

ECE 2400 Computer Systems Programming

Fall 2017

Topic 18: Tables

School of Electrical and Computer Engineering
Cornell University

revision: 2017-11-29-12-14

1	Tables Concepts	2
2	Lookup Tables	3
2.1.	Set of Integers as a Lookup Table	3
2.2.	Set of Strings as a Lookup Table	5
2.3.	Strengths and Weaknesses of Lookup Tables	7
3	Hash Tables	8

1. Tables Concepts

- Tables have rows and columns
- Rows are usually identified by a name (number or even a string)
- Columns are usually identified by a name (number or even a string)
- Rows may have multiple columns of data
- Tables are indexed to access row data at a certain column

	A	B	C
0			
1			
2			

Table as a Set ADT

Table as a Map ADT

	A	B	C
0	✓	✓	✓
1	✓	✓	✓
2	✓	✓	✓
3	✓	✓	✓
4	✓	✓	✓
5	✓	✓	✓
6	✓	✓	✓
7	✓	✓	✓
8	✓	✓	✓
9	✓	✓	✓
10	✓	✓	✓
11	✓	✓	✓
12	✓	✓	✓
13	✓	✓	✓
14	✓	✓	✓
15	✓	✓	✓
16	✓	✓	✓
17	✓	✓	✓
18	✓	✓	✓
19	✓	✓	✓
20	✓	✓	✓
21	✓	✓	✓
22	✓	✓	✓
23	✓	✓	✓
24	✓	✓	✓
25	✓	✓	✓
26	✓	✓	✓
27	✓	✓	✓
28	✓	✓	✓
29	✓	✓	✓
30	✓	✓	✓
31	✓	✓	✓
32	✓	✓	✓
33	✓	✓	✓
34	✓	✓	✓
35	✓	✓	✓
36	✓	✓	✓
37	✓	✓	✓
38	✓	✓	✓
39	✓	✓	✓
40	✓	✓	✓
41	✓	✓	✓
42	✓	✓	✓
43	✓	✓	✓
44	✓	✓	✓
45	✓	✓	✓
46	✓	✓	✓
47	✓	✓	✓
48	✓	✓	✓
49	✓	✓	✓
50	✓	✓	✓
51	✓	✓	✓
52	✓	✓	✓
53	✓	✓	✓
54	✓	✓	✓
55	✓	✓	✓
56	✓	✓	✓
57	✓	✓	✓
58	✓	✓	✓
59	✓	✓	✓
60	✓	✓	✓
61	✓	✓	✓
62	✓	✓	✓
63	✓	✓	✓
64	✓	✓	✓
65	✓	✓	✓
66	✓	✓	✓
67	✓	✓	✓
68	✓	✓	✓
69	✓	✓	✓
70	✓	✓	✓
71	✓	✓	✓
72	✓	✓	✓
73	✓	✓	✓
74	✓	✓	✓
75	✓	✓	✓
76	✓	✓	✓
77	✓	✓	✓
78	✓	✓	✓
79	✓	✓	✓
80	✓	✓	✓
81	✓	✓	✓
82	✓	✓	✓
83	✓	✓	✓
84	✓	✓	✓
85	✓	✓	✓
86	✓	✓	✓
87	✓	✓	✓
88	✓	✓	✓
89	✓	✓	✓
90	✓	✓	✓
91	✓	✓	✓
92	✓	✓	✓
93	✓	✓	✓
94	✓	✓	✓
95	✓	✓	✓
96	✓	✓	✓
97	✓	✓	✓
98	✓	✓	✓
99	✓	✓	✓
100	✓	✓	✓

