

15-122: Principles of Imperative Computation,
Spring 2011
Assignment 4: Data Structure Design,
Memory Layout, and Amortization

William Lovas (wlovas@cs) Bill Zorn (waz@andrew)

Out: Thursday, February 24, 2011

Due: Thursday, March 3, 2011

(Written part: before lecture,
Programming part: 11:59 pm)

1 Written: Potpourri (20 points)

The written portion of this week's homework will give you some practice working with hash tables, amortized analysis, memory management, debugging, and recursion. You can either type up your solutions or write them *neatly* by hand, and you should submit your work in class on the due date just before lecture begins. Please remember to *staple* your written homework before submission.

1.1 Hash Tables

Exercise 1 (4 pts). In Java, strings are hashed using the following function:

$$(s[0] * 31^{n-1} + s[1] * 31^{n-2} + \dots + s[n-2] * 31^1 + s[n-1] * 31^0) \% m$$

where $s[i]$ is the ASCII code for the i th character of string s , n is the length of the string, and m is the number of chains in the hash table.

- If 15122 strings were stored in a hash table with 4401 chains, what would the load factor of the table be? If the strings above were equally distributed in the hash table, what does the load factor tell you?
- Using the hash function above with a table size of 4401, give an example of two strings that would "collide" in the hash table and would be stored in the same chain. Show your work. (Think of short strings please!)

- (c) Why did the Java designers use 31 in the formula above? Investigate several sources for your answer and report on the validity of the sources. (For this question, you are free to use the Internet, but be sure to cite your sources!)

1.2 Amortized Analysis

Exercise 2 (4 pts). A ternary counter is a data structure that stores an integer as powers of three. We can think of one exactly like the binary representation of an integer, except that each *trit* (ternary digit) holds a value between 0 and 2 instead of 0 and 1, and the place value for trit i is given by 3^i instead of 2^i . For example, a ternary counter in the state 1210 stores the value $1 \times 3^3 + 2 \times 3^2 + 1 \times 3^1 + 0 \times 3^0 = 48$.

For purposes of analysis, assume that modifying one trit of the counter costs exactly one instruction.

- (a) Consider a ternary counter with an *increment* operation, which increases the value of the counter by 1, and a *decrement* operation, which decreases the value of the counter by 1. Show that for a k -trit counter, the worst-case cost of a series of n operations is $O(nk)$ instructions.
- (b) Consider a ternary counter with only an increment operation but with a *reset* operation that sets all of the trits to 0. Using amortized analysis, show that for a k -trit counter, the worst-case cost of a series of n operations is $O(n)$ instructions.

1.3 Memory Management

Constant factors can influence the memory efficiency of data structures and algorithms, even if the asymptotic complexity is the same. For example, a stack can be internally implemented in multiple ways without changing the interface functions you use to manipulate it. The implementation we saw in class uses linked lists:

```
struct list_stack {
    struct list * top;
    struct list * bottom;
};

struct list {
    elem data;
    struct list * next;
};
```

Another possible implementation uses unbounded arrays instead:

```
struct uba_stack {
    struct ubarray * U;
};
```

```

struct ubarray {
    int limit;          /* limit > 0 */
    int size;          /* 0 <= size && size <= limit */
    elem[] elems;     /* \length(elems) == limit */
};

```

Assume for the moment that both stack implementations store characters as their data elements: `typedef char elem` has been defined previously in the code. Recall that in C_0 , a `char` is represented using 8 bits, a pointer is represented using 64 bits, and an `int` is represented using 32 bits. An array of characters of size n takes exactly $8n + 64$ bits— $8n$ for the characters and 32 bits each for two ints: one to record the length of the array and one to record the size of its elements. The size of a struct is the sum of the sizes of its elements.

Exercise 3 (4 pts). Consider some program that uses a stack of characters. If the program can allocate only 1 MB (2^{20} bytes, or 8×2^{20} bits) of memory for this stack, what is the maximum number of characters that it can store using a struct `list_stack`? How about using a struct `uba_stack`? (Show your work.)

