

Lecture 13

Sets

15-122: Principles of Imperative Computation (Spring 2017)
Frank Pfenning, Rob Simmons

1 Introduction

In this lecture, we will discuss the data structure of hash tables further and use hash tables to implement a very basic interface of sets. With this lecture, we will also begin to discuss a new separation of concerns. Previously, we have talked a great deal about the distinction between a library's interface (which the client can rely on) and a library's implementation (which should be able to change without affecting a correctly-designed client).

The interface defines not only the *types*, but also the available *operations* on them and the pre- and post-conditions for these operations. For general data structures it is also useful to note the asymptotic complexity of the operations so that potential clients can decide if the interface serves their purpose.

One wrinkle we have not yet discussed is that, in order for a library to provide its services, it may in turn require some operations supplied by the client. Hash tables provide an excellent example for this complication, so we will discuss the interface to hash tables in details before giving the hash table implementation.

For the purposes of this lecture we call the data structures and the operations on them provided by an implementation the *library* and code that uses the library the *client*.

Relating to our learning goals, we have

Computational Thinking: We discuss the separation of client interfaces and client implementations.

Algorithms and Data Structures: We discuss algorithms for hashing strings.

Programming: We revisit the **char** data type and use it to consider string hashing.

2 Generic Data Structures

So far, all the data structures that we've considered, have always had particular type information that seemed irrelevant. In the implementation of queues, why is it important that we have a queue of *strings* in particular?

```
// typedef _____* queue_t;
bool queue_empty(queue_t Q)                /* 0(1) */
    /*@requires Q != NULL; @*/;
queue_t queue_new()                        /* 0(1) */
    /*@ensures \result != NULL; @*/;
void enq(queue_t Q, string x)              /* 0(1) */
    /*@requires Q != NULL; @*/;
string deq(queue_t S)                      /* 0(1) */
    /*@requires Q != NULL && !queue_empty(S); @*/ ;
```

It's both wasteful and a potential source of errors to have to rewrite our code if we want our program to use integers (or chars, or pointers to structs, or arrays of strings, ...) instead of strings. A way we deal with this is by creating a type, *elem*, that is used by the library but not defined in the library:

```
/** Client interface */
// typedef _____ elem;    // Supplied by client

/** Library interface */
// typedef _____* queue_t;
bool queue_empty(queue_t Q)                /* 0(1) */
    /*@requires Q != NULL; @*/;
queue_t queue_new()                        /* 0(1) */
    /*@ensures \result != NULL; @*/;
void enq(queue_t Q, elem x)                /* 0(1) */
    /*@requires Q != NULL; @*/;
elem deq(queue_t Q)                       /* 0(1) */
    /*@requires Q != NULL && !queue_empty(S); @*/ ;
```

The underscores in the library interface, before *queue_t*, mean that the client doesn't know how the abstract type *queue_t* is implemented beyond knowing that it is a pointer. The library is therefore free to change this implementation without breaking any (interface-respecting) client code. The underscores in the *client* interface mean that the *library* doesn't know how the abstract type *elem* is implemented, which means that the client is free

to change this implementation without breaking the library. The library's implementation just refers to the `elem` type, which it expects the client to have already defined, whenever it needs to refer to client data.

This approach is still not perfect, because any given program only supports a single type of queue element. We'll start working on that problem in the next lecture.

3 Generic Hash Sets

When we implement the set interface with a hash table, we'll call it a *hash set* or `hset`. When we implement the dictionary interface with a hash table, we'll call it a *hash dictionary* or `hdict`. Our hash set implementation will be *generic*; it will work regardless of the type of keys or elements to be stored in the table.

