of stacks seems slightly better than the contiguous implementation..

Review Questions

1. The `pop()` operation in the stack ADT returns a stack, while the `pop()` operation in the `Stack` interface returns a value of type T .. Why are these so different?

2. Should the `size()` operation from the Container interface return the capacity of a Stack or the number of elements currently in a Stack? What value should be returned by this operation in terms of the ArrayStack implementation?
3. The nodes in a LinkedStack hold a reference for every stack element, increasing the space needed to store data.. Does this fact invalidate the claim that a LinkedStack uses space more efficiently than an ArrayStack?

Exercises

1. State three more axioms about the stack of T ADT..
2. Suppose that an ArrayStack is implemented so that the top elements is always stored at `store[0]`.. What are the advantages or disadvantages of this approach?
3. What should happen if a precondition of a Stack operation is violated?
4. How can a programmer who is using an ArrayStack or a LinkedStack make sure that his code will not fail because it violates a precondition?
5. Should a LinkedStack have a count attribute? Explain why or why not..
6. Suppose a LinkedStack has a count attribute.. State a class invariant relating the count and `topNode` attributes..
7. Could the top element be stored at the tail of a LinkedStack? What consequences would this have for the implementation?
8. A LinkedStack could be implemented using a doubly-linked list.. What are the advantages or disadvantages of this approach?
9. As noted before, every Ruby value is an object, so every Ruby value has the same type (in a broad sense).. What consequences does this have in implementing the stack of T ADT?
10. Implement the Stack interface and the ArrayStack and LinkedStack classes in Ruby..

Review Question Answers

1. The stack ADT `pop()` operation is a mathematical function that shows how elements of the carrier set of the ADT are related to one another, specifically, it returns the stack that is the result of removing the top element of the stack that is its argument.. The `pop()` operation of the Stack interface is an operation that alters the data stored in a stack container by removing the top element stored; the new value of the ADT is represented by the data in the container, so it need not be returned as a result of the operation.. For convenience, `pop()` returns the value that is removes..
2. The Container `size()` operation, which is inherited by the Stack interface and must thus be implemented in all Stacks, should return the number of elements currently stored in a Stack.. If a Stack has an unspecified capacity (such as a resizable ArrayStack or a LinkedStack), then the capacity of the Stack is not even well defined, so it would not make sense for the `size()` operation to return the capacity..

3. Each node in a `LinkedList` contains both an element and a reference, so a `LinkedList` does use more space (perhaps twice as much space) as an `ArrayStack` to store a single element. On the other hand, an `ArrayStack` typically allocates far more space than it uses at any given moment to store data—usually there will be at least as many unused elements of the store array than used elements. This is because the `ArrayStack` must have enough capacity to accommodate the largest number of elements that could ever be pushed on the stack, but most of the time relatively little of this space is actually in use. On balance, then, `ArrayStacks` will typically use more space than `LinkedLists`.

Queues

Introduction

Queues are what we usually refer to as lines, as in “please get in line for a free lunch..” The essential features of a queue are that it is ordered and that access to it is restricted to its ends: things can enter a queue only at the rear and leave the queue only at the front..

Queue: A dispenser holding a sequence of elements that allows insertions only at one end, called the **back** or **rear**, and deletions and access to elements at the other end, called the **front**..

Queues are also called first-in-first-out, or FIFO, lists.. Queues are important in computing because of the many cases where resources provide service on a first-come-first-served basis, such as jobs sent to a printer, or processes waiting for the CPU in an operating system..

The Queue ADT

Queues are containers, and they hold values of some type.. We must therefore speak of the ADT *queue of T*, where *T* is the type of the elements held in the queue.. The carrier set of this type is the set of all queues holding elements of type *T*.. The carrier set thus includes the empty queue, the queues with one element of type *T*, the queues with two elements of type *T*, and so forth.. The essential operations of the type are the following (*q* is a queue of *T* and *e* is a *T* value)..

enter(q,e)—Return a new queue just like *q* except that *e* has been added at the rear of *q*..

leave(q)—Remove the front element of *q* and return the resulting (shorter) queue..

Attempting to remove an element from an empty queue gives an undefined result; a precondition of the *leave()* operation is that *q* is not empty..

empty?(q)—Return the Boolean value true just in case *q* is empty..

front(q)—Return the front element of *q* without removing it.. Like *leave()*, this operation has the precondition that *q* is not empty..

The queue ADT operations thus transform queues into one another.. When implemented, a single data structure stores the values in the queue, so there is generally no need to return queues once they have been changed.. Hence the operations of a queue type usually have slightly different parameters and return types, as we will see next..

The Queue Interface

A Queue interface is a sub-interface of *Dispenser*, which is a sub-interface of *Container*, so it already contains an *empty?()* operation that it has inherited from *Container*.. The Queue interface need only add operations for entering elements, removing elements, and peeking at the front element of the queue.. The diagram in Figure 1 shows the Queue interface..

Note that a generic or template is used to generalize the interface for any element type..

Note also that preconditions have been added for the operations that need them..




Figure 1: The Queue

Interface Using Queues—An Example

When sending a document to a printer, it may have to be held until the printer finishes whatever job it is working on.. Holding jobs until a printer is free is generally the responsibility of a print spooler (a program that manages the input to a printer).. Print spoolers hold jobs in a Queue until the printer is free.. This provides fair access to the printer and guarantees that no print job will be held forever.. The pseudocode in Figure 2 describes the main activities of a print spooler..

```

Queue(Job) queue;

spool( Document d ) {
    queue.enter( Job.new(d) );
}

run() {
    while ( true ) {
        if ( printer.isFree && !queue.empty? )
            printer.print( queue.leave );
    }
}
  
```

Figure 2: Using A Queue to Spool Pages for Printing

The print spooler has a job queue.. A client can ask the spooler to print a document for it using the spool() operation and the spooler will add the document to its job queue.. Meanwhile, the spooler is continuously checking the printer and its job queue; whenever the printer is free and there is a job in the queue, it will remove the job from the queue and send it to the printer..

Contiguous Implementation of the Queue ADT

There are two approaches to implementing the carrier set for the queue ADT: a contiguous implementation using arrays, and a linked implementation using singly linked lists; we consider each in turn..

Implementing queues of elements of type T using arrays requires a T array to hold the elements of the queue and some way to keep track of the front and the rear of the queue. We might, for example, decide that the front element of the queue (if any) would always be stored at location 0, and record the size of the queue, implying that the rear element would be at location $\text{size}-1$. This approach requires that the data be moved forward in the array every time an element leaves, which is not very efficient.

A clever solution to this problem is to allow the data in the array to “float” upwards as elements enter and leave, and then wrap around to the start of the array when necessary. It is as if the locations in the array are in a circular rather than a linear arrangement. Figure 3 illustrates this solution. Queue elements are held in the store array. The variable `frontIndex` keeps track of the array location holding the element at the front of the queue, and `count` holds the number of elements in the queue. The capacity is the size of the array and hence the number of elements that can be stored in the queue.

In Figure 3, data occupies the regions with an X in them: there are three elements in the queue, with the front element at `store[2]` (as indicated by the `frontIndex` variable) and the

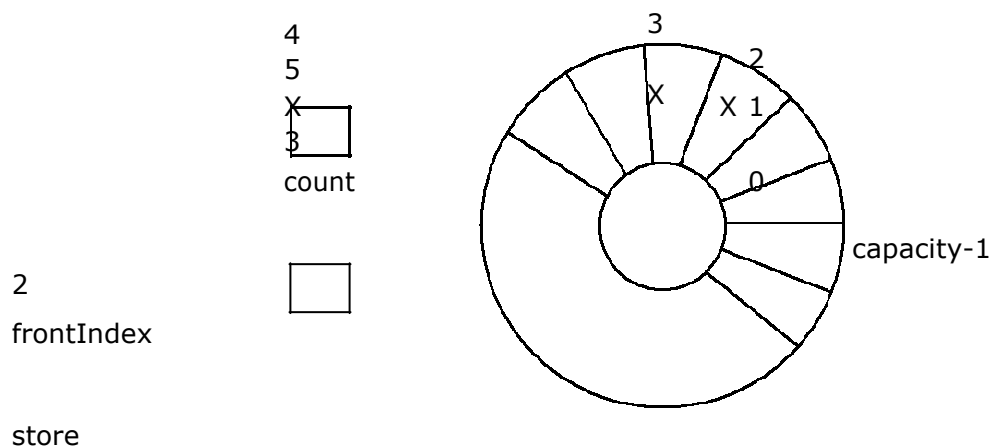


Figure 3: Implementing a Queue With a Circular Array

rear at `store[4]` (because $\text{frontIndex} + \text{count} - 1$ is 4). The next element entering the queue would be placed at `store[5]`; in general, elements enter the queue at $\text{store}[(\text{frontIndex} + \text{count}) \bmod \text{capacity}]$

The modular division is what makes the queue values wrap around the end of the array to its beginning. This trick of using a circular array is the standard approach to implementing queues in contiguous locations.

If a static array is used, then the queue can become full; if a dynamic array is used, then the queue is essentially unbounded. Usually resizing an array is an expensive operation because new space must be allocated, the contents of the array copied, and the old space deallocated, so this flexibility is acquired at a cost. Care must also be taken to move elements properly to the expanded array—remember that the front of the queue may be somewhere in the middle of the full array, with elements wrapping around to the front.

Queues

Implementing the operations of the queue ADT using this data structure is quite straightforward.. For example, to implement the leave() operation, a check is first made that the precondition of the operation (that the queue is not empty) is not violated by testing whether count equals 0.. If not, then the value at store[frontIndex] is saved in a temporary variable, frontIndex is set to (frontIndex+1) mod capacity, count is decremented, and the value stored in the temporary variable is returned..

A class realizing this implementation might be called ArrayQueue(T).. It would implement the Queue(T) interface and have the store array and the frontIndex and count variables as private attributes.. The queue operations would be public methods.. Its constructor would create an empty queue.. The value of capacity could be a constant defined at compile time, or a constructor parameter (if dynamic arrays are available)..

Linked Implementation of the Queue ADT

A linked implementation of a queue ADT uses a linked data structure to represent values of the ADT carrier set.. A singly linked list is all that is required, so list nodes need only contain a value of type T and a reference to the next node.. We could keep a reference to only the head of the list, but this would require moving down the list from its head to its tail whenever an operation required manipulation of the other end of the queue, so it is more efficient to keep a reference to each end of the list.. We will thus use both frontNode and rearNode references to keep track of both ends of the list..

If rearNode refers to the head of the list and frontNode to its tail, it will be impossible to remove elements from the queue without walking down the list from its head; in other words, we will have gained nothing by using an extra reference.. Thus we must have frontNode refer to the head of the list, and rearNode to its tail.. Figure 4 illustrates this data




Figure 4: Implementing a Queue With a Linked List

structure.. Each node has data field (for values of type T) and a link field (for the references).. The figure shows a queue with three elements.. The queue is empty when the frontNode and rearNode references are null; the queue never becomes full (unless memory is exhausted)..

Implementing the operations of the Queue ADT using a linked data structure is quite simple, though some care is needed to keep the two references synchronized.. For example, to implement the leave() operation, a check is first made that the queue is not empty.. If not, then the value field of the front node is assigned to a temporary variable.. The frontNode variable is then assigned the link field of the first node, which removes the first node from

the list.. If frontNode is null, then the list has become empty, so rearNode must also be assigned null.. Finally, the value saved in the temporary variable is returned..

A class realizing this implementation might be called `LinkedList(T)`.. It would implement the `Queue(T)` interface and have `frontNode` and `rearNode` as private attributes and the queue operations as public methods.. It might also have a private inner `Node` class for node instances.. Its constructor would create an empty queue.. As new nodes are needed when values enter the queue, new `Node` instances would be allocated and used in the list..

Summary and Conclusion

Both queue implementations are simple and efficient, but the contiguous implementation either places a size restriction on the queue or uses an expensive reallocation technique if a queue grows too large.. If contiguously implemented queues are made extra large to make sure that they don't overflow, then space may be wasted..

A linked implementation is essentially unbounded, so the queue never becomes full.. It is also very efficient in its use of space because it only allocates enough nodes to store the values actually kept in the queue..

Overall, then, the linked implementation of queues seems slightly better than the contiguous implementation..

Review Questions

1. Which operations of the queue ADT have preconditions? Do these preconditions translate to the `Queue` interface?
2. Why should storage be thought of as a circular rather than a linear arrangement of storage locations when implementing a queue using contiguous memory locations?
3. Why is there a reference to both ends of the linked list used to store the elements of a queue?

Exercises

1. In the contiguous storage implementation of a queue, is it possible to keep track of only the location of the front element (using a variable `frontIndex`) and the rear element (using a variable `rearIndex`), with no count variable? If so, explain how this would work..
2. Suppose that an `ArrayQueue` is implemented so that the array is reallocated when a client attempts to `enter()` an element when the array is full.. Assume that the reallocated array is twice as large as the full array, and write Ruby code for the `enter()` operation that includes arranging the data where it needs to be in the newly allocated array..
3. Write a class invariant for a `LinkedList` class whose attributes are `frontNode`, `rearNode`, and `count`..
4. Write a version of the `LinkedList` class that does not have a count attribute..
5. A **circular singly linked list** is a singly linked list in which the last node in the list holds a references to the first element rather than null.. It is possible to implement a

LinkedList efficiently using only a single reference into a circular singly linked list rather than two references into a (non-circular) singly linked list as we did in the text.. Explain how this works and write Ruby code for the `enter()` and `leave()` operations to illustrate this technique..

6. A LinkedList could be implemented using a doubly-linked list.. What are the advantages or disadvantages of this approach?
7. Implement the Queue interface and the ArrayQueue and LinkedList classes in Ruby..

Review Question Answers

1. The `leave()` and `front()` operations both have as their precondition that the queue q not be empty.. This translates directly into the precondition of the `leave()` and `front()` operations of the Queue interface that the queue not be empty..
2. If storage is linear, then the data in a queue will “bump up” against the end of the array as the queue grows, even if there is space at the beginning of the array for queue elements.. This problem can be solved by copying queue elements downwards in the array (inefficient), or by allowing the queue elements to wrap around to the beginning of the array, which effectively treats the array as a circular rather than a linear arrangement of memory locations..
3. In a queue, alterations are made to both ends of the container.. It is not efficient to walk down the entire linked list from its beginning to get to the far end when an alteration must be made there.. Keeping a reference to the far end of the list obviates this inefficiency and makes all queue operations very fast..

Stacks and Recursion

Introduction

Before moving on from discussing dispensers to discussing collections, we should briefly discuss the strong connection between stacks and recursion.. Recall that recursion involves operations that call themselves..

Recursive operation: An operation that either calls itself directly, or calls other operations that call it..

Recursion and stacks are intimately related in the following ways:

- Every recursive operation (or group of mutually recursive operations) can be rewritten without recursion using a stack..
- Every algorithm that uses a stack can be rewritten without a stack using one or more recursive operations..

To establish the first point, note that computers do not support recursion at the machine level—most processors can move data around, do a few simple arithmetic and logical operations, and compare values and branch to different points in a program based on the result, but that is all.. Yet many programming language support recursion.. How is this possible? At runtime, compiled programs use a stack that stores data about the current state of execution of a sub-program, called an *activation record*.. When a sub-program is called, a new activation record is pushed on the stack to hold the sub-program’s arguments, local variables, return address, and other book-keeping information.. The activation record stays on the stack until the sub-program returns, when it is popped off the stack.. Because every call of a sub-program causes a new activation record to be pushed on the stack, this mechanism supports recursion: every recursive call of a sub-program has its own activation record, so the data associated with a particular sub-program call is not confused with that of other calls of the sub-program.. Thus the recursive calls can be “unwound” onto the stack, and a non-recursive machine can implement recursion..

The second point is not quite so easy to establish, but the argument goes like this: when an algorithm would push data on a stack, a recursive operation can preserve the data that would go on the stack in local variables and then call itself to continue processing.. The recursive call returns just when it is time to pop the data off the stack, so processing can continue with the data in the local variables just as it would if the data had been popped off the stack..

In some sense, then, recursion and stacks are equivalent.. It turns out that some algorithms are easier to write with recursion, some are easier to write with stacks, and some are just as easy (or hard) one way as the other.. But in any case, there is always a way to write algorithms either entirely recursively without any stacks, or without any recursion using stacks..

In the remainder of this chapter we will illustrate the theme of the equivalence of stacks and recursion by considering a few examples of algorithms that need either stacks or recursion, and we will look at how to implement these algorithms both ways..

Balanced Brackets

Because of its simplicity, we begin with an example that doesn't really need a stack or recursion to solve, but illustrates how both can be used with equal facility: determining whether a string of brackets is balanced or not.. The strings of balanced brackets are defined as follows:

1. The empty string and the string “[”] are both strings of balanced brackets..
2. If A is a string of balanced brackets, then so is “[” A ”]..
3. If A and B are strings of balanced brackets, then so is AB ..

So, for example, “[[[[]]]]” is a string of balanced brackets, but “[[[[]]]][[]]” is not..