1.4 Debugging

The following helper function is intended to take a structure cell containing a stack of strings and print it out, along with some general identifying information about it. You may assume that the code compiles, that type `elem` is defined as `string`, and that stacks are implemented using linked lists as in lecture.

```

typedef struct wordcell * cell;
struct wordcell {
    int ID;
    int size;
    stack words;
};

bool is_cell(cell W) {
    if (W == NULL) return false;
    if (!(0 <= W->ID && W->ID < 1000)) return false;

    return (W->words == NULL && W->size == 0)
        || (is_stack(W->words) && stack_size(W->words) == W->size);
}

```

```

void printcell(cell W)
//@requires is_cell(W);
{
    print("Cell: "); printint(W->ID); print("\n");
    print("size: "); printint(W->size); print("\n");

    list place = W->words->top;

    while (place != W->words->bottom) {
        println(place->data);
        place = place->next;
    }
}

```

Exercise 4 (4 pts).

- (a) Devise a test case that causes `printcell` to segmentation fault. You may write out this case as a short `main()` function. Assume that the program is compiled with the `-d` flag, and that a contract failure is different from a segmentation fault. **Hint:** What are all the ways a segmentation fault can occur in a C_0 program?
- (b) Change `printcell` so that this case does not cause a segmentation fault. The specification function `is_cell` is correct—do not change it! **Hint:** Your change should not affect more than a few lines.

1.5 Recursion

Exercise 5 (4 pts). Recursive functions can be used to construct data structures of recursive type, like linked lists.

- (a) Write a recursive function matching the following specification:

```

struct list * from(int start, int end)
    //@requires start <= end;
    //@ensures list_length(\result) == end - start;
    ;

```

The result of a call like `from(1, n)` should be a list containing the numbers from 1 to $n - 1$ in order, terminated by a NULL pointer:



- (b) Prove that if the pre-condition holds, the returned result satisfies the post-condition. When reasoning inside the body of a function f , if you *know* that the pre-condition of a recursive call to f holds before the call, you may *assume* the post-condition of that call holds after the call.

2 Programming: Implementing a Text Editor (30 points)

In the programming portion of this week's assignment, you'll practice a bit of data structure design by implementing the core data structure for a text editor: the edit buffer. Using contracts to guide your coding is a major theme, so be sure to *read the specifications carefully*—you'll thank yourself later!

You should submit your code electronically by 11:59 pm on the due date. Detailed submission instructions can be found below.

Starter code. Download the file `hw4-starter.zip` from the course website. It contains definitions of the edit buffer data structures (`gapbuf.c0` and `tbuf.c0`) and some code for visualizing them (`visuals.c0`), which you may find useful for testing your implementation.

Compiling and running. Compile and test your code using either `cc0` or `coin`. Don't forget to include the `-d` switch to enable dynamic checking of contracts!

For details on how we will compile your code, see the file `COMPILING.txt` included in the starter code. **Warning:** *You will lose credit if your code does not compile.*

Submitting. Once you've completed some files, you can submit them by running the command

```
handin -a hw4 <file1>.c0 ... <fileN>.c0
```

You can submit files as many times as you like and in any order. When we grade your assignment, we will consider the most recent version of each file submitted before the due date. If you get any errors while trying to submit your code, you should contact the course staff immediately.

The submission verifier framework for this assignment will check *only* that your code compiles: you are responsible for testing your code thoroughly!

Annotations. Be sure to include appropriate `//@requires`, `//@ensures`, `//@assert`, and `//@loop_invariant` annotations in your program. You should write these as you are writing the code rather than after you're done: documenting your code as you go along will help you reason about what it should be doing, and thus help you write code that is both clearer and more correct. **Annotations are part of your score for the programming problems; you will not receive maximum credit if your annotations are weak or missing.**

Style. Strive to write code with *good style*: indent every line of a block to the same level, use descriptive variable names, keep lines to 80 characters or fewer, document your code with comments, etc. We will read your code when we grade it, and good style is sure to earn our good graces. Feel free to ask on the course bboard (`academic.cs.15-122`) if you're unsure of what constitutes good style.