We need to think carefully about which types and functions are provided by the client of the hash set, and which are provided by the library itself. Clearly, the library should determine the type of hash sets:

```
/* library side types */
// typedef _____* hset_t;
```

That is really the only type provided by the implementation. In addition, the library interface is supposed to provide a few functions:

```
/* library side functions */
hset_t hset_new(int capacity)          /* 0(1) */
/*@requires capacity > 0; @*/
/*@ensures \result != NULL; @*/ ;

bool hset_contains(hset_t H, elem x)  /* 0(1) avg. */
/*@requires H != NULL; @*/ ;

void hset_add(hset_t H, elem x)        /* 0(1) avg. */
/*@requires H != NULL; @*/
/*@ensures hset_contains(H, x); @*/ ;
```

The function `hset_new(int capacity)` takes the initial capacity of the hash table as an argument (which must be strictly positive) and returns a new hash set without any elements.

The function `hset_contains(hset_t H, elem x)` answers the question of whether the element `x` has been added to the set already. The last function, `hset_add(hset_t H, elem x)`, ensures that `x` is now in the set.

From these decisions we can see that the *client* must provide the type of elements. Only the client can know what these might be in any particular use of the library. In this implementation, we don't need to know anything about the type `elem`, which we indicate by using only underscores:

```
/* client-side types */  
// typedef _____ elem;      // Supplied by client
```

Does the client also need to provide any functions? Yes! The hash table implementation needs functions that can operate on values of the types `elem` so that it can hash elements and so that it can determine whether they are equal. Since the library is supposed to be generic, the library implementer cannot write these functions; we require the client to provide them.

There are two of these “client-side” functions. First, and most obviously, we need a hash function which maps keys to integers.

```
/* client-side functions */  
int elem_hash(elem x);
```

The result, the *hash value*, can be any integer, so our hash table implementation will have to take both this arbitrary integer and m , the size of the hash table's table, into consideration when figuring out which index of the table the element hashes to. For the hash table implementation to achieve its advertised (average-case) asymptotic complexity, the hash function should have the property that its results are evenly distributed between 0 and m . The hash set implementation will work correctly (albeit slowly) even if it maps every key to 0.

Hash table operations also need to be able to check for the equality of elements in order to be able to tell whether two objects that collide are actually the same or not.

```
/* client-side functions */  
bool elem_equiv(elem x, elem y);
```

This completes the interface which we now summarize.

```

/*****
*** Client interface ***
*****/
// typedef _____* elem;    // Supplied by client
bool elem_equiv(elem x, elem y);
int elem_hash(elem x);

/*****
*** Library interface ***
*****/
// typedef _____* hset_t;

hset_t hset_new(int capacity)
/*@requires capacity > 0; @*/
/*@ensures \result != NULL; @*/ ;

elem hset_contains(hset_t H, elem x)
/*@requires H != NULL; @*/ ;

void hset_add(hset_t H, elem x)
/*@requires H != NULL; @*/
/*@ensures hset_contains(H, x); @*/ ;
```

4 A Tiny Client

One sample application is to count word occurrences — say, in a corpus of Twitter data or in the collected works of Shakespeare. In this application, the keys are the words, represented as strings. Data elements are pairs of words and word counts, the latter represented as integers.

```
1 /*****  
2  /* client-side implementation */  
3  *****/  
4  struct wcount {  
5      string word;  
6      int count;  
7  };  
8  
9  typedef struct wcount* elem;  
10  
11 int elem_hash(elem x)  
12 //@requires x != NULL;  
13 {  
14     return hash_string(x->word);    /* from hash-string.c0 */  
15 }  
16  
17 bool elem_equiv(elem x1, elem x2)  
18 //@requires x1 != NULL && x2 != NULL;  
19 {  
20     return string_equiv(x1->word, x2->word);  
21 }
```

5 A Universal Hash Function

One question we have to answer is how to hash strings, that is, how to map strings to integers so that the integers are evenly distributed no matter how the input strings are distributed.