	A	B	C
0	✓	✓	✓
1	✓	✓	✓
2	✓	✓	✓
3	✓	✓	✓
4	✓	✓	✓
5	✓	✓	✓
6	✓	✓	✓
7	✓	✓	✓
8	✓	✓	✓
9	✓	✓	✓
10	✓	✓	✓
11	✓	✓	✓
12	✓	✓	✓
13	✓	✓	✓
14	✓	✓	✓
15	✓	✓	✓
16	✓	✓	✓
17	✓	✓	✓
18	✓	✓	✓
19	✓	✓	✓
20	✓	✓	✓
21	✓	✓	✓
22	✓	✓	✓
23	✓	✓	✓
24	✓	✓	✓
25	✓	✓	✓
26	✓	✓	✓
27	✓	✓	✓
28	✓	✓	✓
29	✓	✓	✓
30	✓	✓	✓
31	✓	✓	✓
32	✓	✓	✓
33	✓	✓	✓
34	✓	✓	✓
35	✓	✓	✓
36	✓	✓	✓
37	✓	✓	✓
38	✓	✓	✓
39	✓	✓	✓
40	✓	✓	✓
41	✓	✓	✓
42	✓	✓	✓
43	✓	✓	✓
44	✓	✓	✓
45	✓	✓	✓
46	✓	✓	✓
47	✓	✓	✓
48	✓	✓	✓
49	✓	✓	✓
50	✓	✓	✓
51	✓	✓	✓
52	✓	✓	✓
53	✓	✓	✓
54	✓	✓	✓
55	✓	✓	✓
56	✓	✓	✓
57	✓	✓	✓
58	✓	✓	✓
59	✓	✓	✓
60	✓	✓	✓
61	✓	✓	✓
62	✓	✓	✓
63	✓	✓	✓
64	✓	✓	✓
65	✓	✓	✓
66	✓	✓	✓
67	✓	✓	✓
68	✓	✓	✓
69	✓	✓	✓
70	✓	✓	✓
71	✓	✓	✓
72	✓	✓	✓
73	✓	✓	✓
74	✓	✓	✓
75	✓	✓	✓
76	✓	✓	✓
77	✓	✓	✓
78	✓	✓	✓
79	✓	✓	✓
80	✓	✓	✓
81	✓	✓	✓
82	✓	✓	✓
83	✓	✓	✓
84	✓	✓	✓
85	✓	✓	✓
86	✓	✓	✓
87	✓	✓	✓
88	✓	✓	✓
89	✓	✓	✓
90	✓	✓	✓
91	✓	✓	✓
92	✓	✓	✓
93	✓	✓	✓
94	✓	✓	✓
95	✓	✓	✓
96	✓	✓	✓
97	✓	✓	✓
98	✓	✓	✓
99	✓	✓	✓
100	✓	✓	✓

2. Lookup Tables

- A lookup table is a table in which all potential data is known
- A lookup table is accessed by simple array indexing in O(1) time
- Lookup tables may be multidimensional

2.1. Set of Integers as a Lookup Table

- A set of unsigned integers in the range [0, 4]
- Implemented internally as a simple array of booleans

Index	0	1	2	3	4
Present?					

Interface

```
1 class LookupTableInt
2 {
3     public:
4         LookupTableInt();
5
6         void add      ( unsigned int v );
7         void remove   ( unsigned int v );
8         bool contains ( unsigned int v );
9
10    private:
11        static const unsigned int size = 5;
12        bool m_table[size];
13 }
```

Implementation

```
1  LookupTableInt::LookupTableInt()
2  {
3
4      void LookupTableInt::add( unsigned int v )
5      {
6          m_table[v] = true;
7      }
8
9      void LookupTableInt::remove( unsigned int v )
10     {
11         m_table[v] = false;
12     }
13
14     bool LookupTableInt::contains( unsigned int v )
15     {
16         return m_table[v];
17     }
```

Main

```
1  int main( void )
2  {
3      LookupTableInt table;
4
5      table.add(3);      // Add value 3
6      table.remove(3); // Remove value 3
7
8      return 0;
9  }
```

2.2. Set of Strings as a Lookup Table

- A set of strings limited to "Apple", "Banana", "Cherry"
- Implemented internally as a simple array of booleans
- Must correspond each possible string to an internal array index

Index	Apple	Banana	Cherry
Present?			