The recursive algorithm in Figure 1, written in Ruby, checks whether a string of brackets is balanced..

```
def recursive_balanced?(string)
  return true if string.empty?
  source = StringEnumerator.new(string)
  check_balanced?(source) && source.empty?
end

def check_balanced?(source)
  return false unless source.current == '['
  loop do
    if source.next == '['
      return false if !check_balanced?(source)
    end
    return false unless source.current ==
      ']' break if source.next != '['
  end
  true
end
```

Figure 1: Recursive Algorithm For Checking String of Balanced Brackets

The `recursive_balanced?()` operation has a string argument containing only brackets.. It does some initial and final processing, but most of the work is done by the recursive helper function `check_balanced?()`.. Note that a `StringEnumerator` is created from the string argument and passed to the `check_balanced?()` operation.. We will discuss enumerators later, but for now it suffices to say that a `StringEnumerator` is a class that provides characters from a string one by one when asked for them.. It also indicates when all the characters have been provided, signalling the end of the string.. The algorithm is based on the recursive definition of balanced brackets.. If the string is empty, then it is a string of balanced brackets.. This is checked right away by `recursive_balanced?()`.. If the string is not empty, then `check_balanced?()` is

called to check the string.. It first checks the current character to see whether it is a left bracket, and returns false if it is not.. It then considers the next character.. If it is another left bracket, then there is a nested string of balanced brackets, which is checked by a recursive call.. In any case, a check is then made for the right bracket matching the initial left bracket, which takes care of the other basis case in the recursive definition.. The loop is present to take care of the case of a sequence of balanced brackets, as allowed by the recursive definition.. Finally, when `check_balanced?()` returns its result to `recursive_balanced?()`, the latter checks to make sure that the string has all been consumed, which guarantees that there are no stray brackets at the end of the string..

This same job could be done just as well with a non-recursive algorithm using a stack.. In the code in Figure 2 below, again written in Ruby, a stack is used to hold left brackets as they are encountered.. If a right bracket is found for every left bracket on the stack, then the string of brackets is balanced.. Note that the stack must be checked to make sure it is not empty as we go along (which would mean too many right brackets), and that it is empty when the entire string is processed (which would mean too many left brackets)..

```
def stack_balanced?(string)
  stack = LinkedStack.new
  string.chars do | ch |
    case
    when ch == '['
      stack.push(ch)
    when ch == ']'
      return false if stack.empty?
      stack.pop
    else
      return false
    end
  end
  stack.empty?
end
```

Figure 2: Non-Recursive Algorithm For Checking Strings of Balanced Brackets

In this case the recursive algorithm is about as complicated as the stack-based algorithm.. In the examples below, we will see that sometimes the recursive algorithm is simpler, and sometimes the stack-based algorithm is simpler, depending on the problem..

Infix, Prefix, and Postfix Expressions

The arithmetic expressions we learned in grade school are infix expressions, but other kinds of expressions, called prefix or postfix expressions, might also be used..

Infix expression: An expression in which the operators appear between their operands..

Prefix expression: An expression in which the operators appear before their operands..

Postfix expression: An expression in which the operators appear after their operands..

In a prefix expression, the operands of an operator appear immediately to its right, while in a postfix expression, they appear immediately to its left.. For example, the infix expression $(4 + 5) * 9$ can be rewritten in prefix form as $* + 4 5 9$ and in postfix form as $4 5 + 9 *$.. An advantage of pre- and postfix expressions over infix expressions is that the latter don't need parentheses..

Many students are confused by prefix and postfix expressions the first time they encounter them, so lets consider a few more examples.. In the expressions in the table below, all numbers are one digit long and the operators are all binary.. All the expressions on a row are equivalent..

| Infix | Prefix | Postfix |
|----------------------------------|--------------------------|--------------------------|
| $(2 + 8) * (7 \% 3)$ | $* + 2 8 \% 7 3$ | $2 8 + 7 3 \% *$ |
| $((2 * 3) + 5) \% 4$ | $\% + * 2 3 5 4$ | $2 3 * 5 + 4 \%$ |
| $((2 * 5) \% (6 / 4)) + (2 * 3)$ | $+ \% * 2 5 / 6 4 * 2 3$ | $2 5 * / 6 4 \% 2 3 * +$ |
| $1 + (2 + (3 + 4))$ | $+ 1 + 2 + 3 4$ | $1 2 3 4 + + +$ |
| $((1 + 2) + 3) + 4$ | $+ + + 1 2 3 4$ | $1 2 + 3 + 4 +$ |

Note that all the expressions have the digits in the same order.. This is necessary because order matters for the subtraction and division operators.. Also notice that the order of the operators in a prefix expression is not necessarily the reverse of its order in a postfix expression; sometimes operators are in the opposite order in these expressions, but not always.. The systematic relationship between the operators is that the main operator always appears within the fewest number of parentheses in the infix expression, is first in the prefix expression, and is last in the postfix expression.. Finally, in every expression, the number of constant arguments (digits) is always one more than the number of operators..

Let's consider the problem of evaluating prefix and postfix expressions.. It turns out that sometimes it is much easier to write a recursive evaluation algorithm, and sometimes it is much easier to write a stack-based algorithm.. In particular,

- It is very easy to write a recursive prefix expression evaluation algorithm, but somewhat harder to write this algorithm with a stack..
- It is very easy to write a stack-based postfix expression evaluation algorithm, but very hard to write this algorithm recursively..

To establish these claims, we will consider a few of the algorithms.. An algorithm in Ruby to evaluate prefix expressions recursively appears in Figure 3 below.. The main operation `recursive_eval_prefix()` accepts a string as an argument.. Its job is to create a `StringEnumeration` object to pass along to the recursive helper function, and to make sure that the string has all been read (if not, then there are extra characters at the end of the expression).. The real work is done by the `eval_prefix()` operation, which is recursive..

It helps to consider the recursive definition of a prefix expression to understand this algorithm:

A prefix expression is either a digit, or if A and B are prefix expressions and op is an operator, then an expression of the form $op A B$.

The `eval_prefix()` operation first checks to see whether the string is exhausted and throws an exception if it is (because the empty string is not a prefix expression). Otherwise, it fetches the current character and advances to the next character to prepare for later processing. If the current character is a digit, this is the basis case of the recursive definition of a prefix expression, so it simply returns the integer value of the digit. Otherwise, the current character is an operator. According to the recursive definition, the operator should be followed by two prefix expressions, so the algorithm applies this operator to the result of recursively evaluating the following left and right arguments. If these arguments are not there, or are ill-formed, then one of these recursive calls will throw an exception that is propagated to the caller. The `evaluate()` operation is a helper function that simply applies the operation indicated in its `op` argument to its `left_arg` and `right_arg` values.

```
def recursive_eval_prefix(string)
  source = StringEnumerator.new(string)
  result = eval_prefix(source)
  raise "Too many arguments" unless source.empty?
  result
end

def eval_prefix(source)
  raise "Missing argument" if source.empty?
  ch = source.current
  source.next
  if ch =~ /\d/
    return ch.to_i
  else
    left_arg = eval_prefix(source)
    right_arg = eval_prefix(source)
    return evaluate(ch, left_arg, right_arg)
  end
end
```

Figure 3: Recursive Algorithm to Evaluate Prefix Expressions

This recursive algorithm is extremely simple, yet it does a potentially very complicated job. In contrast, algorithms to evaluate prefix expressions using a stack are quite a bit more complicated. One such an algorithm is shown below in Figure 4. This algorithm has two stacks: one for (integer) left arguments, and one for (character) operators.

Stacks and Recursion

```
def stack_eval_prefix(string)
  raise "Missing expression" if string.empty?
  op_stack = LinkedStack.new
  left_arg_stack = LinkedStack.new
  left_arg = right_arg = nil
  string.chars do | ch |
    case
    when ch =~ /\d/
      if left_arg == nil
        left_arg = ch.to_i
      else
        right_arg = ch.to_i
        loop do
          raise "Missing operator" if op_stack.empty?
          right_arg = evaluate(op_stack.pop, left_arg, right_arg)
          if left_arg_stack.empty?
            left_arg = right_arg
            break
          else
            left_arg = left_arg_stack.pop
          end
        end
      end
    when ch =~ /[+\-*/%]/
      op_stack.push(ch)
      if left_arg != nil
        left_arg_stack.push(left_arg)
        left_arg = nil
      end
    end
  end
  raise "Missing argument" if !op_stack.empty?
  raise "Too many arguments" unless left_arg_stack.empty?
  raise "Missing expression" unless left_arg
  left_arg
end
```

Figure 4: Stack-Based Algorithm to Evaluate Prefix Expressions

The strategy of this algorithm is to process each character from the string in turn, pushing operators on the operator stack as they are encountered and left arguments on the left argument stack as necessary.. When a right argument is encountered, as many operations are applied as possible until arguments run out.. Once the string is exhausted, the result

value should be stored in the `left_arg` variable, and the left argument and operator stacks should both be empty—if not, then there are either too many arguments or too many operators..

Clearly, this stack-based evaluation algorithm is much more complicated than the recursive evaluation algorithm.. In contrast, a stack-based evaluation algorithm for postfix expressions is quite simple, while a recursive algorithm is quite complicated.. To illustrate, consider the stack-based postfix expression evaluation algorithm in Figure 5 below..

```
def stack_eval_postfix(string)
  stack = LinkedStack.new
  string.chars do | ch |
    case
    when ch =~ /\d/
      stack.push(ch.to_i)
    when ch =~ /[+|-*\|\/%]/
      raise "Missing argument" if stack.empty?
      right_arg = stack.pop
      raise "Missing argument" if stack.empty?
      left_arg = stack.pop
      stack.push( evaluate(ch, left_arg, right_arg) )
    end
  end
  raise "Missing expresion" if stack.empty?
  raise "Too many arguments" unless stack.size ==
  1 stack.pop
end
```

Figure 5: Stack-Based Algorithm to Evaluate Postfix Expressions

The strategy of this algorithm is quite simple: there is a single stack that holds arguments, and values are pushed on the stack whenever they are encountered in the input string.. Whenever an operator is encountered, the top two values are popped of the stack, the operator is applied to them, and the result is pushed back on the stack.. This continues until the string is exhausted, at which point the final value should be on the stack.. If the stack becomes empty along the way, or there is more than one value on the stack when the input string is exhausted, then the input expression is not well-formed..

The recursive algorithm for evaluating postfix expressions is quite complicated.. The strategy is the remember arguments in local variables, making recursive calls as necessary until an operator is encountered.. We leave this algorithm as a challenging exercise..

The lesson of all these examples is that although it is always possible to write an algorithm using either recursion or stacks, in some cases a recursive algorithm is easier to develop, and in other cases a stack-based algorithm is easier.. Each problem should be explored by

sketching out both sorts of algorithms, and then choosing the one that appears easiest for detailed development..

Tail Recursive Algorithms

We have claimed that every recursive algorithms can be replaced with a non-recursive algorithm using a stack.. This is true, but it overstates the case: sometimes a recursive algorithm can be replaced with a non-recursive algorithm that does not even need to use a stack.. If a recursive algorithm is such that at most one recursive call is made as the final step in each execution of the algorithm's body, then the recursion can be replaced with a loop.. No stack is needed because data for additional recursive calls is not needed—there are no additional recursive calls.. A very simple example is a recursive algorithm to compute the factorial function, like the one shown in Figure 6..

```
def recursive_factorial(n)
  raise ArgumentError if n < 0
  (n <= 1) ? 1 : n * recursive_factorial(n-1)
end
```

Figure 6: A Recursive Factorial Algorithm

The recursion in this algorithm can be replaced with a simple loop as shown in Figure 7..

```
def factorial(n)
  raise ArgumentError if n < 0
  product = 1
  n.downto(1).each { | i | product *= i }
  product
end
```

Figure 7: A Non-Recursive Factorial Algorithm

Algorithms that only call themselves at most once as the final step in every execution of their bodies, like the factorial algorithm, are called *tail-recursive*..

Tail recursive algorithm: A recursive algorithm that calls itself at most once as the last step in every execution of its body..

Recursion can always be removed from tail-recursive algorithms without using a stack.. **Summary and Conclusion**

Algorithms that use recursion can always be replaced by algorithms that use a stack, and vice versa, so stacks and recursion are in some sense equivalent.. However, some algorithms are much easier to write using recursion, while others are easier to write using a stack.. Which

is which depends on the problem.. Programmers should evaluate both alternatives when deciding how to solve individual problems..

Review Questions

1. Which of the algorithms for determining whether a string of brackets is balanced is easiest to for you to understand?
2. What characteristics do prefix, postfix, and infix expressions share?
3. Which is easier: evaluating a prefix expression with a stack or using recursion?
4. Which is easier: evaluating a postfix expression with a stack or using recursion?
5. Is the recursive algorithm to determine whether a string of brackets is balanced tail recursive? Explain why or why not..

Exercises

1. We can slightly change the definition of strings of balanced brackets to exclude the empty string.. Restate the recursive definition and modify the algorithms to check strings of brackets to see whether they are balanced to incorporate this change..
2. Fill in the following table with equivalent expressions in each row..

| Infix | Prefix | Postfix |
|-------------------------------|----------------------|--------------------------|
| $((2 * 3) - 4) * (8 / 3) + 2$ | | |
| | $\% + 8 * 2 6 - 8 4$ | |
| | | $8 2 - 3 * 4 5 + 8 \% /$ |

3. Write a recursive algorithm to evaluate postfix expressions as discussed in this chapter..
4. Write a recursive algorithm to evaluate infix expressions.. Assume that operators have equal precedence and are left-associative, so that without parentheses, operations are evaluated from left to right.. Parentheses alter the order of evaluation in the usual way..
5. Write a stack-based algorithm to evaluate infix expressions as defined in the last exercise..
6. Which of the algorithms for evaluating infix expressions is easier to develop?
7. Write a non-recursive algorithm that does not use a stack to determine whether a string of brackets is balanced.. Hint: count brackets..

Review Question Answers

1. This answer depends on the individual, but most people probably find the stack-based algorithm a bit easier to understand because its strategy is so simple..
2. Prefix, postfix, and infix expressions list their arguments in the same order.. The number of operators in each is always one less than the number of constant arguments.. The main operator in each expression and sub-expression is easy to find: the main operator in an infix expression is the one inside the fewest number of parentheses; the main operator of

a prefix expression is the first operator; the main operator of a postfix expression is the last operator..

3. Evaluating a prefix expression recursively is much easier than evaluating it with a stack..
4. Evaluating a postfix expression with a stack is much easier than evaluating it recursively..
5. The recursive algorithm to determine whether a string of brackets is balanced calls itself at most once on each activation, but the recursive call is not the last step in the execution of the body of the algorithm—there must be a check for the closing right bracket after the recursive call.. Hence this operation is not tail recursive and it cannot be implemented without a stack.. (There is a non-recursive algorithm to check for balanced brackets without using a stack, but it uses a completely different approach from the recursive algorithms—see exercise 7)..

Lists

Introduction

Lists are simple linearly ordered collections.. Some things we refer to in everyday life as lists, such as shopping lists or laundry lists, are really sets because their order doesn't matter.. Order matters in lists.. A to-do list really is a list if the tasks to be done are in the order in which they are supposed to be completed (or some other order)..

List: An ordered collection..

Because order matters in lists, we must specify a location, or index, of elements when we modify the list.. Indices can start at any number, but we will follow convention and give the first element of a list index 0, the second index 1, and so forth..

The List ADT

Lists are collections of values of some type, so the ADT is *list of T*, where *T* is the type of the elements in the list.. The carrier set of this type is the set of all sequences or ordered tuples of elements of type *T*.. The carrier set thus includes the empty list, the lists with one element of type *T* (one-tuples), the lists with two elements of type *T* (ordered pairs), the lists with three elements of type *T* (ordered triples), and so forth.. Hence the carrier set of this ADT is the set of all tuples of type *T*, including the empty tuple..

There are many operations that may be included in a list ADT; the following operations are common (*s* and *t* are lists of *T*, *i* is an index, and *e* is a *T* value, *n* is a length, and the value *nil* is a special value that indicates that a result is undefined)..

size(t)—Return the length of list *t*..

insert(t,i,e)—Return a list just like *t* except that *e* has been inserted at index *i*, moving elements with larger indices up in the list, if necessary.. The precondition of this operation is that *i* be a valid index position: $-size(t) \leq i \leq size(t)$.. When *i* is negative, elements are counted backwards from the end of the list starting at -1, and *e* is inserted after the element; when *i* is 0, *e* is inserted at the front of the list; and when *i* is *size(t)*, *e* is appended to the end of the list..

delete_at(t,i)—Remove the element at index *i* of *t* and return the resulting (shorter) list.. The precondition of this operation is that *i* be a valid index position, with negative indices indicating a location relative to the end of the list: $-size(t) \leq i < size(t)$..

t[i]—Return the value at index *i* of *t*.. Its precondition is that *i* be a valid index position: $-size(t) \leq i < size(t)$..

t[i]=e—Replace the element at index *i* of *t* and return the resulting (same sized) list.. The precondition is that *i* be a valid index position: $-size(t) \leq i < size(t)$..

index(t,e)—Return the index of the first occurrence of *e* in *t*.. If *e* is not in *t*, return *nil*..

slice(t, i, n)—Return a new list that is a sub-list of *t* whose first element is the value at index *i* of *t* and whose length is *n*.. The precondition of this operation is that *i* is valid

and n is non-negative: $-size(t) \leq i < size(t)$ and $0 \leq n$.. The sub-list does not extend past the end of t no matter how large n is..

$t==s$ —Return true if and only if s and t have the same elements in the same order..

As with the ADTs we have studied before, an object-oriented implementation of these operations as instance methods will include the list as an implicit parameter, so the signatures of these operations will vary somewhat when they are implemented..

The List Interface

A List interface is a sub-interface of Collection, which is a sub-interface of Container, so it has several operations that it has inherited from its ancestors.. The diagram in Figure 1 shows the List interface..