We can get access to the individual characters in a string with the function `string_charat(s, i)`, and we can get the integer ASCII value of a **char** with the function `char_ord(c)`; both of these are defined in the C0 **string** library. Therefore, our general picture of hashing strings looks like this:

```
23 int hash_string(string s) {  
24     int len = string_length(s);  
25     int h = 0;  
26     for (int i = 0; i < len; i++)  
27         //@loop_invariant 0 <= i;  
28     {  
29         int ch = char_ord(string_charat(s, i));  
30         // Do something to combine h and ch  
31     }  
32     return h;  
33 }
```

Now, if we don't add anything to replace the comment, the function above will still allow the hash table to work correctly, it will just be very slow because the hash value of every string will be zero.

A slightly better idea is combining h and ch with addition or multiplication:

```
26     for (int i = 0; i < len; i++)  
27         //@loop_invariant 0 <= i;  
28     {  
29         int ch = char_ord(string_charat(s, i));  
30         h = h + ch;  
31     }
```

This is still pretty bad, however. We can see how bad by entering the $n = 45,600$ news vocabulary words from the Collected Works of William Shakespeare, say, into a table with $m = 22,800$ chains (load factor is 2) and running `ht_stats`:

Hash table distribution: how many chains have size...

```
...0: 21217  
...1: 239  
...2: 132  
...3: 78  
...4: 73  
...5: 55  
...6: 60  
...7: 46  
...8: 42  
...9: 23  
...10+: 835
```

Longest chain: 176

Most of the chains are empty, and many of the chains are very, very long. One problem is that most strings are likely to have very small hash values when we use this hash function. An even bigger problem is that rearranging the letters in a string will always produce another string with the same hash value — so we know that "cab" and "abc" will always collide in a hash table. Hash collisions are inevitable, but when we can easily predict that two strings have the same hash value, we should be suspicious that something is wrong.

To address this, we can manipulate the value h in some way before we combine it with the current value. Some versions of Java use this as their default string hashing function.

```
26  for (int i = 0; i < len; i++)
27  // @loop_invariant 0 <= i;
28  {
29      int ch = char_ord(string_charat(s, i));
30      h = 31*h;
31      h = h + ch;
32  }
```

This does much better when we add all the news vocabulary strings into the hash table:

Hash table distribution: how many chains have size...

```
...0: 3057
...1: 6210
...2: 6139
...3: 4084
...4: 2151
...5: 809
...6: 271
...7: 53
...8: 21
...9: 4
...10+: 1
```

Longest chain: 10

We can try adding a bit of randomness into this function in a number of different ways. For instance, instead of multiplying by 31, we could multiply by a number generated by the pseudo-random number generator from C0's library:


```
26  rand_t r = init_rand(0x1337BEEF);
27  for (int i = 0; i < len; i++)
28      //@loop_invariant 0 <= i;
29      {
30          int ch = char_ord(string_charat(s, i));
31          h = rand(r) * h;
32          h = h + ch;
33      }
```

If we look at the performance of this function, it is comparable to the Java hash function, though it is not actually quite as good — more of the chains are empty, and more are longer.

Hash table distribution: how many chains have size...

```
...0: 3796
...1: 6214
...2: 5424
...3: 3589
...4: 2101
...5: 1006
...6: 455
...7: 145
...8: 48
...9: 15
...10+: 7
```

Longest chain: 11

Many other variants are possible; for instance, we could try directly applying the linear congruential generator to the hash value at every step:

```
26  for (int i = 0; i < len; i++)
27      //@loop_invariant 0 <= i;
28      {
29          int ch = char_ord(string_charat(s, i));
30          h = 1664525 * h + 1013904223;
31          h = h + ch;
32      }
```

The key goals are that we want a hash function that is very quick to compute and that nevertheless achieves good distribution across our hash table. Handwritten hash functions often do not work well, which can significantly affect the performance of the hash table. Whenever possible, the use of randomness can help to avoid any systematic bias.

6 A Fixed-Size Implementation of Hash Tables

The implementation of hash tables we wrote in lecture did not adjust their size. This requires that we can a priori predict a good size, or we will not be able to get the advertised $O(1)$ average time complexity. Choose the size too large and it wastes space and slows the program down due to a lack of locality. Choose the size too small and the load factor will be high, leading to poor asymptotic (and practical) running time.