Interface

```
1 class LookupTableString
2 {
3     public:
4         LookupTableString();
5
6         void add      ( std::string v );
7         void remove   ( std::string v );
8         bool contains ( std::string v );
9
10    private:
11        unsigned int ind( std::string v );
12        static const unsigned int size = 3;
13        bool m_table[size];
14 }
```

Implementation

```
1 LookupTableString::LookupTableString()
2 { }
```

```
4 void LookupTableString::add( std::string v )
5 {
6     m_table[ ind(v) ] = true;
7 }
8
9 void LookupTableString::remove( std::string v )
10 {
11     m_table[ ind(v) ] = false;
12 }
13
14 bool LookupTableString::contains( std::string v )
15 {
16     return m_table[ ind(v) ];
17 }
18
19 unsigned int LookupTableString::ind( std::string v )
20 {
21     if      ( v == "Apple" ) return 0;
22     else if ( v == "Banana" ) return 1;
23     else if ( v == "Cherry" ) return 2;
24     else throw std::out_of_range( "No such fruit!" );
25 }
```

Main

```
1 int main( void )
2 {
3     LookupTableString table;
4
5     table.add( "Cherry" );    // Add value "Cherry"
6     table.remove( "Cherry" ); // Remove value "Cherry"
7
8     return 0;
9 }
```

2.3. Strengths and Weaknesses of Lookup Tables

- We have implemented a set of integers as a lookup table
- We have implemented a set of strings as a lookup table
- Recall that lists, vectors, and even trees can also implement sets!

Set ADT as a List



Set ADT as a Tree



	add	remove	contains
list/vector	$O(N)$	$O(N)$	$O(N)$
tree	$O(\log(N))$	$O(\log(N))$	$O(\log(N))$
lut	$O(1)$	$O(1)$	$O(1)$

Strengths of Lookup Tables

- Lookup tables have great time complexity!

Weaknesses of Lookup Tables

- I want a set for a larger range of numbers (e.g., from 0 to 1,000,000)
 - Lookup tables are _____
- I want a set of other types (e.g., strings)
 - Lookup tables are _____

2.4. Mitigating the Weaknesses of Lookup Tables

- Lookup tables have space complexity that is linear with k , the space of all possible unique inputs.

$k = 4$ and $N = 2$

$k = 1,000,000$ and $N = 2$

- How can we make the space complexity instead scale with N , the number of unique elements currently in the set?

Reduction Operations

- Define a reduce function that converts a large number into the range accessible by the internal array.
 - When the internal array fills up, double the size of the array.

Member function to reduce a large number

```
1 unsigned int reduce( unsigned int v ) {
```

Interface

```
1  class LookupTableIntOneMillion
2  {
3      public:
4          LookupTableIntOneMillion();
5
6          void add      ( unsigned int v );
7          void remove   ( unsigned int v );
8          bool contains ( unsigned int v );
9
10     private:
11         unsigned int reduce ( unsigned int v );
12         static const unsigned int size = 10;
13         bool m_table[size];
14     };

```

Implementation

```
1  LookupTableIntOneMillion::LookupTableIntOneMillion()
2  {
3
4      void LookupTableIntOneMillion::add( unsigned int v )
5      {
6          m_table[ reduce(v) ] = true;
7      }
8
9      void LookupTableIntOneMillion::remove( unsigned int v )
10     {
11         m_table[ reduce(v) ] = false;
12     }
13
14     bool LookupTableIntOneMillion::contains( unsigned int v )
15     {
16         return m_table[ reduce(v) ];
17     }

```

Implementation of New Member Function

```
18 unsigned int LookupTableIntOneMillion::reduce( unsigned int v
19 {
20     return v % size;
21 }
```

Main

```
1 int main( void )
2 {
3     LookupTableIntOneMillion table;
4
5     table.add( 967 ); // Reduces to index 7
6     table.add( 333 ); // Reduces to index 3
7     table.add( 555 ); // Reduces to index 5
8
9     return 0;
10 }
```

Space Complexity with a Reduction Function

- Ignoring conflicts, we now have space complexity linear with N
- Time complexity remains at O(1) access
- Achieving this space complexity cost us...
 - Increased work in order to double the size of the internal array
 - Increasing the work on each access
- Any other issues?