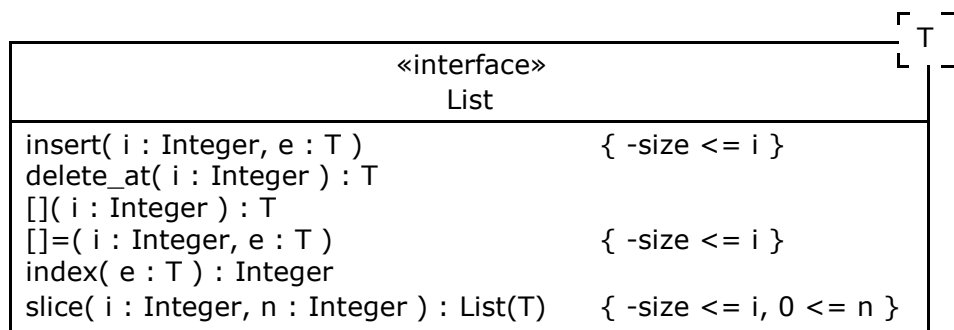


Figure 1: The List Interface

Note that a template parameter is used to generalize the interface for any element type.. Note also that preconditions have been added, though they differ slightly from those in the ADT.. In particular, when querying a list, if an index is out of bounds, then nil is returned.. When inserting or replacing an element in a list, if the index is beyond the upper bound, then the list is extended with nil values until it is large enough, and then the operation is performed.. If delete_at() is given an out of bounds index, nothing is changed..

Using Lists—An Example

Suppose a calendar program has a to-do list whose elements are ordered by precedence, so that the first element on the list must be done first, the second next, and so forth.. The items in the to-do list are all visible in a scrollable display; items can be added, removed, or moved around in the list freely.. Mousing over list items displays details about the item, which can be changed.. Users can ask to see or print sub-lists (like the last ten items on the list), or they can ask for details about a list item with a certain precedence (like the fifth element)..

Clearly, a List is the right sort of container for to-do lists.. Iteration over a List to display its is easy with an internal or external iterator.. The insert() and delete_at() operations allow items to be inserted into the list, removed from it, or moved around in it.. The [] operation can be used to obtain a list element for display during a mouse-over event, and the []= operation can replace a to-do item's record if it is changed.. The slice() operation produces portions of the list for display or printing, and the index() operation can determine where an item lies in the list..

Contiguous Implementation of the List ADT

Lists are very easy to implement with arrays.. Static arrays impose a maximum list size, while dynamic arrays allow lists to grow without limit.. The implementation is similar in either case.. An array is allocated to hold the elements of the list, with the elements placed into the array in order so that the element at index i of the list is at index i of the array.. A counter maintains the current size of the list.. Elements added at index i require that the elements in the slice from i to the end of the list be moved upwards in the array to make a “hole” into which the inserted value is placed.. When element i is removed, the slice from $i+1$ to the end of the list is copied down to close the “hole” left when the value at index i is removed..

Static arrays have their size set at compile time, so an implementation using a static array cannot accommodate lists of arbitrary size.. In contrast, an implementation using a dynamic array can allocate a larger array if a list exceeds the capacity of the current array during execution.. Reallocating the array is an expensive operation because the new, larger array must be created, the elements of the old, smaller array must be copied into the new array, and the old array must be deallocated.. To avoid this expense, the number of array reallocations should be kept to a minimum.. One popular approach is to double the size of the array whenever it needs to be made larger.. For example, suppose a list begins with a capacity of 10.. As it expands, its capacity is changed to 20, then 40, then 80, and so forth.. The array never becomes smaller..

Iterators for lists implemented with an array are also very easy to code.. The iterator object need merely keep a reference to the list and the current index during iteration, which acts as a cursor marking the current element during iteration..

Cursor: A variable marking a location in a data structure..

Accessing the element at index i of an array is almost instantaneous, so the `[]` and `[]=` list operations are very fast using a contiguous implementation.. But adding and removing elements requires moving slices of the list up or down in the array, which can be very slow.. Hence for applications of lists where elements are often accessed but not too often added or removed, the contiguous implementation of lists will be very efficient; applications that have the opposite behavior, will be much less efficient, especially if the lists are long..

Linked Implementation of the List ADT

A linked implementation of the list ADT uses a linked data structure to represent values of the ADT carrier set.. A singly or multiply linked list may be used, depending on the needs of clients.. We will consider using singly or doubly linked lists to illustrate implementation alternatives..

Suppose a singly-linked list is used to hold list elements.. It consists of a reference, traditionally called head, holding null when the list is empty, and a reference to the first node in the list otherwise.. The length of list is also typically recorded..

Most list operations take an index i as an argument, so most algorithms to implement these operations will have to begin by walking down the list to locate the node at position i or position $i-1$.. Why $i-1$? Because for addition and removal, the link field in the node

preceding node i will have to be changed.. In any case, if lists are long and there are many changes towards the end of the list, much processing will be done simply finding the right spots in the list to do things..

This difficulty can be alleviated by keeping a special cursor consisting of an index number and a pointer into the list.. The cursor is used to find the node that some operation needs to do its job.. The next time an operation is called, it may be able to use the existing value of the cursor, or use it with slight changes, thus saving time.. For example, suppose that a value is added at the end of the list.. The cursor is used to walk down to the end of the list and make the addition; when this task is done, the cursor marks the node at, let us say, location $size-2$ in the list.. If another addition is made, the cursor only needs to be moved forward one node to the end of the list to mark the node whose link field must be changed—the walk down the list from its beginning has been avoided..

It may be useful to maintain a pointer to the end of the list.. Then operations at the end of the list can be done quickly in exchange for the slight extra effort of maintaining the extra pointer.. If a client does many operations at the end of a list, the extra work will be justified..

Another way to make list operations faster is to store elements in a doubly linked list in which each node (except those at the ends) has a reference to both its successor and its predecessor nodes.. Figure 2 below illustrates this setup.. Using a cursor with a doubly linked list can speed things up considerably because the links make it possible to move both backwards and forwards in the list.. If an operation needs to get to node i and the cursor marks node j , which is closer to node i than node i is to the head of the list, following links from the cursor can get to node i more quickly than following links from the head of the list.. Keeping a pointer to the end of the list makes things faster still: it is possible to start walking toward a node from three points: the front of the list, the back of the list, or the cursor, which is often somewhere in the middle of the list..




Figure 2: A Doubly Linked List

Another trick is to make the list **circular**: have the pred pointer of the first node point to the last node rather than containing nil, and have the succ pointer of the last node point to the first node rather than containing nil.. Then there is no need for a separate pointer to the end of the list: the head pointer can be used to get to both the front and the rear of the list..

Iterators for the linked implementation of lists must obtain a pointer to the head of the list, which can be passed to the instance constructor.. Then it is merely a question of maintaining a cursor and walking down the list as the `Iterator.next()` operation is called..

Modifying linked implementations of lists is very fast once the nodes to operate on have been found, and using doubly linked lists with cursors can make node finding fairly fast..

A linked implementation is a good choice when a variety of list operations are needed..

Implementing Lists in Ruby

The Ruby array is already very similar to a contiguously implemented list; it lacks only `contains?()` and `iterator()` operations.. Hence the simplest way to implement an `ArrayList` in Ruby is to make the `ArrayList` a sub-class of the built-in Ruby `Array` class and add the two missing operations.. Of course, the `iterator()` factory method must return an instance of an `Iterator` that traverses the `ArrayList`, so an `ArrayListIterator` class must also be written..

A `LinkedList` can be implemented in Ruby by making linked lists using a `Node` class containing data and reference fields and writing operations to perform list operations as discussed above.. In this case `LinkedList` is a sub-class of `List`, an interface that is a sub-class of `Collection`, which in turn is a sub-class of `Container`..

Summary and Conclusion

The contiguous implementation of lists is very easy to program and quite efficient for element access, but quite slow for element addition and removal.. The linked implementation is considerably more complex to program, and can be slow if operations must always locate nodes by walking down the list from its head, but if double links and a cursor are used, list operations can be quite fast across the board.. The contiguous implementation is preferable for lists that don't change often but must be accessed quickly, while the linked implementation is better when these characteristics don't apply..

Review Questions

1. Does it matter where list element numbering begins?
2. What does the list ADT $index(t,e)$ operation return when e is not in list t ?
3. What is a cursor?
4. Under what conditions does the contiguous implementation of the list ADT not perform well?
5. What advantage does a doubly linked list provide over a singly linked list?
6. What advantage does a circular doubly-linked list provide over a non-circular list?
7. What advantage does a cursor provide in the linked implementation of lists?
8. In an application where long lists are changed infrequently but access to the middle of the lists are common, would a contiguous or linked implementation be better?

Exercises

1. Do we really need iterators for lists? Explain why or why not..
2. Would it be worthwhile to maintain a cursor for a contiguously implemented list? Explain why or why not..
3. What should happen if a precondition of a `List` operation is violated?

Lists

4. In the List interface, why do the `delete_at()` and `[]()` operations not have preconditions while the `insert()`, `[]=()`, and `slice()` operations do have preconditions?
5. An Iterator must have a reference to the Collection with which it is associated.. How is such a connection made if a concrete iterator class is separate from its associated collection class? For example, suppose the `ArrayListIterator` is a class entirely separate from the `ArrayList` class.. How can an `ArrayListIterator` instance obtain a connection to the particular `ArrayList` instance it is supposed to traverse?
6. A `ListIterator` is a kind of `Iterator` that (a) allows a client to start iteration at the end of list and go through it backwards, (b) change the direction of iteration during traversal, and (c) obtain the index as well as the value of the current element.. Write an interface for `ListIterator` that is a sub-interface of `Iterator`.
7. Write `ArrayList` and `ArrayListListIterator` implementations in Ruby.. Note that the List interface is very almost entirely implemented by the Ruby `Array` class: all that is missing is the `iterator()` and `contains?()` operations..
8. Write an `ArrayListListIterator` to go along with the implementation in the last exercise..
9. Write a `LinkedList` implementation in Ruby that uses a singly-linked list, no reference to the end of the list, and a cursor.. Write a `LinkedListIterator` to go along with it..
10. Write a `LinkedListListIterator` to go along with the `LinkedList` class in the previous exercise..
11. Write a `LinkedList` implementation in Ruby that uses a doubly-linked circular list and a cursor.. Write a `LinkedListIterator` to go along with it..
12. Write a `LinkedListListIterator` to go along with the `LinkedList` class in the previous exercise..

Review Question Answers

1. It does not matter where list element numbering begins: it may begin at any value.. However, it is usual in computing to start at zero, and in everyday life to start at one, so one of these values is preferable..
2. The list ADT $index(t,e)$ operation cannot return an index when e is not in list t .. It should return a value indicating that the result is undefined, namely *nil*..
3. A cursor is a variable marking a location in a data structure.. In the case of a List, a cursor is a data structure marking a particular element in the List.. For an `ArrayList`, a cursor might be simply an index.. For a `LinkedList`, it is helpful for the cursor to hold both the index and a reference to the node where the item is stored..
4. A contiguous implementation of the list ADT does not perform well when the list is long and it is changed near its beginning often.. Every change near the beginning of a long contiguously implemented list requires that many list elements be copied up or down the list, which is expensive..

Lists

5. A doubly linked list makes it faster to move around than in a singly linked list, which can speed up the most expensive part of linked list operations: finding the node in the list where the operation must do its job..
6. A circular doubly-linked list makes it possible to follow links quickly from the list head to the end of the list.. This can only be done in a non-circular list if a reference to the end of the list is maintained..
7. A cursor helps to speed up linked list operations by often making it faster to get to the nodes where the operations must do their work.. Even in a circular doubly-linked list, it is expensive to get to the middle of the list.. If a cursor is present and it ends up near the middle of a list after some list operation, then it can be used to get to a node in the middle of the list very quickly..
8. In an application where long lists are changed infrequently but are accessed near their middle often, a contiguous implementation will likely be better than a linked implementation because access to the middle of a contiguously implemented list (no matter how long) is instantaneous, while access to the middle of a linked list will almost always be slower, and could be extremely slow (if link following must begin at the end of the list)..

13: Basic Sorting Algorithms

Introduction

Sorting is one of the most fundamental and important data processing tasks..

Sorting algorithm: An algorithm that rearranges records in lists so that they follow some well-defined ordering relation on values of keys in each record..

An *internal* sorting algorithm works on lists in main memory, while an *external* sorting algorithm works on lists stored in files.. Some sorting algorithms work much better as internal sorts than external sorts, but some work well in both contexts.. A sorting algorithm is *stable* if it preserves the original order of records with equal keys..

Many sorting algorithms have been invented; in this chapter we will consider the simplest sorting algorithms.. In our discussion in this chapter, all measures of input size are the length of the sorted lists (arrays in the sample code), and the basic operation counted is comparison of list elements (also called *keys*)..

Bubble Sort

One of the oldest sorting algorithms is Bubble sort.. The idea behind the sort is to make repeated passes through the list from beginning to end, comparing adjacent elements and swapping any that are out of order.. After the first pass, the largest element will have been moved to the end of the list; after the second pass, the second largest will have been moved to the penultimate position; and so forth.. The idea is that large values “bubble up” to the top of the list on each pass..

A Ruby implementation of Bubble sort appears in Figure 1..

```
def bubble_sort(array) (array.size-
  1).downto(1).each do | j |
  1.upto(j).each do | i |
    if array[i] < array[i-1]
      array[i], array[i-1] = array[i-1], array[i]
    end
  end
end
return array
end
```

Figure 1: Bubble Sort

It should be clear that the algorithm does exactly the same key comparisons no matter what the size of the array, so we need only consider the basic complexity of the algorithm..

On the first pass through the data, every element in the array but the first is compared with its predecessor, so $n-1$ comparisons are made.. On the next pass, one less comparison is made,

so $n-2$ comparisons are made.. This continues until the last pass, where only one comparison is made.. The total number of comparisons is thus given by the following summation..

$$C(n) = \sum_{i=1 \text{ to } n-1} i = n(n-1)/2$$

Clearly, $n(n-1)/2 \in O(n^2)$..

Bubble sort is not very fast.. Various suggestions have been made to improve it.. For example, a Boolean variable can be set to false at the beginning of each pass through the list and set to true whenever a swap is made.. If the flag is false when the pass is completed, then no swaps were done and the array is sorted, so the algorithm can halt.. This gives exactly the same worst case complexity, but a best case complexity of only n .. The average case complexity is still in $O(n^2)$, however, so this is not much of an improvement..

Selection Sort

The idea behind Selection sort is to make repeated passes through the list, each time finding the largest (or smallest) value in the unsorted portion of the list, and placing it at the end (or beginning) of the unsorted portion, thus shrinking the unsorted portion and growing the sorted portion.. Thus, the algorithm works by repeatedly “selecting” the item that goes at the end of the unsorted portion of the list..

A Ruby implementation of Selection sort appears in Figure 2..

```
def selection_sort(array)
  0.upto(array.size-2).each do | j |
    min_index = j
    (j+1).upto(array.size-1).each do | i |
      min_index = i if array[i] < array[min_index]
    end
    array[j], array[min_index] = array[min_index], array[j] end
  return array
end
```

Figure 2: Selection Sort

This algorithm finds the minimum value in the unsorted portion of the list $n-1$ times and puts it where it belongs.. Like Bubble sort, it does exactly the same thing no matter what the of the array, so we need only consider its basic complexity..

On the first pass through the list, Selection sort makes $n-1$ comparison; on the next pass, it makes $n-2$ comparisons; on the third, it makes $n-3$ comparisons, and so forth.. It makes $n-1$ passes altogether, so its complexity is

$$C(n) = \sum_{i=1 \text{ to } n-1} i = n(n-1)/2$$

As noted before, $n(n-1)/2 \in O(n^2)$..

Although the number of comparisons that Selection sort makes is identical to the number that Bubble sort makes, Selection sort usually runs considerably faster.. This is because Bubble sort typically makes many swaps on every pass through the list, while Selection sort makes only one.. Nevertheless, neither of these sorts is particularly fast..

Insertion Sort

Insertion sort works by repeatedly taking an element from the unsorted portion of a list and inserting it into the sorted portion of the list until every element has been inserted.. This algorithm is the one usually used by people when sorting piles of papers..

A Ruby implementation of Insertion sort appears in Figure 3..

```
def insertion_sort(array)
  1.upto(array.size-1).each do | j |
    element = array[j]
    i = j
    while 0 < i && element < array[i-1]
      array[i] = array[i-1]
      i -= 1
    end
    array[i] = element
  end
  return array
end
```

Figure 3: Insertion Sort

A list with only one element is already sorted, so the elements inserted begin with the second element in the array.. The inserted element is held in the `element` variable and values in the sorted portion of the array are moved up to make room for the inserted element in the same loop where the search is done to find the right place to make the insertion.. Once it is found, the loop ends and the inserted element is placed into the sorted portion of the array..