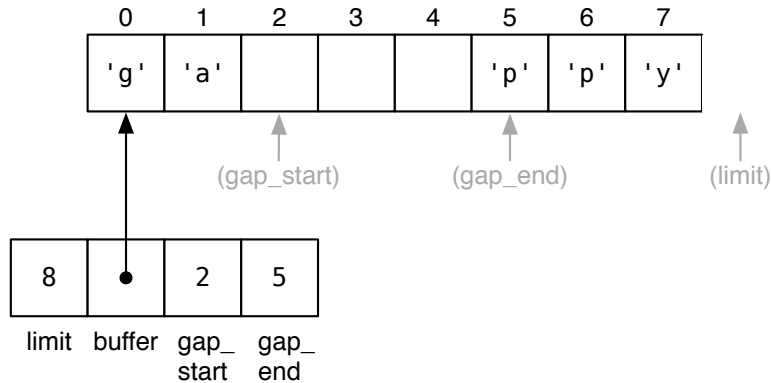


Figure 1: A gap buffer in memory.

2.1 Gap Buffers (`gapbuf.c0`)

One typical data structure used to implement an edit buffer is a *gap buffer*. A gap buffer is a generalization of an unbounded array: whereas an unbounded array allows for efficient addition and deletion of elements from the end only, a gap buffer allows efficient addition and deletion of elements from a *gap* somewhere in the middle. In addition, the gap buffer offers efficient operations to move the gap to the left and to the right.

A gap buffer is represented in memory by an array of elements stored along with its size (`limit`) and two integers representing the beginning (inclusive, `gap_start`) and end (exclusive, `gap_end`) of the gap (see Figure 1).

```
typedef struct gapbuf * gapbuf;
struct gapbuf {
    int limit;           /* limit > 0 */
    char[] buffer;     /* \length(buffer) == limit */
    int gap_start;     /* 0 <= gap_start */
    int gap_end;       /* <= gap_end <= limit */
};
```

A valid gap buffer is non-NULL, has a strictly positive `limit` which correctly describes the size of its array, and has `gap_start` and `gap_end` which are valid for the array.

Task 1 (3 pts). Implement a specification function `bool is_gapbuf(gapbuf G)` formalizing the gap buffer data structure invariants.

For the case of a gap buffer representing an edit buffer, the elements are characters, and we can visualize the gap buffer as a string of characters with a gap in the middle:

the s[.]pace race

The above represents a gap buffer of size 16 containing the string “the space race”, with the gap situated between the “s” and the “p” in “space”. To delete a character, we simply expand the gap:

```
the [...]pace race
```

To move the gap, we copy a character across it:

```
the p[...]ace race
```

To insert a character, we write it into the gap, shrinking it by one:

```
the pe[...]ace race
```

The gap can be at the left end of the buffer,

```
[...]mineshaft gap
```

or at the right end of the buffer,

```
minecraft gap[...]
```

and a buffer can be empty,

```
[.....]
```

or it can be full,

```
mind the ga[]p sir
```

Task 2 (2 pts). Implement the following utility functions on gap buffers:

Function:	Returns true iff...
<code>bool gapbuf_empty(gapbuf G)</code>	the gap buffer is empty
<code>bool gapbuf_full(gapbuf G)</code>	the gap buffer is full
<code>bool gapbuf_at_left(gapbuf G)</code>	the gap is at the left end of the buffer
<code>bool gapbuf_at_right(gapbuf G)</code>	the gap is at the right end of the buffer

Task 3 (5 pts). Implement the following interface functions for manipulating gap buffers:

<code>gapbuf gapbuf_new(int limit)</code>	Create a new gapbuf of size <i>limit</i>
<code>void gapbuf_forward(gapbuf G)</code>	Move the gap forward, to the right
<code>void gapbuf_backward(gapbuf G)</code>	Move the gap backward, to the left
<code>void gapbuf_insert(gapbuf G, char c)</code>	Insert the character <i>c</i> before the gap
<code>void gapbuf_delete(gapbuf G)</code>	Delete the character before the gap

If an operation cannot be performed (e.g., moving the gap backward when it’s already at the left end), it should leave the gap buffer unchanged.

All functions should require and ensure the data structure invariants. Furthermore, the gap buffer returned by `gapbuf_new` should be empty. Use these facts to help you write your code, and document them with appropriate assertions.