2.5. Mitigating the Weaknesses of Lookup Tables (cont.)

- Lookup tables are *inflexible* in storing non-integer types like strings.
- But strings, like any other type represented by a computer, are actually just numbers in the underlying representation!
- The key idea is to add a new function, a *hash* function, which converts a string or other arbitrary type into an integer.
 - Values of `char` type are just numbers!
 - The ASCII char ‘a’ is 97, ‘b’ is 98, etc...

Member function to convert a `char` array into an `unsigned integer`

```
1 unsigned int hash( char v[] ) {
```

Good Hash Functions

- What makes a hash function a “good” hash function?
 - Property 1: We want a *valid* hash function
 - Returns the same value on subsequent calls to the same item
 - For any equivalent objects $a == b$, their hashes are also equal

- Property 2: We want a hash function that provides *uniformity*
 - Maps the expected inputs as evenly as possible over the output range
 - Specifically, the hash result should not be a value (e.g., 100) more often

- Property 3: We want a hash function with $O(1)$ time complexity

Interface

```
1 class LookupTableStringHash
2 {
3     public:
4         LookupTableStringHash();
5         void add      ( std::string v );
6         void remove   ( std::string v );
7         bool contains ( std::string v );
8
9     private:
10        unsigned int reduce ( unsigned int v );
11        unsigned int hash   ( std::string v );
12        static const unsigned int size = 10;
13        bool m_table[size];
14 }
```

Implementation

```
1 LookupTableStringHash::LookupTableStringHash()
2 {
3
4     void LookupTableStringHash::add( std::string v )
5     {
6         m_table[ reduce( hash(v) ) ] = true;
7     }
8
9     void LookupTableStringHash::remove( std::string v )
10    {
11        m_table[ reduce( hash(v) ) ] = false;
12    }
13
14     bool LookupTableStringHash::contains( std::string v )
15    {
16        return m_table[ reduce( hash(v) ) ];
17    }
```

Implementation of New Member Functions

```
1 unsigned int LookupTableStringHash::reduce( unsigned int v )
2 {
3     return v % size;
4 }
5
6 unsigned int LookupTableStringHash::hash( std::string v )
7 {
8     unsigned int hashnum = 0;
9     for ( auto& c : v )
10        hashnum += c;
11     return hashnum;
12 }
```

Main

```
1 int main( void )
2 {
3     LookupTableStringHash table;
4
5     table.add( "c" );    // Hash 99 : Reduced to 9
6     table.add( "a" );    // Hash 97 : Reduced to 7
7     table.add( "t" );    // Hash 116 : Reduced to 6
8
9     table.add( "cat" ); // Hash: 99 + 97 + 116 = 312
10                           // Reduced to 2
11
12     return 0;
13 }
```

Summary

- We started off trying to mitigate the weaknesses of lookup tables

	add	remove	contains
list/vector	$O(N)$	$O(N)$	$O(N)$
tree	$O(\log(N))$	$O(\log(N))$	$O(\log(N))$
lut	$O(1)$	$O(1)$	$O(1)$
lut w/ hash & reduce	$O(1)$	$O(1)$	$O(1)$

Strengths of Lookup Tables with hash and reduce Functions

- Time complexity is still $O(1)$

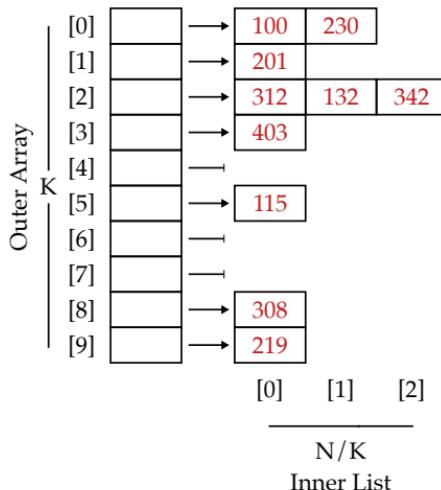
Mitigated Weaknesses of Lookup Tables with hash and reduce