Insertion sort does different things depending on the of the list, so we must consider its worst, best, and average case behavior.. If the list is already sorted, one comparison is made for each of $n-1$ elements as they are “inserted” into their current locations.. So the best case behavior of Insertion sort is

$$B(n) = n-1$$

The worst case occurs when every inserted element must be placed at the beginning of the already sorted portion of the list; this happens when the list is in reverse order.. In this case, the first element inserted requires one comparison, the second two, the third three, and so forth, and $n-1$ elements must be inserted.. Hence

$$W(n) = \sum_{i=1 \text{ to } n-1} i = n(n-1)/2$$

To compute the average case complexity, let's suppose that the inserted element is equally likely to end up at any location in the sorted portion of the list, as well as the position it initially occupies.. When inserting the element with index j , there are $j+1$ locations where the element may be inserted, so the probability of inserting into each location is $1/(j+1)$.. Hence the average number of comparison to insert the element with index j is given by the following expression..

$$\begin{aligned} & 1/(j+1) + 2/(j+1) + 3/(j+1) + \dots + j/(j+1) + j/(j+1) \\ &= 1/(j+1) \cdot \sum_{i=1 \text{ to } j} i + j/(j+1) \\ &= 1/(j+1) \cdot j(j+1)/2 + j/(j+1) \\ &= j/2 + j/(j+1) \end{aligned}$$

$$\approx j/2 + 1$$

The quantity $j/(j+1)$ is always less than one, and very close to one for large values of j , so we simplify the expression as noted above to produce a close upper bound for the count of the average number of comparisons done when inserting the element with index j .. We will use this simpler expression in our further computations because we know that the result will always be a close upper bound on the number of comparisons..

We see that when inserting an element into the sorted portion of the list, we have to make comparisons with about half the elements in that portion of the list, which makes sense..

Armed with this fact, we can now write down an equation for the approximate average case complexity:

$$\begin{aligned} A(n) &= \sum_{j=1 \text{ to } n-1} (j/2 + 1) \\ &= \frac{1}{2} \sum_{j=1 \text{ to } n-1} j + \sum_{j=1 \text{ to } n-1} 1 \\ &= \frac{1}{2} (n(n-1)/2) + (n-1) \\ &= (n^2 + 3n - 4)/4 \end{aligned}$$

In the average case, Insertion sort makes about half as many comparisons as it does in the worst case.. Unfortunately, both these functions are in $O(n^2)$, so Insertion sort is not a great sort.. Nevertheless, Insertion sort is quite a bit better than Bubble and Selection sort on average and in the best case, so it is the best of the three $O(n^2)$ sorting algorithms..

Insertion sort has one more interesting property to recommend it: it sorts nearly sorted lists very fast.. A *k-nearly sorted list* is a list all of whose elements are no more than k positions from their final locations in the sorted list.. Inserting any element into the already sorted portion of the list requires at most k comparisons.. A close upper bound on the worst case complexity of Insertion sort on a k -nearly sorted list is:

$$W(n) = \sum_{i=1 \text{ to } n-1} k = k \cdot (n-1)$$

Because k is a constant, $W(n)$ is in $O(n)$, that is, Insertion sort always sorts a nearly sorted list in linear time, which is very fast indeed..

Basic Sorting Algorithms

Although our analysis of Insertion sort only counted comparisons of keys, there are also many comparisons of indices.. We can lessen this overhead cost and speed up the algorithm quite a bit by placing a sentinel value at the start of the list..

Sentinel value: A special value placed in a data structure to mark a boundary..

Extreme values can sometimes be found to use for sentinels, but often using the minimum or maximum value in a data structure, as we do here, works well..

In this case the sentinel value needed is the smallest value in the list.. By placing it in the first location, the algorithm need not check the index marching backwards through the list to make sure that it does not go beyond the start of the list—comparison with the sentinel value will always fail, obviating the index check.. Ruby code for this revised version of Insertion sort appears in Figure 4..

```
def sentinel_insertion_sort(array)
  # first put the minimum value in location
  0 min_index = 0;
  1.upto(array.size-1).each do | index |
    min_index = index if array[index] <
array[min_index] end
  array[0], array[min_index] = array[min_index], array[0]

  # now insert elements into the sorted portion
  2.upto(array.size-1).each do | j |
    element = array[j]
    i = j
    while element < array[i-1]
      array[i] = array[i-1]
      i -= 1
    end
    array[i] = element
  end
  return array
end
```

Figure 4: Insertion Sort with a Sentinel Value

Insertion sort with a sentinel value still has average and worst case complexity in $O(n^2)$, best case complexity in $O(n)$, and it still sorts nearly sorted lists in linear time, but it is perhaps 40% faster than the unmodified Insertion sort algorithm most of the time..

Shell Sort

Shell sort is an interesting variation of Insertion sort invented by Donald Shell in 1959.. It works by Insertion sorting the elements in a list that are h positions apart for some h , then decreasing h and doing the same thing over again until $h = 1$..

Basic Sorting Algorithms

A version of Shell sort in Ruby appears in Figure 5..

```
def shell_sort(array)
  # compute the starting value of h
  h = 1;
  h = 3*h + 1 while h < a.size/9

  # insertion sort using decreasing values of h
  while 0 < h do
    h.upto(array.size-1).each do | j |
      element = array[j]
      i = j
      while 0 < i && element < array[i-h]
        array[i] = array[i-h]
        i -= h
      end
      array[i] = element
    end
    h /= 3
  end
  return array
end
```

Figure 5: Shell Sort

Although Shell sort has received much attention over many years, no one has been able to analyze it yet! It has been established that for many sequences of values of h (including those used in the code above), Shell sort never does more than $n^{1.5}$ comparisons in the worst case.. Empirical studies have shown that it is quite fast on most lists.. Hence Shell sort is the fastest sorting algorithm we have considered so far..

Summary and Conclusion

For small lists of less than a few hundred elements, any of the algorithms we have considered in this chapter are adequate.. For larger lists, Shell sort is usually the best choice, except in a few special cases:

- If a list is nearly sorted, use Insertion sort;
- If a list contains large records that are very expensive to move, use Selection sort because it does the fewest number of data moves (of course, the fast algorithms we study in a later chapter are even better)..

Never use Bubble sort: it makes as many comparisons as any other sort, and usually moves more data than any other sort, so it is generally the slowest of all..

Basic Sorting Algorithms

Review Questions

1. What is the difference between internal and external sorts?
2. The complexity of Bubble sort and Selection sort is exactly the same.. Does this mean that there is no reason to prefer one over the other?
3. What is a sentinel value?
4. Would a sentinel value be useful in Shell sort?

Exercises

1. Rewrite the Bubble sort algorithm to incorporate a check to see whether the array is sorted after each pass, and to stop processing when this occurs..
2. An alternative to Bubble sort is the Cocktail Shaker sort, which uses swaps to move the largest value to the top of the unsorted portion, then the smallest value to the bottom of the unsorted portion, then the largest value to the top of the unsorted portion, and so forth, until the array is sorted..
 - (a) Write code for the Cocktail Shaker sort..
 - (b) What is the complexity of the Cocktail Shaker sort?
 - (c) Does the Cocktail Shaker sort have anything to recommend it (besides its name)?
3. Adjust the Selection sort algorithm presented above to sort using the maximum rather than the minimum element in the unsorted portion of the array..
4. Every list of length n is n -nearly sorted.. Using the formula for the worst case complexity of Insertion sort on a k -nearly sorted list with $k = n$, we get $W(n) = n(n-1)$.. Why is this result different from $W(n) = n(n-1)/2$, which we calculated elsewhere?
5. Rewrite Shell sort using a sentinel value to avoid having to check to make sure that invalid array accesses are not attempted..
6. The sorting algorithms presented in this chapter are written for ease of analysis and do not take advantage of all the features of Ruby.. Rewrite the sorting algorithms using as many features of Ruby as possible to shorten the algorithms or make them faster..
7. A certain data collection program collects data from seven remote stations that it contacts over the Internet.. Every minute, the program sends a message to the remote stations prompting each of them to collect and return a data sample.. Each sample is time stamped by the remote stations.. Because of transmission delays, the seven samples do not arrive at the data collection program in time stamp order.. The data collection program stores the samples in an array in the order in which it receives them.. Every 24 hours, the program sorts the samples by time stamp and stores them in a database.. Which sorting algorithm should the program use to sort samples before they are stored: Bubble, Selection, Insertion, or Shell sort? Why?

Review Question Answers

1. An internal list processes lists stored in main memory, while an external sorts processes lists stored in files..

Basic Sorting Algorithms

2. Although the complexity of Bubble sort and Selection sort is exactly the same, in practice they behave differently.. Bubble sort tends to be significantly slower than Selection sort, especially when list elements are large entities, because Bubble sort moves elements into place in the list by swapping them one location at a time, while Selection sort merely swap one element into place on each pass.. Bubble sort makes $O(n^2)$ swaps on average, while Selection sort $O(n)$ swaps in all cases; had we chosen swaps as a basic operation, this difference would have been reflected in our analysis..
3. A sentinel value is a special data value placed in a data structure to mark a boundary..
4. Because Shell sort is abased on Insertion sort, and a sentinel value slightly speeds up Insertion sort, it is reasonable to assume that using a sentinel value in Shell sort would speed it up a bit.. Unfortunately, sentinels must be found for each of the sub-lists, which requires many passes over the data, and actually slows down Shell sort slightly..

Merge sort and Quicksort

Introduction

The sorting algorithms we have looked at so far are not very fast, except for Shell sort and Insertion sort on nearly-sorted lists.. In this chapter we consider two of the fastest sorting algorithms known: merge sort and quicksort..

Merge Sort

Merge sort is a classic example of a *divide and conquer algorithm* that solves a large problem by dividing it into parts, solving the resulting smaller problems, and then combining these solutions into a solution to the original problem.. The strategy of merge sort is to sort halves of a list (recursively) then merge the results into the final sorted list.. Merging is a pretty fast operation, and breaking a problem in half repeatedly quickly gets down to lists that are already sorted (lists of length one or less), so this algorithm performs well.. A Ruby implementation of merge sort appears in Figure 1 below..

```
def merge_sort(array)
  merge_into(array.dup, array, 0, array.size)
  return array
end

def merge_into(src, dst, lo, hi)
  return if hi-lo < 2
  m = (lo+hi)/2
  merge_into(dst, src, lo, m)
  merge_into(dst, src, m, hi)
  j = lo; k = m
  (lo..hi-1).each do | i |
    if j < m && k < hi
```

```
        if src[j] < src[k]
            dst[i] = src[j]; j += 1
        else
            dst[i] = src[k]; k += 1
        end
    elsif j < m
        dst[i] = src[j]; j += 1
    else # k < lo
        dst[i] = src[k]; k += 1
    end
end
end
end
```

Figure 1: Merge Sort

Merge Sort and Quicksort

Merging requires a place to store the result of merging two lists, and duplicating the original list provides space for merging.. Hence this algorithm duplicates the original list and passes the duplicate to `merge_into()`. This operation recursively sorts the two halves of the auxiliary list and then merges them back into the original list.. Although it is possible to sort and merge in place, or to merge using a list only half the size of the original, the algorithms to do merge sort this way are complicated and have a lot of overhead—it is simpler and faster to use an auxiliary list the size of the original, even though it requires a lot of extra space..

In analyzing this algorithm, the measure of the size of the input is, of course, the length of the list sorted, and the operations counted are key comparisons.. Key comparison occurs in the merging step: the smallest items in the merged sub-lists are compared, and the smallest is moved into the target list.. This step is repeated until one of the sub-lists is exhausted, in which case the remainder of the other sub-list is copied into the target list..

Merging does not always take the same amount of effort: it depends on the of the sub-lists.. In the best case, the largest element in one sub-list is always smaller than the smallest element in the other (which occurs, for example, when the input list is already sorted).. If we make this assumption, along with the simplifying assumption that $n = 2^k$, then the recurrence relation for the number of comparisons in the best case is

$$\begin{aligned} B(n) &= n/2 + 2 \cdot B(n/2) \\ &= n/2 + 2 \cdot (n/4 + 2 \cdot B(n/4)) = 2 \cdot n/2 + 4 \cdot B(n/4) \\ &= 2 \cdot n/2 + 4 \cdot (n/8 + 2 \cdot B(n/8)) = 3 \cdot n/2 + 8 \cdot B(n/8) \\ &= \dots \\ &= i \cdot n/2 + 2_i \cdot B(n/2_i) \end{aligned}$$

The initial condition for the best case occurs when n is one or zero, in which case no comparisons are made.. If we let $n/2_i = 1$, then $i = k = \lg n$.. Substituting this into the equation above, we have

$$B(n) = \lg n \cdot n/2 + n \cdot B(1) = (n \lg n)/2$$

Thus, in the best case, merge sort makes only about $(n \lg n)/2$ key comparisons, which is quite fast.. It is also obviously in $O(n \lg n)$..

In the worst case, making the most comparisons occurs when merging two sub-lists such that one is exhausted when there is only one element left in the other.. In this case, every merge operation for a target list of size n requires $n-1$ comparisons.. We thus have the following recurrence relation:

$$\begin{aligned} W(n) &= n-1 + 2 \cdot W(n/2) \\ &= n-1 + 2 \cdot (n/2-1 + 2 \cdot W(n/4)) = n-1 + n-2 + 4 \cdot W(n/4) \\ &= n-1 + n-2 + 4 \cdot (n/4-1 + 2 \cdot W(n/8)) = n-1 + n-2 + n-4 + 8 \cdot W(n/8) \\ &= \dots \\ &= n-1 + n-2 + n-4 + \dots + n-2_{i-1} + 2_i \cdot W(n/2_i) \end{aligned}$$

The initial conditions are the same as before, so we may again let $i = \lg n$ to solve this recurrence..

$$\begin{aligned}
W(n) &= n-1 + n-2 + n-4 + \dots + n-2^{i-1} + 2^i \cdot W(n/2^i) \\
&= n-1 + n-2 + n-4 + \dots + n-2^{\lg n - 1} \\
&= \sum_{j=0}^{\lg n - 1} n - 2^j \\
&= \sum_{j=0}^{\lg n - 1} n - \sum_{j=0}^{\lg n - 1} 2^j \\
&= n \sum_{j=0}^{\lg n - 1} 1 - (2^{\lg n - 1 + 1} - 1) \\
&= n \lg n - n + 1
\end{aligned}$$

The worst case behavior of merge sort is thus also in $O(n \lg n)$.

As an average case, let's suppose that each comparison of keys from the two sub-lists is equally likely to result in an element from one sub-list being moved into the target list as from the other. This is like flipping coins: it is as likely that the element moved from one sub-list will win the comparison as an element from the other. And like flipping coins, we expect that in the long run, the elements chosen from one list will be about the same as the elements chosen from the other, so that the sub-lists will run out at about the same time. This situation is about the same as the worst case behavior, so on average, merge sort will make about the same number of comparisons as in the worst case.

Thus, in all cases, merge sort runs in $O(n \lg n)$ time, which means that it is significantly faster than the other sorts we have seen so far. Its major drawback is that it uses $O(n)$ extra memory locations to do its work.

Quicksort

The most widely studied, widely used (by professionals), and fastest of all algorithms that sort by comparison of keys is quicksort, invented by C. A. R. Hoare in 1960. A Ruby implementation of quicksort appears in Figure 2 below.

Quicksort is a divide and conquer algorithm. It works by selecting a single element in the list, called the *pivot element*, and rearranging the list so that all elements less than or equal to the pivot are to its left, and all elements greater than or equal to it are to its right. This operation is called *partitioning*. Once a list is partitioned, the algorithm calls itself recursively to sort the sub-lists left and right of the pivot. Eventually, the sub-lists have length one or less, at which point they are sorted, ending the recursion.

The heart of quicksort is the partitioning algorithm. This algorithm must choose a pivot element and then rearrange the list as quickly as possible so that the pivot element is in its final position, all values greater than the pivot are to its right, and all values less than it are to its left. Although there are many variations of this algorithm, the general approach is to choose an arbitrary element as the pivot, scan from the left until a value greater than the pivot is found, and from the right until a value less than the pivot is found. These values are then swapped, and the scans resume. The pivot element belongs in the position where the scans meet. Although it seems very simple, the quicksort partitioning algorithm is quite subtle and hard to get right. For this reason, it is generally a good idea to copy it from a source that has tested it extensively.

Merge Sort and Quicksort

```
def quick(array, lb, ub)
  return if ub <= lb
  pivot = array[ub]
  i, j = lb-1, ub
  loop do
    loop do i += 1; break if pivot <= array[i]; end
    loop do j -= 1; break if j <= lb || array[j] <= pivot;
    end array[i], array[j] = array[j], array[i] break if j <=
    i
  end
  array[j], array[i], array[ub] = array[i], pivot, array[j]
  quick(array, lb, i-1)
  quick(array, i+1, ub)
end

def quicksort(array)
  quick(array, 0, array.size-1)
  return array
end
```

Figure 2: Quicksort

We analyze this algorithm using the list size as the measure of the size of the input, and using comparisons as the basic operation.. Quicksort behaves very differently depending on the of the list it sorts.. In the best case, the pivot always ends up right in the middle of the partitioned sub-lists.. We assume, for simplicity, that the original list has 2^{k-1} elements.. The partitioning algorithm compares the pivot value to every other value, so it makes $n-1$ comparisons on a list of size n .. This means that the recurrence relation for the number of comparison is

$$B(n) = n-1 + 2 \cdot B((n-1)/2)$$

The initial condition is $B(n) = 0$ for $n = 0$ or 1 because no comparisons are made on lists of size one or empty lists.. We may solve this recurrence as follows:

$$\begin{aligned} B(n) &= n-1 + 2 \cdot B((n-1)/2) \\ &= n-1 + 2 \cdot ((n-1)/2 - 1 + 2 \cdot B(((n-1)/2 - 1)/2)) \\ &= n-1 + n-3 + 4 \cdot B((n-3)/4) \\ &= n-1 + n-3 + 4 \cdot ((n-3)/4 - 1 + 2 \cdot B(((n-3)/4 - 1)/2)) \\ &= n-1 + n-3 + n-7 + 2 \cdot B((n-7)/8) \\ &= \dots \\ &= n-1 + n-3 + n-7 + \dots + (n-(2^{i-1})) + 2^i \cdot B((n-(2^{i-1}))/2^i) \end{aligned}$$

If we let $(n-(2^{i-1}))/2^i = 1$ and solve for i , we get $i = k-1$.. Substituting, we have

$$\begin{aligned}
B(n) &= n-1 + n-3 + n-7 + \dots + n-(2^{k-1}-1) \\
&= \sum_{i=0}^{k-1} (n - (2^i - 1)) \\
&= \sum_{i=0}^{k-1} n - \sum_{i=0}^{k-1} (2^i - 1) \\
&= k \cdot (n+1) - (2^k - 1) \\
&= (n+1) \lg (n+1) - n
\end{aligned}$$

Thus the best case complexity of quicksort is in $O(n \lg n)$.