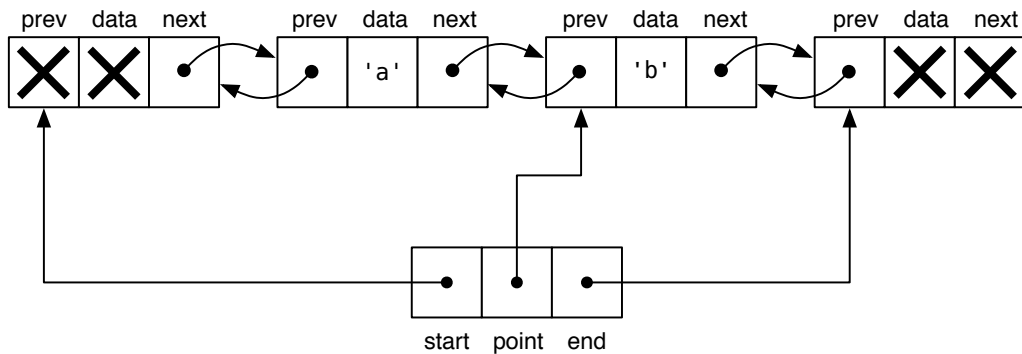


Figure 2: An editable sequence as a doubly-linked list in memory.

2.2 Doubly-Linked Lists (`tbuf.c0`)

Another data structure that could be used to represent an edit buffer is a *doubly-linked list*. We have seen singly-linked lists used to represent stacks and queues—sequences of nodes, each node containing some data and a pointer to the next node—but such a structure cannot provide constant time insertions or deletion in the middle of the list. The nodes of a doubly-linked list contain a data field just like those of a singly-linked list, but in contrast, the doubly-linked nodes contain *two* pointers: one to the next element (*next*) and one to the *previous* (*prev*).

An editable sequence is represented in memory by a doubly-linked list and three pointers: one to the start of the sequence, one to the end of the sequence, and one to the distinguished point node where updates may take place (see Figure 2). We employ our usual trick of terminating the list with “dummy” nodes whose contents we never inspect.

```
typedef struct dll * dll;
struct dll {
    elem data;
    dll next;
    dll prev;
};

typedef struct tbuf * tbuf;
struct tbuf {
    dll start;
    dll point;
    dll end;
};
```

We can visualize a well-formed doubly-linked list as the sequence of its data elements

with terminator nodes at the end and one distinguished element.

```
** <--> 'a' <--> 'b' <--> **
```

For now, we do not concern ourselves with the type of the data elements: basic doubly-linked list functions are agnostic to it anyway.

The key invariant in a well-formed doubly-linked list is that the next links proceed from the start node to the end node, passing point node along the way, and that the prev links mirror the next links. Additionally, we require that the point be a distinct node from both the start and the end nodes, i.e., that the list be non-empty.¹

Task 4 (3 pts). Write a specification function `bool is_linked(tbuf B)` formalizing the linking invariants on a doubly-linked list text buffer. (You are not required to check for circularity, but you may find it useful to do so.)

Task 5 (3 pts). Implement the following utility functions on doubly-linked text buffers:

Function:	Returns true iff...
<code>bool tbuf_at_left(tbuf B)</code>	the point is at the far left end
<code>bool tbuf_at_right(tbuf B)</code>	the point is at the far right end

and the following interface functions for manipulating doubly-linked text buffers:

<code>void tbuf_forward(tbuf B)</code>	Move the point forward, to the right
<code>void tbuf_backward(tbuf B)</code>	Move the point backward, to the left
<code>void tbuf_delete_point(tbuf B)</code>	Remove the point node from the list

As above, if an operation cannot be performed, it should leave the text buffer unchanged. When deleting the point, the new point may be either to the right or to the left of the old one.

These functions should require and preserve the linking invariant you wrote above, and you should both document this fact and use it to help write the code. (Be especially careful when implementing deletion!)

2.3 Putting It Together (tbuf.c0)

While either gap buffers or doubly-linked lists of characters could be used to represent an edit buffer—and your implementations could easily be extended to implement the general edit buffer interface—neither strategy is particularly realistic. One large edit buffer requires the entire file contents to be stored in a single, contiguous block of memory, which can be difficult for large files, and furthermore, although all operations are amortized constant time, long pauses could still hamper interactivity. A doubly-linked list of single characters, while offering truly constant-time operations

¹And of course, for any of these to hold, several pointers must be non-NULL!

and allowing the file to be split across several chunks of memory, involves far too much space overhead at two pointers per character.