- I want a set for a larger range of numbers (e.g., from 0 to 1,000,000)
 - Lookup tables were *space-inefficient*
 - Lookup tables with `hash` and `reduce` have space efficiency of $O(N)$
- I want a set of other types (e.g., strings)
 - Lookup tables were *inflexible*
 - Lookup tables with `hash` and `reduce` can handle any type as long as it provides a valid `hash` function
- What do we do about collisions?

3. Hash Tables

- Hash tables are lookup tables with hash and reduce functions that also handle collisions.
- Two common approaches for handling collisions
 - Separate chaining (usually with linked lists)
 - Open addressing (usually with linear probing)

3.1. Collision Handling with Separate Chaining



- The internal array elements are lists in a separate chaining scheme
- With a good hash function, the N elements are spread evenly over the K buckets.
- Adding a new element to the hash table is now a push back on a list
- What does this do to our time complexity?

```
1  class List
2  {
3      public:
4
5      class Itr
6      {
7          public:
8              Itr( Node* node_p );
9              void         next();
10             std::string& get();
11             bool        eq( const Itr& itr ) const;
12
13     private:
14         friend class List;
15         Node* m_node_p;
16     };
17
18     Itr begin();
19     Itr end();
20
21     List();
22     ~List();
23     void push_front( std::string& v );
24
25     private:
26
27     struct Node
28     {
29         std::string value;
30         Node*       next_p;
31     };
32
33     Node* m_head_p;
34 }
```

```
1 List::Itr    operator++( List::Itr& itr, std::string );
2 List::Itr&  operator++( List::Itr& itr );
3 std::string& operator* ( List::Itr& itr );
4 bool        operator!=( const List::Itr& itr0,
5                           const List::Itr& itr1 );
```

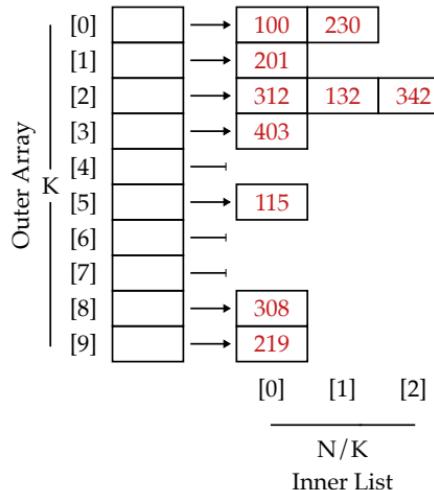
- With the above List class (provided exactly as discussed in previous lectures), implement the add function of a hash table that handles collisions with chaining using linked lists.
- Note that a templated argument **Hasher** is often used to provide a class with a hash function.
 - The following class “has-a” **Hasher** named **hash**.
 - The **Hasher** has an overloaded call operator (i.e., **operator()**) that does the hash (e.g., calling **hash** with a string returns an unsigned integer).
- Below is the interface of the hash table. Assume that calling the **Hasher** object using **hash()** will create some collisions.

```
1 template < typename Hasher >
2 class HashTableChaining
3 {
4     public:
5         HashTableChaining();
6         void add      ( const std::string& v );
7         void remove   ( const std::string& v );
8         bool contains ( const std::string& v );
9
10    private:
11        unsigned int reduce( unsigned int v );
12        Hasher hash;
13        static const unsigned int size = 10;
14        List m_table[size];
15    };
```

```
1 void add( const std::string& v ) {
```

Summary

- How do hash tables compare to previous implementations of sets?



	add	remove	contains
list/vector	O(N)	O(N)	O(N)
tree	O(log(N))	O(log(N))	O(log(N))
lut	O(1)	O(1)	O(1)
lut w/ hash & reduce	O(1)	O(1)	O(1)
hash table	_____	_____	_____