Quicksort's worst case behavior occurs when the pivot element always ends up at one end of the sub-list, meaning that sub-lists are not divided in half when they are partitioned, but instead one sub-list is empty and the other has one less element than the sub-list before partitioning. If the first or last value in the list is used as the pivot, this occurs when the original list is already in order or in reverse order. In this case the recurrence relation is

$$W(n) = n-1 + W(n-1)$$

This recurrence relation is easily solved and turns out to be $W(n) = n(n-1)/2$, which of course we know to be $O(n^2)$!

The average case complexity of quicksort involves a recurrence that is somewhat hard to solve, so we simply present the solution: $A(n) = 2(n+1) \cdot \ln 2 \cdot \lg n \approx 1.39 (n+1) \lg n$. This is not far from quicksort's best case complexity. So in the best and average cases, quicksort is very fast, performing $O(n \lg n)$ comparisons; but in the worst case, quicksort is very slow, performing $O(n^2)$ comparisons.

Improvements to Quicksort

Quicksort's worst case behavior is abysmal, and because it occurs for sorted or nearly sorted lists, which are often encountered, this is a big problem. Many solutions to this problem have been proposed, but perhaps the best is called the *median-of-three improvement*, and it consists of using the median of the first, last, and middle values in each sub-list as the pivot element. Except in rare cases, this technique produces a pivot value that ends up near the middle of the sub-list when it is partitioned, especially if the sub-list is sorted or nearly sorted. A version of quicksort with the median-of-three improvement appears in Figure 3 below. The median finding process also allows sentinel values to be placed at the ends of the sub-list, which speeds up the partitioning algorithm a little bit as well. From now on, we will assume that quicksort includes the median-of-three improvement.

Other improvements to quicksort have been proposed, and each speeds it up slightly at the expense of making it a bit more complicated. Among the suggested improvements are the following:

- Use Insertion sort for small sub-lists (ten to fifteen elements). This eliminates a lot of recursive calls on small sub-lists, and takes advantage of Insertion sort's linear behavior on nearly sorted lists. Generally this is implemented by having quicksort stop when it gets down to lists of less than ten or fifteen elements, and then Insertion sorting the whole list at the end.
- Remove recursion and use a stack to keep track of sub-lists yet to be sorted.

This removes function calling overhead.

- Partition smaller sub-lists first, which keeps the stack a little smaller..

Even without these further refinements, empirical studies have shown that quicksort is, on average, about twice as fast as any other $O(n \lg n)$ sorting algorithm.. This is because there is minimal overhead in the partitioning algorithm loops, and because quicksort does relatively few swaps (look at all the assignments that merge sort does by comparison, for example)..

```
def quick_m3(array, lb, ub)
  return if ub <= lb

  # find sentinels and the median for the
  pivot m = (lb+ub)/2
  array[lb],array[m]=array[m],array[lb] if array[m] < array[lb]
  array[m],array[ub]=array[ub],array[m] if array[ub] < array[m]
  array[lb],array[m]=array[m],array[lb] if array[m] < array[lb]

  # if the sub-array is size 3 or less, it is now sorted
  return if ub-lb < 3

  # put the median just shy of the end of the list
  array[ub-1], array[m] = array[m], array[ub-1]

  pivot = array[ub-1]
  i, j = lb, ub-1
  loop do
    loop do i += 1; break if pivot <= array[i]; end
    loop do j -= 1; break if j <= lb || array[j] <= pivot;
    end array[i], array[j] = array[j], array[i] break if j <=
    i
  end
  array[j], array[i], array[ub-1] = array[i], pivot, array[j]
  quick_m3(array,lb,i-1)
  quick_m3(array,i+1,ub)
end
private :quick_m3

def quicksort_m3(array)
  quick_m3(array, 0, array.size-1)
  return array
end
```

Figure 3: Quicksort with the Median-of-Three Improvement

Merge Sort and Quicksort

Summary and Conclusion

Merge sort is a fast sorting algorithm whose best, worst, and average case complexity are all in $O(n \lg n)$, but unfortunately it uses $O(n)$ extra space to do its work.. Quicksort has best and average case complexity in $O(n \lg n)$, but unfortunately its worst case complexity is in $O(n^2)$.. The median-of-three improvement makes quicksort's worst case behavior extremely unlikely, however, and quicksort's unmatched speed in practice make it the sorting algorithm of choice when sorting by comparison of keys..

Review Questions

1. Why does merge sort need extra space?
2. What stops recursion in merge sort and quicksort?
3. What is a pivot value in quicksort?
4. What changes have been suggested to improve quicksort?
5. If quicksort has such bad worst case behavior, why is it still used so widely?

Exercises

1. Explain why the merge sort algorithm first copies the original list into the auxiliary array..
2. Write a non-recursive version of merge sort that uses a stack to keep track of sub-lists that have yet to be sorted.. Time your implementation against the unmodified merge sort algorithm and summarize the results..
3. The merge sort algorithm presented above does not take advantage of all the features of Ruby.. Write a version of merge sort that uses features of Ruby to make it simpler or faster than the version in the text.. Time your algorithm against the one in the text to determine which is faster..
4. Modify the quicksort algorithm with the median-of-three improvement so that it does not sort lists smaller than a dozen elements, and calls insertion sort to finish sorting at the end.. Time your implementation against the unmodified quicksort with the median-of-three improvement and summarize the results..
5. Modify the quicksort algorithm with the median-of-three improvement so that it uses a stack rather than recursion, and works on smaller sub-lists first.. Time your implementation against the unmodified quicksort with the median-of-three improvement and summarize the results..
6. The quicksort algorithm presented above does not take advantage of all the features of Ruby.. Write a version of quicksort that takes advantage of Ruby's features to make the algorithm shorter or faster.. Time your version against the one in the text to determine which one is faster..
7. Write the fastest quicksort you can.. Time your implementation against the unmodified quicksort and summarize the results..

Merge Sort and Quicksort

8. The Ruby `Array` class has a `sort()` method that uses a version of the quicksort algorithm.. If you time any of the quicksort algorithms we have presented, or one you write yourself, against the `Array.sort()` method, you will find that the latter is much faster.. Why is this so?

Review Question Answers

1. Merge sort uses extra space because it is awkward and slow to merge lists without using extra space..
2. Recursion in merge sort and quicksort stops when the sub-lists being sorted are either empty or of size one—such lists are already sorted, so no work needs to be done on them..
3. A pivot value in quicksort is an element of the list being sorted that is chosen as the basis for rearranging (partitioning) the list: all elements less than the pivot are placed to the left of it, and all elements greater than the pivot are placed to the right of it.. (Equal values may be placed either left or right of the pivot, and different partitioning algorithms may make different choices)..
4. Among the changes that have been suggested to improve quicksort are (a) using the median of the first, last, and middle elements in the list as the pivot value, (b) using insertion sort for small sub-lists, (c) removing recursion in favor of a stack, and (d) sorting small sub-lists first to reduce the depth of recursion (or the size of the stack)..
5. Quicksort is still used widely because its performance is so good on average: quicksort usually runs in about half the time of other sorting algorithms, especially when it has been improved in the ways discussed in the chapter.. Its worst case behavior is relatively rare if it incorporates the median-of-three improvement.. Innovations like Introspective sort (see exercise 6) can even deal effectively with quicksort's worst case behavior without sacrificing its amazing speed..

Trees, Heaps, and Heapsort

Introduction

Trees are the basis for several important data types and data structures.. There are also several sorting algorithms based on trees.. One of these algorithms is heapsort, which uses a complete binary tree represented in an array for fast in-place sorting..

Basic Terminology

A trees is a special type of graph..

Graph: A collection of *vertices*, or *nodes*, and *edges* connecting the nodes.. An edge may be thought of a pair of vertices.. Formally, a graph is an ordered pair $\langle V, E \rangle$ where V is a set of vertices, and E is a set of pairs of elements of V ..

Simple path: A list of distinct vertices such that successive vertices are connected by edges..

Tree: A graph with a distinguished vertex r , called the *root*, such that there is exactly one simple path between each vertex in the tree and r ..

We usually draw trees with the root at the top and the nodes and edges descending below it.. Figure 1 illustrates a tree..

r




Figure 1: A Tree

Node r is the root.. The root has three *children*: a , b , and c .. The root is the *parent* of these vertices.. These vertices are also *siblings* of one another because they have the same parent Vertex a has child d , and vertex c has children e and f .. The ancestors of a vertex are the vertices on the path between it and the root; the *descendents* of a node are all the nodes of which it is an ancestor.. Thus node f has ancestors f , c , and r , and c has descendents c , e , and f .. A vertex without children is a *terminal node* or a *leaf*; those with children are *non-terminal nodes* or *internal nodes*.. The tree in Figure 1 has three internal nodes (r , a , and c), and four leaf nodes (b , d , e , and f).. The graph consisting of a vertex in a tree, all its descendents, and the edges connecting them, is a *subtree* of the tree..

A graph consisting of several trees is a *forest*.. The *level* of a node in a tree is the number of nodes in the path from the node to the root, not including itself.. In Figure 1, node r is at level 0, nodes a , b , and c are at level 1, and nodes d , e , and f are at level 2.. The *height* of a tree is the maximum level in the tree.. The height of the tree in Figure 1 is 2..

An *ordered tree* is one in which the order of the children of each node is specified.. Ordered trees are not drawn in any special way—some other mechanism must be used to specify whether a tree is ordered..

Binary Trees

Binary trees are especially important for making data structures..

Binary tree: An ordered tree whose vertices have at most two children.. The children are distinguished as the *left child* and *right child*.. The subtree whose root is the left (right) child of a vertex is the *left (right) subtree* of that vertex..

A *full binary tree* is one in which every level is full, except possibly the last.. A *complete binary tree* is a full binary tree in which only the right-most nodes at the bottom level are missing.. Figure 2 illustrates these notions..




Figure 2: Binary Trees

The trees in Figure 2 are binary trees.. In the tree on the left, node *a* has a left child, node *b* has a right child, and node *c* has no children.. This tree is neither full nor complete.. The middle tree is full but not complete, and the right tree is complete..

Trees have several interesting and important properties, the following among them..

- A tree with n nodes has $n-1$ edges..
- A complete binary tree with n internal nodes has either n or $n+1$ leaves..
- The height of a full binary tree with n nodes is $\text{floor}(\lg n)$..

Heaps

A node in a binary tree has the *heap-order property* if the value stored at the node is greater than or equal to the values stored at its descendants..

Heap: A complete binary tree whose every node has the heap-order property..

An arbitrary complete binary tree can be made into a heap quite easily as follows:

- Every leaf already has the heap-order property, so the subtrees whose roots are leaves are heaps..
- Starting with the right-most internal node v at the next-to-last level, and working left across levels and upwards in the tree, do the following: if node v does not have

the heap-order property, swap its value with the largest of its children, then do the same with the modified node, until the subtree rooted at v is a heap..

It is fairly efficient to make complete binary trees into heaps because each subtree is made into a heap by swapping its root downwards in the tree as far as necessary.. The height of a complete binary tree is $\text{floor}(\lg n)$, so this operation cannot take very long..

Heaps can be implemented in a variety of ways, but the fact that they are complete binary trees makes it possible to store them very efficiently in contiguous memory locations.. Consider the numbers assigned to the nodes of the complete binary tree in Figure 3.. Note that numbers are assigned left to right across levels, and from top to bottom of the tree..




Figure 3: Numbering Nodes for Contiguous Storage

This numbering scheme can be used to identify each node in a complete binary tree: node zero is the root, node one is the left child of the root, node two is the right child of the root, and so forth.. Note in particular that

- The left child of node k is node $2k+1$..
- The right child of node k is node $2k+2$..
- The parent node k is node $\text{floor}((k-1)/2)$..
- If there are n nodes in the tree, the last one with a child is node $\text{floor}(n/2) - 1$..

Now, if we let these node numbers be array indices, then each array location is associated with a node in the tree, and we can store the values at the nodes of the tree in the array: the value of node k is stored in array location k .. The correspondence between array indices and node locations thus makes it possible to represent complete binary trees in arrays.. The fact that the binary tree is complete means that every array location stores the value at a node, so no space is unused in the array..

Heapsort

We now have all the pieces we need to for an efficient and interesting sorting algorithm based on heaps.. Suppose we have an array to be sorted.. We can consider it to be a complete binary tree stored in an array as explained above.. Then we can

- Make the tree into a heap as explained above..
- The largest value in a heap is at the root, which is always be at array location zero.. We can swap this value with the value at the end of the array and pretend the array is one element shorter.. Then we have a complete binary tree that is almost a heap—we just need to sift the root value down the tree as far as necessary to make it one.. Once we do, the tree will once again be a heap..
- We can then repeat the process again and again until the entire array is sorted..

This sorting algorithm, called *heapsort*, is shown in the Ruby code in Figure 4 below..

```

def heapify(array, i, max_index)
  tmp = array[i]
  j = 2*i + 1
  while j <= max_index
    j += 1 if j < max_index && array[j] < array[j+1]
    break if array[j] <= tmp
    array[i] = array[j]

    i, j = j, 2*i + 1
  end
  array[i] = tmp
end

def heap_sort(array)
  # make the entire array into a heap
  max_index = array.size-1 ((max_index-
  1)/2).downto(0).each do | i |
    heapify(array,i,max_index)
  end

  # repeatedly remove the root and remake the
  heap loop do
    array[0],array[max_index] = array[max_index],array[0]
    max_index -= 1
    break if max_index <= 0
    heapify(array, 0, max_index)
  end
  return array
end

```

Figure 4: Heapsort

A somewhat complex analysis that we will not reproduce here shows that the number of comparisons done by heapsort in both the best, worst, and average cases are all in $O(n \lg n)$.. Thus heapsort joins merge sort and quicksort in our repertoire of fast sorting algorithms..

Trees, Heaps, and Heapsort

Empirical studies have shown that while heapsort is not as fast as quicksort, it is not much slower than merge sort, with the advantage that it does not use any extra space, and it does not have bad worst case complexity..

Summary and Conclusion

A tree is a special sort of graph that is important in computing.. One application of trees is for sorting: an array can be treated as a complete binary tree and then transformed into a heap.. The heap can then be manipulated to sort the array in place in $O(n \lg n)$ time.. This algorithm is called heapsort, and is a good algorithm to use when space is at a premium and respectable worst case complexity is required..

Review Questions

1. In Figure 1, what are the descendents of r ? What are the ancestors of r ?
2. How can you tell from a diagram whether a tree is ordered?
3. Is every full binary tree and a complete binary tree? Is every complete binary tree a full binary tree?
4. Where is the largest value in a heap?
5. Using the heap data structure numbering scheme, which nodes are the left and right children of node 27? Which node is the parent of node 27?
6. What is the worst case behavior of heapsort?

Exercises

1. Represent the three trees in Figure 2 as sets of ordered pairs according to the definition of a graph..
2. Could the graph in Figure 1 still be a tree if b was its root? If so, redraw the tree in the usual way (that is, with the root at the top) to make clear the relationships between the nodes..
3. Draw a complete binary tree with 12 nodes, placing arbitrary values at the nodes.. Use the algorithm discussed in the chapter to transform the tree into a heap, redrawing the tree at each step..
4. Suppose that we change the definition of the heap-order property to say that the value stored at the node is less than or equal to the values stored at its descendents.. If we use the heapsort algorithm on trees that are heaps according to this definition, what will be the result?
5. In the heapsort algorithm in Figure 4, the `heapify()` operation is applied to nodes starting at `max_index-1`.. Why does the algorithm not start at `max_index`?
6. Draw a complete binary tree with 12 node, placing arbitrary values at the nodes.. Use the heapsort algorithm to sort the tree, redrawing the tree at each step, and placing removed values into a list representing the sorted array as they are removed from the tree..

Trees, Heaps, and Heapsort

7. Write a program to sort arrays of various sizes using heapsort, merge sort, and quicksort.. Time your implementations and summarize the results..
8. Introspective sort is a quicksort-based algorithm recently devised by David Musser.. Introspective sort works like quicksort except that it keeps track of the depth of recursion (or of the stack), and when recursion gets too deep (about $2 \cdot \lg n$ recursive calls), it switches to heapsort to sort sub-lists.. This algorithm does $O(n \lg n)$ comparisons even in the worst case, sorts in place, and usually runs almost as fast as quicksort on average.. Write an introspective sort operation, time your implementation against standard quicksort, and summarize the results..

