

Name: _____

Start work early.

This assignment focuses on *hashing*, described in Chapter 20 of the textbook.¹ You should submit (1) a hardcopy of this assignment with your answers filled in neatly and (2) a well-packaged electronic submission with your answers to the programming components (as in previous assignments).

1. (1 pt.) Write your name in the space provided above.
2. (1 pt.) Package and submit your solutions to the programming questions as described above and in class. *After* submitting your work, complete the following:
File name: _____ Size, in bytes: _____
MD5 checksum: _____
Timestamp: _____
3. (6 pts.) Compare search trees and hashing as data structures for collections of keys by naming (1) three operations that are supported by both and (2) three operations that are supported by search trees, but not by hashing.
4. (6 pts.) Refer to the simple method for storing a set of 16-bit integers described at the beginning of Section 20.1. How much storage (in bytes) would be required if that method were used for 64-bit integers? Explain how you arrive at your answer.

¹Mark Allen Weiss, *Data Structures and Problem Solving Using Java*, 4th edition (Addison-Wesley, 2010).

5. (6 pts.) The book notes that a hash table may be used to implement a set in constant time per operation. Name three set operations that a hash-table may implement in constant time. Name three set operations that a hash table cannot easily implement in constant time. Justify your answers.

6. (6 pts.) What is the maximum number of collisions that may result from hashing n keys? Provide an explicit hash function that produces this number of collisions.

7. (6 pts.) If a hash table H of size m contains n keys, what are the minimum and maximum number of probes required to insert a new key into H ? Justify your answer.

8. (8 pts.) Using the method of the fourth paragraph of Section 20.1, map each of the 11 space-separated strings [from *A Fit of Rime Against Rime*² by Ben Johnson]:

Rime, the lack of finest wits,
That expresseth but by fits,

to integers. You may use a calculator or computer *and must show your work for at least the string **finest***. (Showing your work for the others is recommended, but not required. Note that there can be no partial credit awarded if you arrive at a wrong answer and do not show your work.) Assume a 7-bit ASCII character encoding; for convenience, a table of codes appears below.

0 NUL	16 DLE	32	48 0	64 @	80 P	96 ‘	112 p
1 SOH	17 DC1	33 !	49 1	65 A	81 Q	97 a	113 q
2 STX	18 DC2	34 "	50 2	66 B	82 R	98 b	114 r
3 ETX	19 DC3	35 #	51 3	67 C	83 S	99 c	115 s
4 EOT	20 DC4	36 \$	52 4	68 D	84 T	100 d	116 t
5 ENQ	21 NAK	37 %	53 5	69 E	85 U	101 e	117 u
6 ACK	22 SYN	38 &	54 6	70 F	86 V	102 f	118 v
7 BEL	23 ETB	39 ’	55 7	71 G	87 W	103 g	119 w
8 BS	24 CAN	40 (56 8	72 H	88 X	104 h	120 x
9 HT	25 EM	41)	57 9	73 I	89 Y	105 i	121 y
10 LF	26 SUB	42 *	58 :	74 J	90 Z	106 j	122 z
11 VT	27 ESC	43 +	59 ;	75 K	91 [107 k	123 {
12 FF	28 FS	44 ,	60 <	76 L	92 \	108 l	124
13 CR	29 GS	45 -	61 =	77 M	93]	109 m	125 }
14 SO	30 RS	46 .	62 >	78 N	94 ^	110 n	126 ~
15 SI	31 US	47 /	63 ?	79 O	95 _	111 o	127 DEL

²*Purely for pleasure*: You are encouraged, but not required, to read the rest of the poem. See also the commentary by Robert Pinsky, “I, Too, Dislike It,” *Slate* (2005), http://www.slate.com/articles/arts/culturebox/2005/04/i_too_dislike_it.html.

9. (8 pts.) Repeat Question 8 using the hash function of Figure 20.1 with `tableSize = 17`. As in that question, show your work at least for the string `finest`, here by noting the value of `hashVal` after each iteration of the for loop.

10. (8 pts.) Repeat Question 9 for the hash function of Figure 20.2. Justify the textbook's claim "... the result obtained by allowing overflow... not the same... every step" by providing a concrete example of different results.

11. (8 pts.) Repeat Question 10 for the hash function of Figure 20.3.

12. (4 pts.) Suppose we are writing code that must run in a programming environment that supports only addition and shifting of 16-bit binary numbers (no multiplication or other arithmetic operations). Provide efficient, detailed pseudocode for this environment that is equivalent to line 12 of Figure 20.2 in the textbook.

13. (8 pts.) Using Figure 20.5 in the textbook as a template, depict the result of inserting the strings of Question 8 into a hash table of size 17, using each of the three hash functions noted in Questions 9, 10, and 11. Use the linear probing method for collision resolution. It may be convenient to depict the tables in landscape mode on this page.

14. (8 pts.) Repeat Question 13 using the quadratic probing method for collision resolution.

15. (8 pts.) Repeat Question 13 using the separate chaining method for collision resolution, and a hash table of size 11.

For this question only, when using the hash function of Question 9, compute the index into the hash table by taking the remainder of that function's result and 11. (Computing hash-table indices in this way is a bad idea in general.³)

³Use the class newsgroup to explore why.

16. (8 pts.) Consider the following modification of the separate chaining method, yielding a method called *coalesced hashing*: Instead of using separate storage for the linked lists used by separate chaining, the linked lists are stored in the otherwise unused slots of the hash table. With this method, each slot in the hash table contains either a key (as in linear or quadratic probing) or a cell belonging to a linked list that is used for chaining (as in separate chaining). Note that the linked-list cell occupying a hash table slot does not necessarily contain a key that hashes to that slot.

Repeat Question 13 using this coalesced chaining method for collision resolution.

17. (100 pts.) Modify the textbook's implementation of the `HashSet` class⁴ to use the *coalesced hashing* technique of Question 16 instead of quadratic probing. You may use as much or as little of the textbook's code as you wish (with proper attribution, as always) but your implementation of `HashSet` must be *completely interchangeable* with the textbook's implementation. That is, any program that uses the textbook's implementation should work as before if all the files implementing the textbook's `HashSet` class are replaced with those implementing yours (with no other changes).

Some minor changes to this question may be announced in class later (mostly related to input-output and nothing that will take more than a few minutes to change—so finish up your work without waiting).

⁴Weiss, *op. cit.*, §20.4.1.