We start with the type of lists to represent the chains of elements, and the hash table type itself.

```

1  /*****
2  /* library-side implementation */
3  *****/
4  typedef struct chain_node chain;
5  struct chain_node {
6      elem data;
7      chain* next;
8  };
9
10 struct hset_header {
11     int size;          /* size >= 0          */
12     int capacity;      /* capacity > 0        */
13     chain*[] table;    /* \length(table) == capacity */
14 };

```

The first thing after the definition of a data structure is a function to verify its invariants. Besides the invariants noted above we should check the hash value of the key of every element in each chain stored in $A[i]$ is indeed i . (This `is_hset` function is incomplete.)

```

16 bool is_hset(hset H) {
17     return  H != NULL
18           && H->capacity > 0
19           && H->size >= 0
20           && is_table_expected_length(H->table, H->capacity);
21     /* && each element satisfies its own representation invariants */
22     /* && there aren't equal elements */
23     /* && the number of elements matches the size */
24     /* && every element in H->table[i] hashes to i */
25 }

```

Recall that the test on the length of the array must be inside an annotation, because the `\length` function is not available when the code is compiled without dynamic checking enabled.

In order to check that the elements of a hash set hash to the correct index, we need a way of mapping the hash value returned by `elem_hash` to an index of the table. This is a common enough operation that we'll write a helper function:

```

27 int elemhash(hset H, elem x)
28 //@requires H != NULL && H->capacity > 0;
29 //@ensures 0 <= \result && \result < H->capacity;
30 {
31     return abs(elem_hash(x) % H->capacity);
32 }
```

Allocating a hash table is straightforward.

```

34 hset hset_new(int capacity)
35 //@requires capacity > 0;
36 //@ensures is_hset(\result);
37 {
38     hset H = alloc(struct hset_header);
39     H->size = 0;
40     H->capacity = capacity;
41     H->table = alloc_array(chain*, capacity);
42     return H;
43 }
```

Equally straightforward is searching for an element with a given key. We omit an additional loop invariant.

```

45 bool hset_contains(hset H, elem x)
46 //@requires is_hset(H);
47 {
48     int i = elemhash(H, x);
49     for (chain* p = H->table[i]; p != NULL; p = p->next) {
50         if (elem_equiv(p->data, x)) return true;
51     }
52     return false;
53 }
```

Inserting an element follows generally the same structure as search. If we find an element in the correct chain with the same key we replace it. If we find none, we insert a new node at the beginning of the chain.

```
55 void hset_insert(hset H, elem x)
56 // @requires is_hset(H);
57 // @ensures is_hset(H);
58 // @ensures x == hset_contains(H, x);
59 {
60     int i = elemhash(H, x);
61     for (chain* p = H->table[i]; p != NULL; p = p->next)
62     {
63         if (elem_equiv(p->data, x)) {
64             p->data = x;
65             return;
66         }
67     }
68
69     // prepend new element
70     chain* p = alloc(chain);
71     p->data = x;
72     p->next = H->table[i];
73     H->table[i] = p;
74     (H->size)++;
75 }
```

Exercises

Exercise 1. *Extend the hash table implementation so it dynamically resizes itself when the load factor exceeds a certain threshold. When doubling the size of the hash table you will need to explicitly insert every element from the old hash table into the new one, because the result of hashing depends on the size of the hash table.*

Exercise 2. *Redo the library implementation for a different client interface that has a function `elem_hash(key k, int m)` that returns a result between 0 (inclusive) and `m` (exclusive).*

Exercise 3. *Extend the hash table interface with a new function `ht_tabulate` that returns an array with the elements in the hash table, in some arbitrary order.*

Exercise 4. *Extend the hash table interface with a new function to delete an element with a given key from the table. To be extra ambitious, shrink the size of the hash table once the load factor drops below some minimum, similarly to the way we could grow and shrink unbounded arrays.*