Review Question Answers

1. In Figure 1 the descendents of r are all the nodes in the tree.. Node r has no ancestor except itself..
2. You can't tell from a diagram whether a tree is ordered; there must be some other notation to indicate that this is the case..
3. Not every full binary tree is complete because all the leaves of a tree might be on two levels, making it full, but some of the missing leaves at the bottom level might not be on the right, meaning that it is not complete.. Every complete binary tree must be a full binary tree, however..
4. The largest value in a heap is always at the root..
5. Using the heap data structure numbering scheme, the left child of node 27 is node $(2 \cdot 27) + 1 = 55$, the right children of node 27 is node $(2 \cdot 27) + 2 = 56$, and the parent of node 27 is node $\text{floor}((27 - 1) / 2) = 13$..
6. The worst, best, and average case behavior of heapsort is in $O(n \lg n)$..

Binary Trees

Introduction

As mentioned in the last chapter, binary trees are ordered trees whose nodes have at most two children, the left child and the right child.. Although other kinds of ordered trees arise in computing, binary trees are especially common and have been especially well studied.. In this chapter we discuss the binary tree abstract data type and binary trees as an implementation mechanism..

The Binary Tree ADT

Binary trees hold values of some type, so the ADT is *binary tree of T* , where T is the type of the elements in the tree.. The carrier set of this type is the set of all binary trees whose nodes hold a value of type T .. The carrier set thus includes the empty tree, the trees with only a root holding a value of type T , the trees with a root and a left child, the trees with a root and a right child, and so forth.. Operations in this ADT include the following..

size(t)—Return the number of nodes in the tree t ..

height(t)—Return the height of tree t ..

empty?(t)—Return true just in case t is the empty tree..

contains?(t,v)—Return true just in case the value v is present in tree t ..

buildTree(v,t_l,t_r)—Create and return a new binary tree whose root holds the value v and whose left and right subtrees are t_l and t_r ..

emptyTree()—Return the empty binary tree..

rootValue(t)—Return the value of type T stored at the root of the tree t .. Its precondition is that t is not the empty tree..

leftSubtree(t)—Return the tree whose root is the left child of the root of t ..

Its precondition is that t is not the empty tree..

rightSubtree(t)—Return the tree whose root is the right child of the root of t ..

Its precondition is that t is not the empty tree..

This ADT allows us to create arbitrary binary trees and examine them.. For example, consider the binary tree in Figure 1..

a




Figure 1: A Binary Tree

Binary Trees

This tree can be constructed using the expression below..

```
buildTree(a,  
          buildTree(b,  
                    emptyTree(),  
                    buildTree(d,  
                                emptyTree(),  
                                emptyTree()))),  
          buildTree(c,  
                    emptyTree(),  
                    emptyTree()))
```

To extract a value from the tree, such as the bottom-most node *d*, we could use the following expression, where *t* is the tree in Figure 1..

```
rootValue(rightSubtree(leftSubtree(t)))
```

As with the ADTs we have studied before, an object-oriented implementation of these operations as instance methods will include the tree as an implicit parameter, so the signatures of these operations vary somewhat when they are implemented.. Furthermore, there are several operations that are very useful for a binary tree implementation that are not present in the ADT, and several operations in the ADT that are not needed (more about this below)..

The Binary Tree Class

We could treat binary trees as a kind of collection, adding it to our container hierarchy, but we won't do this for two reasons:

- In practice, binary trees are used to implement other collections, not as collections in their own right.. Usually clients are interested in using basic Collection operations, not in the intricacies of building and traversing trees.. Adding binary trees to the container hierarchy would complicate the hierarchy with a container that not many clients would use..
- Although binary trees have a contiguous implementation (discussed below), it is not useful except for heaps.. Providing such an implementation in line with our practice in the container hierarchy to make both contiguous and linked implementations for all interfaces would create a class without much use..

We will make a BinaryTree class, but its role will be to provide an implementation mechanism for other collections.. Thus the BinaryTree does not implement the Collection interface, though it is convenient for it to have several of standard collection operations.. It also includes operations for creating and traversing trees in various ways, as well as several kinds of iterators.. The Binary Tree class is pictured in Figure 2..

Note that there is no *buildTree()* operation and no *emptyTree()* operation in the BinaryTree class, though there is one in the ADT.. The BinaryTree class constructor does the job of these two operations, so they are not needed as separate operations in the class..

Binary Trees

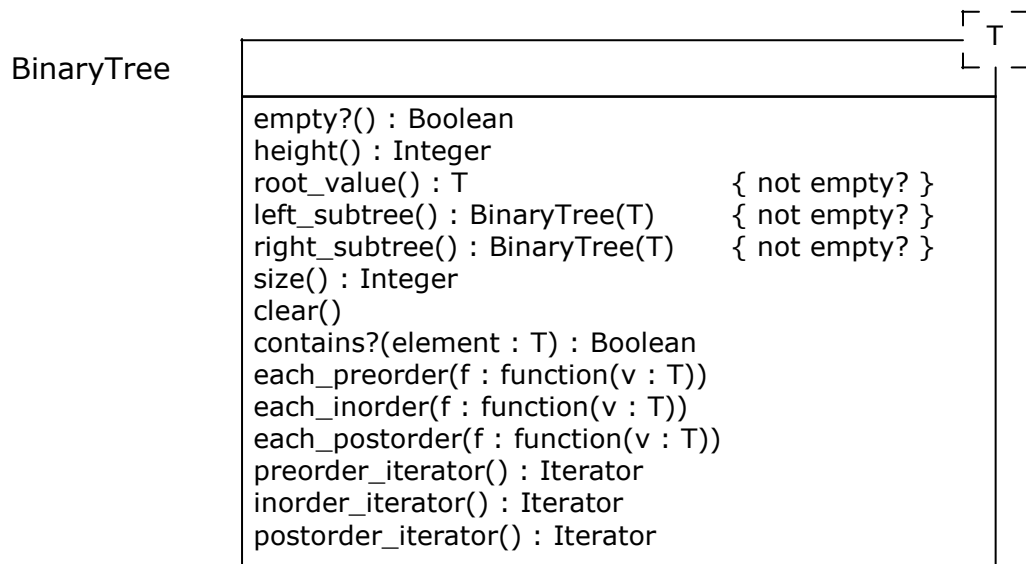


Figure 2: The BinaryTree Class

To *visit* or *enumerate* the nodes of a binary tree is to traverse or iterate over them one at a time, processing the values held in each node.. This requires that the nodes be traversed in some order.. There are three fundamental orders for traversing a binary tree.. All are most naturally described in recursive terms..

Preorder: When the nodes of a binary tree are visited in preorder, the root node of the tree is visited first, then the left sub-tree (if any) is visited in preorder, then the right sub-tree (if any) is visited in preorder..

Inorder: When the nodes of a binary tree are visited inorder, the left sub-tree (if any) is visited inorder, then the root node is visited, then the right sub-tree is visited inorder..

Postorder: When the nodes of a binary tree are visited in postorder, the left sub-tree is visited in postorder, then the right sub-tree is visited in postorder, and then the root node is visited..

To illustrate these traversals, consider the binary tree in Figure 3 below.. An inorder traversal of the tree in Figure 3 visits the nodes in the order *m, b, p, k, t, d, a, g, c, f, h*.. A preorder traversal visits the nodes in the order *d, b, m, k, p, t, c, a, g, f, h*.. A postorder traversal visits the nodes in the order *m, p, t, k, b, g, a, h, f, c, d*..

The BinaryTree class has internal iterators for visiting the nodes of the tree in the three orders listed above and applying the function passed in as an argument to each node of the tree.. For examples, suppose that a print(v : T) operation prints the value v.. If t is a binary tree, then the call t.each_inorder(print) will cause the values in the tree t to be printed out inorder, the call t.each_preorder(print) will cause them to be printed in preorder, and the call t.each_postorder(print) will cause them to be printed in postorder.. In addition, the BinaryTree class has three operations that return external iterators that provide access to the values in the tree in each of the three orders above..

Binary Trees




Figure 3: A Binary Tree

When implementing the `BinaryTree` class in Ruby, it will mix in `Enumerator`, and the `Enumerator.each()` operation will be an alias for the `each_inorder()`, which is the most common way to visit the nodes of a binary tree..

Contiguous Implementation of Binary Trees

We have already considered how to implement binary trees using an array when we learned about heapsort.. The contiguous implementation is excellent for complete or even full binary trees because it wastes no space on pointers, and it provides a quick and easy way to navigate in the tree.. Unfortunately, in most applications binary trees are far from complete, so many array locations are never used, which wastes a lot of space.. Even if our binary trees were always full, there is still the problem of having to predict the size of the tree ahead of time so that an array could be allocated that is big enough to hold all the tree nodes.. The array could be reallocated if the tree becomes too large, but this is an expensive operation..

This is why it is not particularly useful to have a contiguous implementation of binary trees in our container hierarchy, or even to build one as an implementation mechanism used to implement other collections.. Instead we will implement our `BinaryTree` class as a linked data structure, and use it as the linked structure implementation mechanism for several of the collections we will add to our container hierarchy later on..

Linked Implementation of Binary Trees

A linked implementation of binary trees resembles implementations of other ADTs using linked structures.. Binary tree nodes are represented by objects of a `BinaryTreeNode` class that has attributes for the data held at the node and references to the left and right subtrees.. Each `BinaryTree` object holds a reference of type `BinaryTreeNode` to the tree's root node.. Empty trees are represented as null references.. Figure 4 illustrates this approach..

Trees are inherently recursive structures, so it is natural to write many `BinaryTree` class operations recursively.. For example, to implement the `size()` operation, the `BinaryTree` can call an internal `size(r : BinaryTreeNode)` operation on the root node.. This operation returns zero if its parameter is null, and one plus the sum of recursive calls on the left and right sub-

Binary Trees




Figure 4: A Linked Representation of a Binary Tree

trees of its parameter node.. Many other operations, and particularly the traversal operations that apply functions to the data held at each node, can be implemented just as easily..

Implementing iterators is more challenging, however.. The problem is that iterators cannot be written recursively because they have to be able to stop every time a new node is visited to deliver the value at the node to the client.. There are two ways to solve this problem:

- Write a recursive operation to copy node values into a data structure (like a queue) in the correct order, and then extract items from the data structure one at a time as the client requests them..
- Don't user recursion to implement iterators: use a stack instead..

The second alternative, though harder to do, it clearly better because it uses much less space.. **Summary and Conclusion**

The binary tree ADT describes basic operations for building and examining binary trees whose nodes hold values of type T .. A `BinaryTree` class has several operations not in the ADT, in particular, visitor operations for traversing the nodes of the tree and applying a function to the data stored in each node.. Iterators are also made available by this class..

Contiguous implementations of the binary tree ADT are possible, and useful in some special circumstances, such as in heapsort, but the main technique for implementing the binary tree ADT uses a linked representation.. Recursion is a very useful tool for implementing most `BinaryTree` operations, but it cannot be used as easily for implementing iterators.. The `BinaryTree` class implemented using a linked structure will be used as an implementation mechanism for container classes to come..

Review Questions

1. Where does the `BinaryTree` class fit in the Container class hierarchy?
2. Why does the `BinaryTree` class not include a `buildTree()` operation?


Binary Trees

3. Why is the contiguous implementation of binary trees not very useful?
4. What is the relationship between the BinaryTree and BinaryTreeNode classes?

Exercises

1. Write the values of the nodes in the following tree in the order they are visited when the tree is traversed inorder, in preorder, and in postorder..

g



2. Write the size() operation for BinaryTree class in Ruby..
3. Write the height() operation for the BinaryTree class in Ruby..
4. Write the each_preorder(), each_inorder(), and each_postorder() operations for the BinaryTree class in internal iterator operations in Ruby..
5. Write a PreorderIterator class whose instances are iterators over for the BinaryTree class.. The PreorderIterator class will need a stack attribute to hold the nodes that have not yet been visited, with the node holding the current node during iteration at the top of the stack..

Review Question Answers

1. We have decided not to include the BinaryTree class in the Container class hierarchy because it is usually not used as a container in its own right, but rather as an implementation mechanism for other containers..
2. The BinaryTree class does not include a buildTree() operation because it may have a constructor that does the very same job..
3. The contiguous implementation of binary trees is not very useful because it only uses space efficiently if the binary tree is at least full, and ideally complete.. In practice, this is rarely the case, so the linked implementation uses space more efficiently..

Binary Trees

4. The `BinaryTree` class has an attribute that stores a reference to the root of the tree, which is a `BinaryTreeNode` instance.. The root node (if any) stores references to the left and right sub-trees of the root, which are also references to instances of the `BinaryTreeNode` class.. Although the `BinaryTree` and `BinaryTreeNode` classes are not related by inheritance, they are intimately connected, just as the `LinkedList` class is closely connected to the `LinkedListNode` class..

Binary Search and Binary Search Trees

Introduction

Binary search is a much faster alternative to sequential search for sorted lists.. Binary search is closely related to binary search trees, which are a special kind of binary tree.. We will look at these two topics in this chapter, studying the complexity of binary search, and eventually arriving at a specification for a `BinarySearchTree` class..

Binary Search

When people search for something in an ordered list (like a dictionary or a phone book), they do not start at the first element and march through the list one element at a time.. They jump into the middle of the list, see where they are relative to what they are looking for, and then jump either forward or backward and look again, continuing in this way until they find what they are looking for, or determine that it is not in the list..

Binary search takes the same tack in searching for a key value in a sorted list: the key is compared with the middle element in the list.. If it is the key, the search is done; if the key is less than the middle element, then the process is repeated for the first half of the list; if the key is greater than the middle element, then the process is repeated for the second half of the list.. Eventually, either the key is found in the list, or the list is reduced to nothing (the empty list), at which point we know that the key is not present in the list..

This approach naturally lends itself to a recursive algorithm, which we show coded in Ruby below..

```
def rb_search(array, key)
  return nil if array.empty?
  m = array.size/2;
  return m if key == array[m]
  return rb_search(array[0..m],key) if key < array[m]
  index = rb_search(array[m+1..-1],key) index ?
  m+1+index : nil
end
```

Figure 1: Recursive Binary Search

Search algorithms traditionally return the index of the key in the list, or -1 if the key is not found; in Ruby we have a special value for undefined results, so we return `nil` if the key is not in the array.. Note also that although the algorithm has the important precondition that the array is sorted, checking this would take far too much time, so it is not checked..

The recursion stops when the array is empty and the key has not been found.. Otherwise, the element at index `m` in the middle of the array is checked.. If it is the key, the search is done and index `m` is returned; otherwise, a recursive call is made to search the portion of the list before or after `m` depending on whether the key is less than or greater than `array[m]`..

Binary Search and Binary Search Trees

Although binary search is naturally recursive, it is also tail recursive.. Recall that a tail recursive algorithm is one in which at most one recursive call is made in each activation of the algorithm, and that tail recursive algorithms can always be converted to non-recursive algorithms using only a loop and no stack.. This is always more efficient and often simpler as well.. In the case of binary search, the non-recursive algorithm is about equally complicated, as the Ruby code in Figure 2 below shows..

```
def binary_search(array, key)
  lb, ub = 0, array.size-1
  while (lb <= ub)
    m = (ub+lb)/2;
    return m if key == array[m]
    if key < array[m]
      ub = m-1
    else
      lb = m+1
    end
  end
  return nil
end
```

Figure 2: Non-Recursive Binary Search

To analyze binary search, we will consider its behavior on lists of size n and count the number of comparisons between list elements and the search key.. Traditionally, the determination of whether the key is equal to, less than, or greater than a list element is counted as a single comparison, even though it may take two comparisons in most programming languages..

Binary search does not do the same thing on every input of size n .. In the best case, it finds the key at the middle of the list, doing only 1 comparison.. In the worst, case, the key is not in the list, or is found when the sub-list being searched has only one element.. We can easily generate a recurrence relation and initial conditions to find the worst case complexity of binary search, but we will instead use a binary search tree to figure this out..

Suppose that we construct a binary tree from a sorted list as follows: the root of the tree is the element in the middle of the list; the left child of the root is the element in the middle of the first half of the list; the right child of the root is the element in the middle of the second half of the list, and so on.. In other words, the nodes of binary tree are filled according to the order in which the values would be encountered during a binary search of the list.. To illustrate, consider the binary tree in Figure 3 made out of the list $\langle a, b, c, d, e, f, g, h, i, j, k, l \rangle$ in the way just described..

Binary Search and Binary Search Trees

g




Figure 3: A Binary Tree Made from a

List A tree built this way has the following interesting properties:

- It is a full binary tree, so its height is always $\text{floor}(\lg n)$.
- For every node, every element in its left sub-tree (if any) is less than or equal to the element at the node, and every element in its right sub-tree (if any) is greater than or equal to the element at the node.
- If we traverse the tree inorder, we visit the node elements in the order of the original list, that is, in sorted order.

The first property tells us the worst case performance of binary search because a binary search will visit each node from the root to a leaf in the worst case.. The number of nodes on these paths is the height of the tree plus one, so $W(n) = \text{floor}(\lg n) + 1$. We can also calculate the average case by considering each node equally likely to be the target of a binary search, and figuring out the average length of the path to each node.. This turns out to be approximately $\lg n$ for both successful and unsuccessful searches.. Hence, on average and in the worst case, binary search makes $O(\lg n)$ comparisons, which is very good..

Binary Search Trees

The essential characteristic of the binary tree we looked at above is the relationship between the values at a node and the values in its left and right sub-trees.. This is the basis for the definition of binary search trees..