Instead, a more realistic strategy is to combine the two ideas by implementing a doubly-linked list of fixed-size gap buffers. The contents of a text buffer represented in this way is simply the concatenation of the contents of its requisite gap buffers, in order from the start to the end. When a gap buffer is full, we can split it in two:

```

** <--> splitend[] <--> **
insert 's': ** <--> spli[....] <--> tends[...] <--> **

```

and when one becomes empty, we can delete it:

```

** <--> deletio[.] <--> n[.....] <--> **
delete: ** <--> deletio[.] <--> **

```

To move, we use a combination of gap buffer motion and doubly-linked list motion:

```

** <--> just_a[.] <--> j[....]ump <--> **
move ←: ** <--> just_a[.] <--> [....]jump <--> **
move ←: ** <--> just_a[.]_ <--> [....]jump <--> **

```

An invariant that arises from this representation is that a text buffer is always *aligned*: every gap buffer *before* the point is non-empty with its gap to the *right*, and every gap buffer *after* the point is non-empty with its gap to the *left*.

```

** <--> lawfu[...] <--> l_evil_c[] <--> hao[....]t <--> [.]tic_good <--> **

```

Additionally, all the gap buffers are themselves well-formed, and they all have the same size (8 characters in the diagrams here, but you should use 16 for your implementation).

Task 6 (3 pts). Implement a specification function `bool is_aligned(tbuf B)` formalizing the alignment invariant. Ensure that each gap buffer has size 16.

For a text buffer to be fully well-formed, it must either be the empty text buffer:

```

** <--> [.....] <--> **

```

or it must be properly linked and aligned and have a non-empty point.

Task 7 (3 pts). Implement the following specification functions on text buffers:

Function:	Returns true iff...
<code>bool tbuf_empty(tbuf B)</code>	the text buffer is empty
<code>bool is_tbuf(tbuf B)</code>	the text buffer satisfies all invariants

and a text buffer constructor:

`tbuf tbuf_new()` constructs a new, empty text buffer

(**Aside:** At this point, you might consider re-declaring the `tbuf_delete_point` function with a stronger contract: you can now give it a much more precise specification, and doing so will help you later! Think about what further invariants it preserves beyond those specified earlier.)

To split a full gap buffer, we have to copy each half of the character data into one of two new gap buffers, taking special note of where the new gaps should end up. The following diagrams may help you visualize the intended result:

<i>full buffer:</i>	<code>abc[]defghABCDEFGH</code>	<i>full buffer:</i>	<code>stuvwxyzSTUV[]WXYZ</code>
<i>splits into:</i>	<code>abc[.....]defgh</code> <code>[.....]ABCDEFGH</code>	<i>splits into:</i>	<code>stuvwxyz[.....]</code> <code>STUV[.....]WXYZ</code>

We can then link the new gap buffers into the doubly-linked list, taking care to preserve the text buffer invariants.

Task 8 (4 pts). Implement a function `tbuf_split_point(tbuf B)` which takes a valid text buffer whose point is full and turns it into a valid text buffer whose point is not full.

Task 9 (4 pts). Implement the following interface functions for manipulating text buffers:

<code>void forward_char(tbuf B)</code>	Move the point forward, to the right
<code>void backward_char(tbuf B)</code>	Move the point backward, to the left
<code>void insert_char(tbuf B, char c)</code>	Insert the character <i>c</i> before the point
<code>void delete_char(tbuf B)</code>	Delete the character before the point

If an operation cannot be performed (e.g., moving the point backward when it's already at the left end), it should leave the text buffer unchanged.

After you've completed your text buffer implementation and tested it thoroughly, you can try it out interactively by compiling against `e0.c0` from the starter code, a comically minimalist text editor front-end called `E0`. Enjoy the hard-won fruits of your careful programming labor!

2.4 Judges' Prize: Extending the Editor

Extend the `E0` editor implementation with some interesting features. A few suggestions, to pique your imagination: a better display algorithm, line motion, more editing commands, copy and paste—be creative! Feel free to extend the data structures in any way necessary to support your changes effectively. Submit your modified implementation as files named `judges-*.c0` and include a `judges-README` file explaining your work. The coolest extension—both in terms of the interactive editing experience and the supporting data structures and algorithms—will receive a prize!

3 Survey

In order to better serve you, we ask that you submit along with your programming portion a short README file answering the following two questions:

- (a) How long did it take you to complete the written portion of this homework?
- (b) How long did it take you to complete the programming portion of this homework?

You may also include any explanations of your code itself, or any other comments that you feel might be relevant.