Binary search tree: A binary tree whose every node is such that the value at each node is greater than the values in its left sub-tree, and less than the values in its right sub-tree..

Binary search trees are an important data type that retains the property that traversing them in order visits the values in the nodes in sorted order.. However, a binary search tree may not be full, so its height may be greater than $\text{floor}(\lg n)$.. In fact, a binary search tree whose every node but one has only a single child will have height $n-1$..

Binary search trees are interesting because it is fast both to insert elements into them and fast to search them (provided they are not too long and skinny).. This contrasts with most collections, which are usually fast for insertions but slow for searches, or vice versa.. For example, elements can be inserted into an (unsorted) LinkedList quickly, but searching a LinkedList is slow, while a (sorted) ArrayList can be searched quickly with binary search, but inserting elements into it to keep it sorted is slow..

The *binary search tree of T* ADT has as its carrier set the set of all binary search trees whose nodes hold a value of type *T*.. It is thus a subset of the carrier set of the binary tree of *T* ADT.. The operations in this ADT includes all the operations of the binary tree ADT, with the addition of a precondition on *buildTree()*, shown below.. The list below also includes two operations added to the binary search tree ADT..

buildTree(v,t_l,t_r)—Create and return a new binary tree whose root holds the value *v* and whose left and right subtrees are *t_l* and *t_r*.. Its precondition is that *v* is greater than any value held in *t_l* and less than any value held in *t_r*..

add(t,v)—Put *v* into a new node added as a leaf to *t*, preserving the binary search tree property, and return the resulting binary search tree.. If *v* is already in *t*, then *t* is unchanged..

remove(t,v)—Remove the node holding *v* from *t*, if any, while preserving the result as a binary search tree, and return the resulting binary search tree..

This ADT is the basis for a BinarySearchTree class..

The Binary Search Tree Class

A BinarySearchTree is a kind of BinaryTree, so the BinarySearchTree class is a sub-class of BinaryTree. Its constructor needs a precondition to make sure that trees are constructed properly.. It can also override the contains() operation to be more efficient.. Otherwise, it only needs to implement the three operations pictured in Figure 4 below..

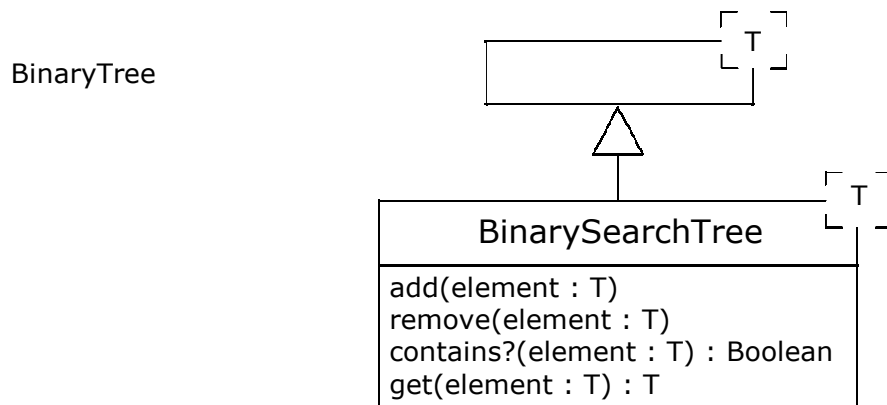


Figure 4: The BinarySearchTree Class

The `add()` operation puts an element into the tree by making a new child node at a spot that preserves the binary search tree's integrity.. If the element is already in the tree, then the element passed in replaces the value currently stored in the tree.. In this way, a new record can replace an old one with the same key (more about this in later chapters)..

Binary Search and Binary Search Trees

The `get()` operation returns the value stored in the tree that is “equal” to the element sent in.. It is intended to fetch a record from the tree with the same key as a dummy record supplied as an argument, thus providing a retrieval mechanism (again, we will discuss this more later)..

The `contains()` and `get()` operations both search the tree by starting at its root and moving down the tree, mimicking a binary search.. If the desired value is at the root (the best case), this requires only one comparison, so $B(n)$ is in $O(1)$.. In the worst case, when the tree is effectively a list, this requires $O(n)$ comparisons.. Empirical studies have shown that when binary search trees are built by a series of insertions of random data, they are more or less bushy, and their height is not too much more than $\lg n$, so the number of comparisons is in $O(\lg n)$ on average..

The `add()` operation takes a path down the tree to the spot where the new element would be found during a search, and adds a new leaf node to hold it.. This again requires $O(1)$ operations in the best case, $O(\lg n)$ comparisons in the average case, and $O(n)$ comparisons in the worst case.. Finally, the `remove()` operation must first find the deleted element, and then manipulate the tree to remove the node holding the element in such a way that it is preserved as a binary search tree.. This operation also takes $O(1)$ time in the best case, $O(\lg n)$ time in the average case, and $O(n)$ time in the worst case..

Binary search trees thus provide very efficient operations except in the worst case.. There are several kinds of balanced binary search trees whose insertion and deletion operations keep the tree bushy rather than long and skinny, thus eliminating the worst case behavior.. We will not have time to study balanced binary search trees..

Summary and Conclusion

Binary search is a very efficient algorithm for searching ordered lists, with average and worst case complexity in $O(\lg n)$.. We can represent the workings of binary search in a binary tree to produce a full binary search tree.. Binary search trees have several interesting properties and provide a kind of collection that features excellent performance for addition, deletion, and search, except in the worst case.. We can also traverse binary search trees inorder to access the elements of the collection in sorted order..

Review Questions

1. Why can recursion be removed from the binary search algorithm without using a stack?
2. If a binary tree is made from an ordered list of 100 names by placing them into the tree to mimic a binary search as discussed in the text, what is the height of the resulting tree?
3. Approximately how many comparisons would be made by binary search when searching a list of one million elements in the best, worst, and average cases?
4. What advantage does a binary search tree have over collections like `ArrayList` and `LinkedList`?

Hashing

Introduction

In an ideal world, retrieving a value from a map would be done instantly by just examining the value's key.. That is the goal of hashing, which uses a hash function to transform a key into an array index, thereby providing instantaneous access to the value stored in an array holding the key-value pairs in the map.. This array is called a hash table..

Hash function: A function that transforms a key into a value in the range of indices of a hash table..

Hash table: An array holding key-value pairs whose locations are determined by a hash function..

Of course, there are a few details to work out..

The Hashing Problem

If a set of key-value pairs is small and we can allocate an array big enough to hold them all, we can always find a hash function that transforms keys to unique locations in a hash table.. For example, in some old programming languages, identifiers consisted of an upper-case letter possibly followed by a digit.. Suppose these are our keys.. There are 286 of them, and it is not too hard to come up with a function that maps each key of this form to a unique value in the range 0.. .285.. But usually the set of keys is too big to make a table to hold all the possibilities.. For example, older versions of FORTRAN had identifiers that started with an upper-case letter, followed by up to five additional upper-case letters or digits.. The number of such identifiers is 1,617,038,306, which is clearly too big for a hash table if we were to use these as keys..

A smaller table holding keys with a large range of values will inevitably require that the function transform several keys to the same table location.. When two or more keys are mapped to the same table location by a hash function we have a collision.. Mechanisms for dealing with them are called *collision resolution schemes*..

Collision: The event that occurs when two or more keys are transformed to the same hash table location..

How serious is the collision problem? After all, if we have a fairly large table and a hash function that spreads keys out evenly across the table, collisions may be rare.. In fact, however, collisions occur surprisingly often.. To see why, lets consider the *birthday problem*, a famous problem from probability theory: what is the chance that a least two people in a group of k people have the same birthday? This turns out to be $p = 1 - (365! / (k! \cdot 365^k))$.. Table 1 below lists some values for this expression.. Surprisingly, in a group of only 23 people there is better than an even chance that two of them have the same birthday!

If we imagine that a hash table has 365 locations, and that these probabilities are the likelihoods that a hash function transforms two values to the same location (a collision), then we can see that we are almost certain to have a collision when the table holds 100 values, and very likely to have a collision with only about 40 values in the table.. Forty is

only about 11% of 365, so we see that collisions are very likely indeed.. Collision resolution schemes are thus an essential part of making hashing work in practice..

| <i>k</i> | <i>p</i> |
|----------|------------|
| 5 | 0..027 |
| 10 | 0..117 |
| 15 | 0..253 |
| 20 | 0..411 |
| 22 | 0..476 |
| 23 | 0..507 |
| 25 | 0..569 |
| 30 | 0..706 |
| 40 | 0..891 |
| 50 | 0..970 |
| 60 | 0..994 |
| 100 | 0..9999997 |

Table 1: Probabilities in the Birthday

Problem An implementation of hashing thus requires two things:

- A hash function for transforming keys to hash table locations, ideally one that makes collisions less likely..
- A collision resolution scheme to deal with the collisions that are bound to occur.. We discuss each of these in turn..

occur.. We discuss each of these in turn..

Hash Functions

A hash function must transform a key into an integer in the range $0 \dots t$, where t is the size of the hash table.. A hash function should distribute the keys in the table as uniformly as possible to minimize collisions.. Although many hash functions have been proposed and investigated, the best hash functions use the division method, which for numeric keys is the following..

$$\text{hash}(k) = k \bmod t$$

This function is simple, fast, and spreads out keys uniformly in the table.. It works best when t is a prime number not close to a power of two.. For this reason, hash table sizes should always be chosen to be a prime number not close to a power of two..

For non-numeric keys, there is usually a fairly simple way to convert the value to a number and then use the division method on it.. For example, the following pseudocode illustrates a way to hash a string..

Hashing

```
def hash_function(string, table_size)
  result = 0
  string.each_byte do |byte|
    result = (result * 151 + byte) %
    table_size end
  return result
end
```

Figure 1: Hash Function for Strings

Thus, making hash functions is not too onerous.. A good rule of thumb is to use prime numbers whenever a constant is needed, and to test the function on a representative set of keys to ensure that it spreads them out evenly across the hash table..

Collision Resolution Schemes

There are two main kinds of collision resolution schemes, with many variations: chaining and open addressing.. In each scheme, an important value to consider is the load factor, $\lambda = n/t$, where n is the number of elements in the hash table and t is the table size..

Chaining

In *chaining* (or *separate chaining*) records whose keys collide are formed into a linked list or chain whose head is in the hash table array.. Figure 2 below shows a hash table with collisions resolved using chaining.. For simplicity, only the keys are listed, and not the elements that go along with them (or, if you like, the key and the value are the same)..

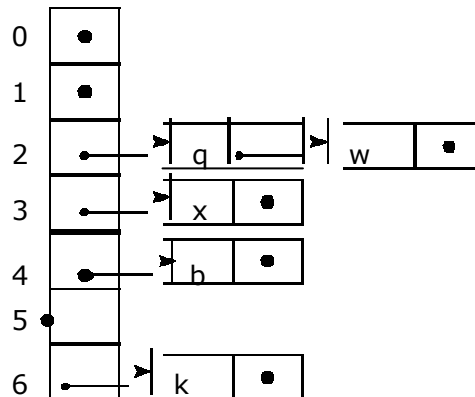


Figure 2: Hash Table with Chaining to Resolve Collisions

In this example, the table has seven locations.. Two keys, q and w , collide at location two, and they are placed in a linked list whose head is at location two.. The keys x , b , and k are hashed to locations three, four, and six, respectively.. This example uses an array of list heads, but an array of list nodes could have been used as well, with some special value in the data field to indicate when a node is unused..

When using chaining, the average chain length is λ . If the chain is unordered, successful searches required about $1 + \lambda/2$ comparisons on average, and unsuccessful searches λ comparisons on average. If the chain is ordered, both successful and unsuccessful searches take about $1 + \lambda/2$ comparisons, but insertions take longer. In the worse case, which occurs when all keys map to a single table location and the search key is at the end of the linked list or not in the list, searches require n comparisons. But this case is extremely unlikely. More complex linked structures (like binary search trees), don't generally improve performance much, particularly if λ is kept fairly small. As a rule of thumb, λ should be kept less than about 10. But performance only degrades gradually as the number of items in the table grows, so hash tables that use chaining to resolve collisions can perform well on wide ranges of values of n .

Open Addressing

In the open addressing collision resolution scheme, records with colliding keys are stored at other free locations in the hash table found by *probing* the table for open locations. Different kinds of open addressing schemes use different *probe sequences*. In all cases, however, the table can only hold as many items as it has locations, that is $n \leq t$, and λ cannot exceed one; this constraint is not present for chaining.

Theoretical studies of random probe sequences have determined ideal levels of performance for open addressing. Performance declines sharply as the load factor approaches one. For example, with a load factor of 0.9, the number of comparisons for a successful search is about 2.6, and for an unsuccessful search is 20. Real open addressing schemes do not do even as well as this, so load factors must generally be kept below about 0.75.

Another point to understand about open addressing is that when a collision occurs, the algorithm proceeds through the probe sequence until either (a) the desired key is found, (b) an open location is found, or (c) the entire table is traversed. But this only works when a marker is left in slots where an element was deleted to indicate that the location may not have been empty before, and so the algorithm should proceed with the probe sequence. In a highly dynamic table, there will be many markers and few empty slots, so the algorithm will need to follow long probe sequences, especially for unsuccessful searches, even when the load factor is low.

Linear probing is using a probe sequence that begins with the hashed table index and increments it by a constant value modulo the table size. If the table size and increment are relatively prime, every slot in the table will appear in the probe sequence. Linear probing performance degrades sharply when load factors exceed 0.8. Linear probing is also subject to *primary clustering*, which occurs when clumps of filled locations accumulate around a location where a collision first occurs. Primary clustering increases the chances of collisions and greatly degrades performance.

Double hashing works by generating an increment to use in the probe sequence using a second hash function on the key. The second hash function should generate values quite different from the first so that two keys that collide will be mapped to different values by the second hash function, making the probe

Hashing

sequences for the two keys different.. Double hashing eliminates primary clustering.. The second hash function must always generate a number that is relatively prime to the table size.. This is easy if the table size is a prime number.. Double hashing works so well that its performances approximates that of a truly random probe sequence.. It is thus the method of choice for generating probe sequences..

Figure 3 below shows an example of open addressing with double hashing.. As before, the example only uses keys for simplicity, not key-value pairs.. The main hash function is $f(x) = x \bmod 7$, and the hash function used to generate a constant for the probe sequence is $g(x) = (x \bmod 5) + 1$.. The values 8, 12, 9, 6, 25, and 22 are hashed into the table..

| | |
|---|----|
| 0 | 22 |
| 1 | 8 |
| 2 | 9 |
| 3 | -- |
| 4 | 25 |
| 5 | 12 |
| 6 | 6 |

Figure 3: Hash Table with Open Addressing and Double Hashing to Resolve Collisions

The first five keys do not collide.. But $22 \bmod 7$ is 1, so 22 collides with 8.. The probe constant for double hashing is $(22 \bmod 5) + 1 = 3$.. We add 3 to location 1, where the collision occurs, to obtain location 4.. But 25 is already at this location, so we add 3 again to obtain location 0 (we wrap around to the start of the array using the table size: $(4+3) \bmod 7 = 0$).. Location 0 is not occupied, so that is where 22 is placed..

Note that some sort of special value must be placed in the unoccupied locations—in this example we used a double dash.. A different value must be used when a value is removed from the table to indicate that the location is free, but that it was not before, so that searches must continue past this value when it is encountered during a probe sequence..

We have noted that when using open addressing to resolve collisions, performance degrades considerably as the load factor approaches one.. Hashing mechanisms that use open addressing must have a way to expand the table, thus lowering the load factor and improving performance.. A new, larger table can be created and filled by traversing the old table and inserting all records into the new table.. Note that this involves hashing every key again because the hash function will generally use the table size, which has now changed.. Consequently, this is a very expensive operation..


Some table expansion schemes work incrementally by keeping the old table around and making all insertions in the new table, all deletions from the old table, and perhaps

moving records gradually from the old table to the new in the course of doing other operations.. Eventually the old table becomes empty and can be discarded..

Binary Search Trees and AVL Trees

Introduction

- Binary search trees are an excellent data structure to implement associative arrays, maps, sets, and similar interfaces.
- The main difficulty, as discussed in last lecture, is that they are efficient only when they are balanced.
- Straightforward sequences of insertions can lead to highly unbalanced trees with poor asymptotic complexity and unacceptable practical efficiency.
- The solution is to dynamically rebalance the search tree during insert or search operations.
- We have to be careful not to destroy the ordering invariant of the tree while we rebalance.
- Because of the importance of binary search trees, researchers have developed many different algorithms for keeping trees in balance, such as AVL trees, red/black trees, splay trees, or randomized binary search trees.
-
- In this lecture we discuss AVL trees, which is a simple and efficient data structure to maintain balance.
- It is named after its inventors, G.M. Adelson-Velskii and E.M. Landis, who described it in 1962.



Height Invariant

Ordering Invariant.

- At any node with key k in a binary search tree, all keys of the elements in the left subtree are strictly less than k , while all keys of the elements in the right subtree are strictly greater than k .


- To describe AVL trees we need the concept of tree height, which we define as the maximal length of a path from the root to a leaf. So the empty tree has height 0, the tree with one node has height 1, a balanced tree with three nodes has height 2.
- If we add one more node to this last tree it will have height 3. Alternatively, we can define it recursively by saying that the empty tree has height 0, and the height of any node is one greater than the maximal height of its two children.
-
- AVL trees maintain a height invariant (also sometimes called a balance invariant).
-
- At any node in the tree, the heights of the left and right subtrees differs by at most 1.

Rotations: How they work

Left Rotation (LL)

Imagine we have this situation:

Figure 1-1




To fix this, we must perform a left rotation, rooted at A. This is done in the following steps:
 b becomes the new root.

a takes ownership of b's left child as its right child, or in this case, null.

b takes ownership of a as its left child.

The tree now looks like this:


Figure 1-2



Right Rotation (RR)

A right rotation is a mirror of the left rotation operation described above. Imagine we have this situation:

Figure 1-3



To fix this, we will perform a single right rotation, rooted at C. This is done in the following steps:


b becomes the new root.

c takes ownership of b's right child, as its left child. In this case, that value is null.

b takes ownership of c, as it's right child.


The resulting tree:

Figure 1-4




Double Rotations (left-right)

Figure 1-6



Our initial reaction here is to do a single left rotation. Let's try that


Figure 1-7




Our left rotation has completed, and we're stuck in the same situation. If we were to do a single right rotation in this situation, we would be right back where we started. What's causing this? The answer is that this is a result of the right subtree having a negative balance. In other words, because the right subtree was left heavy, our rotation was not sufficient. What can we do? The answer is to perform a right rotation on the right subtree. Read that again. We will perform a right rotation on the *right subtree*. We are not rotating on our current root. We are rotating on our right child. Think of our right subtree, isolated from our main tree, and perform a right rotation on it:

After performing a rotation on our right subtree, we have prepared our root to be rotated left.

Here is our tree now:




Looks like we're ready for a left rotation.




Double Rotations (right-left)


A double right rotation, or right-left rotation, or simply RL, is a rotation that must be performed when attempting to balance a tree which has a left subtree, that is right heavy. This is a mirror operation of what was illustrated in the section on Left-Right Rotations, or double left rotations. Let's look at an example of a situation where we need to perform a Right-Left rotation.



In this situation, we have a tree that is unbalanced. The left subtree has a height of 2, and the right subtree has a height of 0. This makes the balance factor of our root node, c, equal to -2. What do we do? Some kind of right rotation is clearly necessary, but a single right rotation will not solve our problem.




Looks like that didn't work. Now we have a tree that has a balance of 2. It would appear that we did not accomplish much. That is true. What do we do? Well, let's go back to the original tree, before we did our pointless right rotation:




The reason our right rotation did not work, is because the left subtree, or 'a', has a positive balance factor, and is thus right heavy. Performing a right rotation on a tree that has a left subtree that is right heavy will result in the problem we just witnessed. What do we do?

The answer is to make our left subtree left-heavy. We do this by performing a left rotation on our left subtree. Doing so leaves us with this situation:



This is a tree which can now be balanced using a single right rotation. We can now perform our right rotation rooted at C.



Increasing the speed by minimizing the height difference is the main of AVL tree.

Operations like insertions, deletions etc can be done in time of $O(\log n)$, even in the worst case.

e.g. In a complete balanced tree, the left and right subtrees of any node would have same height.

Height difference is 0.

Suppose H_l is the height of left sub tree and H_r is the height of right sub tree, then following properties must be satisfied for the tree to be an AVL tree:


1. It must be a BST.

2. Height of left sub-tree - Height of right sub-tree ≤ 1 .

i.e. $|H_l - H_r| \leq 1$

$$H_l - H_r = \{-1, 0, 1\}$$

e.g. (2)



Balancing factor = $H_l - H_r$.

In the above diagram, balancing factor for root node is 2, so, it is not an AVL tree. In such cases, the tree can be balanced by rotations.

There are 4 type of rotations:


1. LL Rotation (left-left)

2. RR Rotation (right-right)


3. LR Rotation (left-right)

4. RL Rotation (right-left)

LL Rotation:




Balance factor of A is +2. Applying LL rotation, we get,




Now, this is a height balanced tree.

RR Rotation:




Balance factor of A is -2. Applying RR rotation, we get,



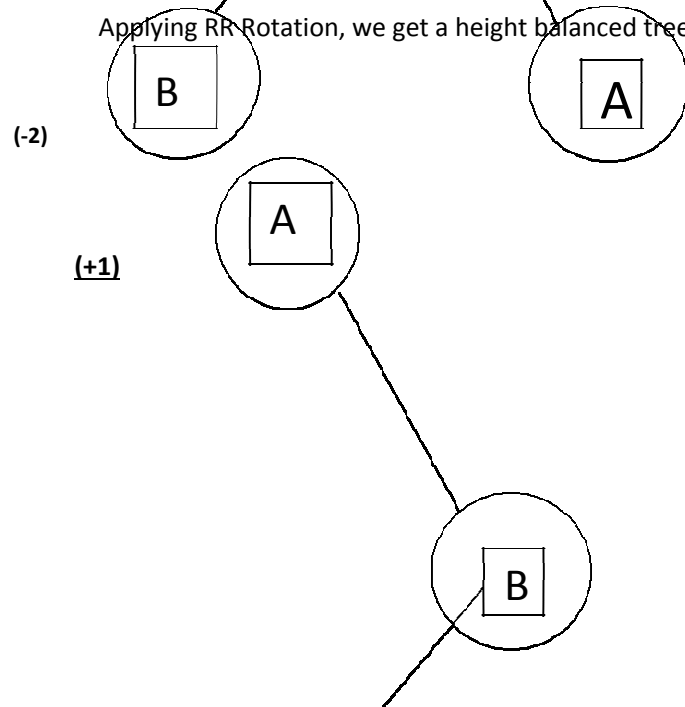
Now, this is a height balanced tree.

LR Rotation:



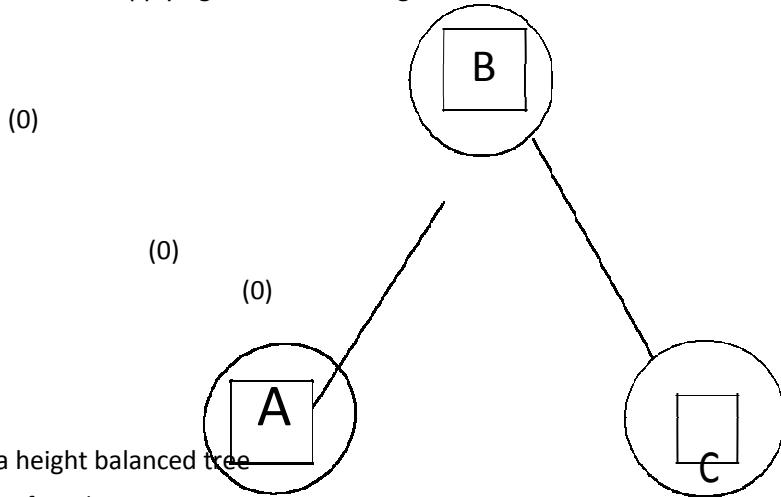
Applying RR Rotation, we get a height balanced tree.

RL Rotation:



(0)

Balance factor of A is -2; applying RL Rotation, we get,

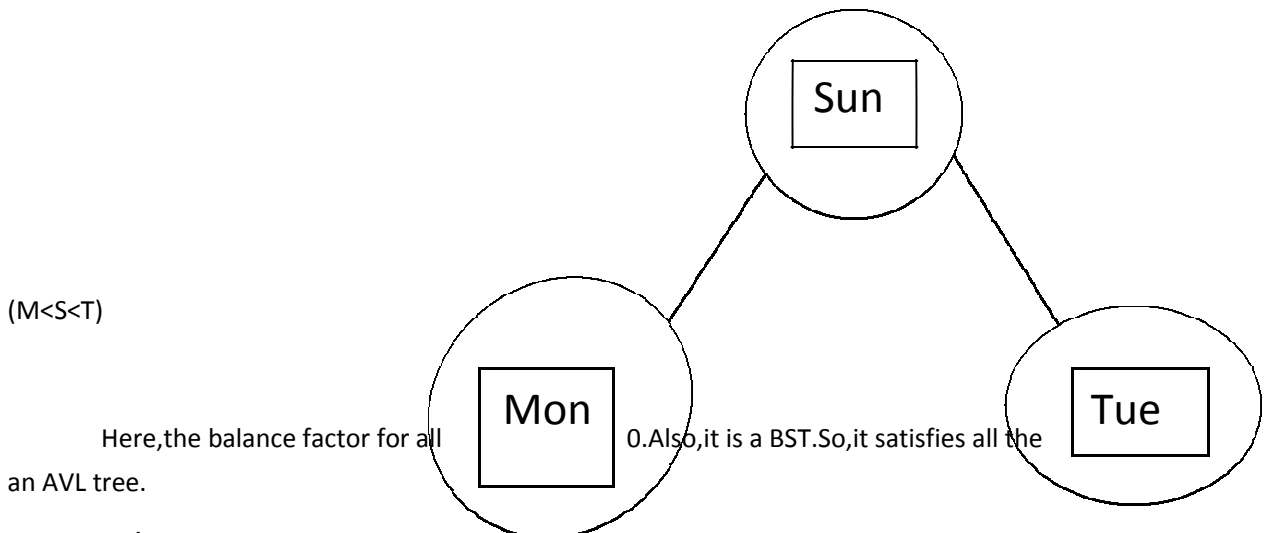


Now, this is a height balanced tree


Q. Arrange days of week:-


Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, Saturday, in an AVL tree.

1st step:- We have to arrange these days alphabetically, and the constructed tree should satisfy the conditions of an AVL Tree. Starting with Sunday (Sun):



2nd step:- Now, Wednesday is to be inserted, as $(W) > (T)$, so, it will be placed at right of Tuesday for satisfying BST conditions:





Now, balance factor of Sunday=-1

Balance factor of Monday=-1


Balance factor of Tuesday=0

Balance factor of Wednesday=-1

Hence, it is an AVL tree.

3rd step:-

Now, Thursday has to be inserted. As alphabetically, Thursday < Tuesday, so it will be at the left of Tuesday.





Here, balance factor of sun=-1
Balance factor of tue, wed, thu=0

So, it is an AVL tree.

4th step:

Now, Friday and Saturday are to be inserted,





Here, balance factor of all the days=0.

So, it is an AVL tree.

Use :

1. Search is $O(\log N)$ since AVL trees are always balanced.
2. Insertion and deletions are also $O(\log n)$
3. The height balancing adds no more than a constant factor to the speed of insertion.
4. AVL Tree height is less than BST.
5. Stores data in balanced way in a disk.

Limitations


1. Difficult to program & debug; more space for balance factor.
2. Asymptotically faster but rebalancing costs time.
3. Most large searches are done in database systems on disk and use other structures (e.g. B-trees).
4. May be OK to have $O(N)$ for a single operation if total run time for many consecutive operations is fast (e.g. Splay trees).

Threaded BST

A **threaded binary tree** defined as follows:

"A binary tree is *threaded* by making all right child pointers that would normally be null point to the inorder successor of the node (if it exists) , and all left child pointers that would normally be null point to the inorder predecessor of the node."

A threaded binary tree makes it possible to traverse the values in the binary tree via a linear traversal that is more rapid than a recursive in-order traversal. It is also possible to discover the parent of a node from a threaded binary tree, without explicit use of parent pointers or a stack, albeit slowly. This can be useful where stack space is limited, or where a stack of parent pointers is unavailable (for finding the parent pointer via DFS).



A threaded tree, with the special threading links shown by dashed arrows.

7.Types of threaded binary trees

1. Single Threaded: each node is threaded towards **either(right)' the in-order predecessor or'** successor.
2. Double threaded: each node is threaded towards **both(left & right)' the in-order predecessor and'** successor.

SPANNING TREE:-

->A tree is a connected undirected graph with no cycles.It is a spanning tree of a graph G if it spans G (that is, it

includes every vertex of G) and is a subgraph of G (every edge in the tree belongs to G).

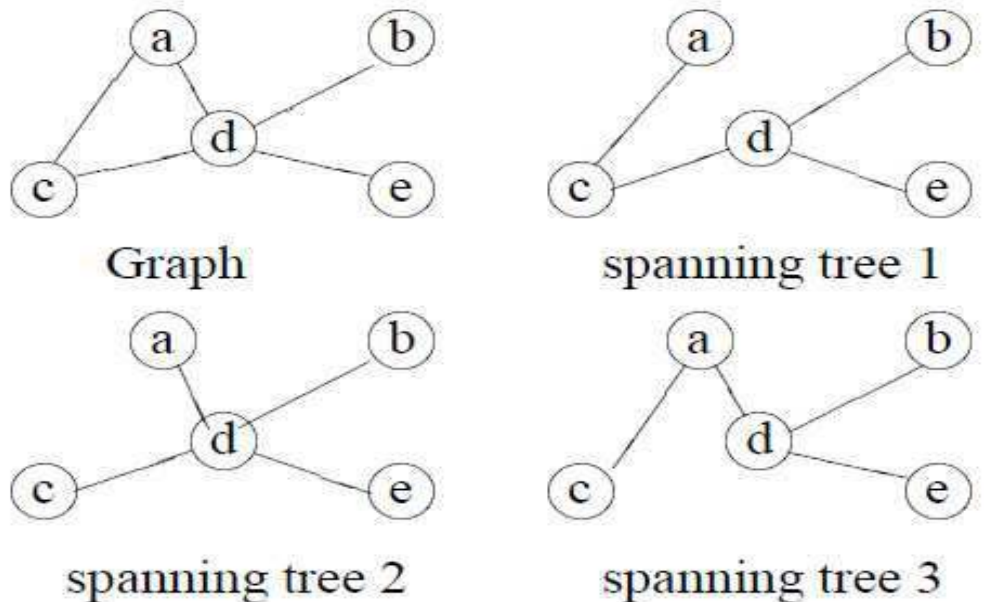
->A spanning tree of a connected graph G can also be defined as a maximal set of edges of G that contains no cycle,

or as a minimal set of edges that connect all vertices.

-> So key to spanning tree is number of edges will be 1 less than the number of nodes.

no. of edges=no. of nodes-1

example:-



Weighted Graphs:- A weighted graph is a graph, in which each edge has a weight (some real number).

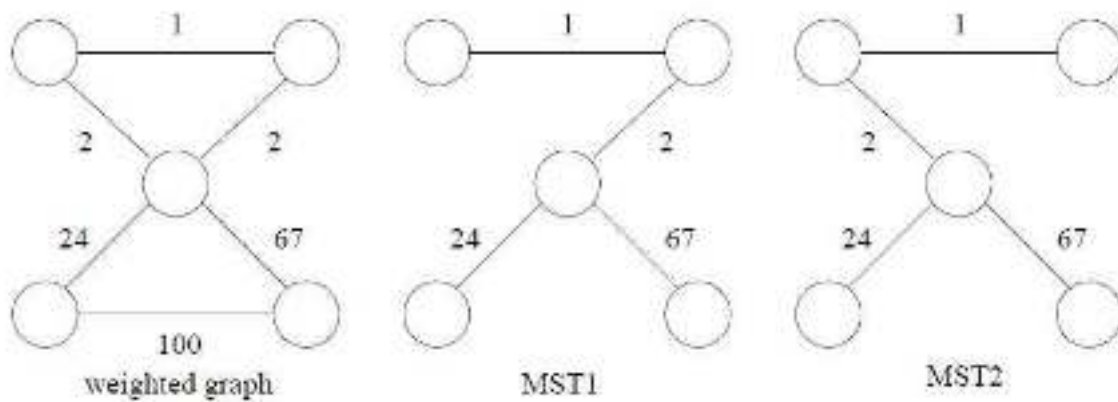
Weight of a Graph:- The sum of the weights of all edges.

Minimum SPANNING TREE:-

Minimum spanning tree is the spanning tree in which the sum of weights on edges is minimum.

NOTE:- The minimum spanning tree may not be unique. However, if the weights of all the edges are pairwise distinct, it is indeed unique.

example:-





There are a number of ways to find the minimum spanning tree but out of them most popular methods are prim's algorithm and kruskal algorithm.

PRIM'S ALGORITHM:-

Prim's algorithm is a greedy algorithm that finds a minimum spanning tree for a connected weighted graph. This means it finds a subset of the edges that forms a tree that includes every vertex, where the total weight of all the edges in the tree is minimized.



steps:-

1. Initialize a tree with a single vertex, chosen arbitrarily from the graph.
2. Grow the tree by one edge: of the edges that connect the tree to vertices not yet in the tree, find the minimum-weight edge, and transfer it to the tree.
3. Repeat step 2 (until all vertices are in the tree).





step-0

choose a.





step-1

since (a,b) is minimum.





step-2

since (b,d) is minimum.





step-3

since weight of (d,c) is minimum.





step-4

since (c,f) is minimum.



step-5

since weight of (f,g) is less than (f,e).



step-6

weight of (g,e) is greater than weight of (f,e).



kruskal algorithm:-

The Kruskal's Algorithm is based directly on the generic algorithm. Unlike Prim's algorithm, we make a different choice of cut.

- Initially, trees of the forest are the vertices (no edges).
- In each step add the cheapest edge that does not create a cycle.
- Observe that unlike Prim's algorithm, which only grows one tree, Kruskal's algorithm grows a collection of trees (a forest).
- Continue until the forest 'merge to' a single tree.

This is a *minimum* spanning tree.

example:-



| | edge | weight |
|---|--------|--------|
| E | {d, c} | 2 |
| | {a, e} | 3 |
| | {a, d} | 5 |
| | {e, d} | 7 |
| | {b, c} | 9 |
| | {a, b} | 10 |
| | {b, d} | 12 |
| | | |

Forest (V, A)

A={ }

