

# Useful data structures and their implementations in C

Julia PHAM BA NIEN

April 21, 2024

# Contents

<b>1</b>	<b>What's a data structure</b>	<b>4</b>
<b>2</b>	<b>Meta structures</b>	<b>4</b>
2.1	Primitives . . . . .	4
2.2	Product types . . . . .	4
2.3	Sum types . . . . .	5
2.3.1	Enums . . . . .	5
2.3.2	Tagged unions . . . . .	5
<b>3</b>	<b>Container structures</b>	<b>5</b>
3.1	Linked lists . . . . .	6
3.2	Array . . . . .	7
3.2.1	Sentinel valued array . . . . .	7
3.2.2	Fat pointers . . . . .	7
3.2.3	Vectors . . . . .	8
3.3	Stacks . . . . .	9
3.4	Queues . . . . .	9
3.5	Hashmaps . . . . .	10
3.5.1	Hashing . . . . .	11
3.5.2	Transposition tables . . . . .	12
3.5.3	Perfect hashmaps . . . . .	12
3.5.4	Sets . . . . .	13
3.6	Bitsets . . . . .	13
3.7	Run length encodings . . . . .	14
<b>4</b>	<b>Graphs</b>	<b>14</b>
4.1	General graphs . . . . .	15
4.1.1	Adjacency lists . . . . .	15
4.1.2	Adjacency matrices . . . . .	15
4.1.3	Actual graph definition . . . . .	16
4.2	Rooted trees . . . . .	16
4.2.1	General rooted trees . . . . .	16
4.2.2	Parent oriented trees . . . . .	17
4.2.3	Complete binary trees . . . . .	17
4.2.4	Heaps . . . . .	18
<b>5</b>	<b>Priority Queues</b>	<b>20</b>
<b>6</b>	<b>Disjoint set unions</b>	<b>20</b>
<b>7</b>	<b>Range queries</b>	<b>25</b>
7.1	Prefix sums . . . . .	25
7.2	Segment trees . . . . .	25
7.3	Fenwick/BIT trees . . . . .	27

<b>8</b>	<b>Memory management</b>	<b>28</b>
8.1	Safe guards . . . . .	28
8.1.1	Rcs . . . . .	28
8.1.2	Mutexes . . . . .	29
8.2	Allocators . . . . .	29
8.2.1	Page allocators . . . . .	30
8.2.2	Fix buffer allocators . . . . .	31
8.2.3	Arena allocators . . . . .	32
8.2.4	General purpose allocators . . . . .	33
<b>9</b>	<b>After words</b>	<b>34</b>

I have decided to show implementation in C, I don't particularly like the Knuth's pseudocode style, and choosing a low level language shows the small memory-related implementations details.

I will use `T` as a type place holder, because data structures often are generics, and I will use a `type<T>` syntax to specify what the underlying `T` would be. (even though it isn't truly valid C), also, I may not use struct names when specifying `struct` with its parameters, and will use `structure_name vals ...` to indicate that it's a slight modification of `structure` with fields being `vals ...` (other than the structure specific ones), also I'll use some non standard functions like `swap`, you know what it means, so do as if I defined it.

Data structures are *incredibly* related to algorithms, I'll assume some familiarity with the  $O$  notation and analysis of algorithms.

Distinction between data structures are from me, their borders are quite blurry, but I think it still makes sense for the most part.

I won't show any proofs for the time complexities, or anything really.

## 1 What's a data structure

A data structure in a way to structure data.

Though, this answer lacks a lot, because the question isn't the right one. (*check-mate person writing this exact guide*)

I think it's important to see data structure by the view of algorithms, and thus exploring *why* they are.

Data structures are ways to structure data to have *efficient* ways to do the operations you want to perform.

Thus, you should choose your data structures depending on *what* you want to perform, and how often you perform it.

That's why there are so many of them, because there are different use cases.

## 2 Meta structures

Meta structures are what the language offers to represent structures.

### 2.1 Primitives

Primitives are structures which represent *values* and wick looking at its parts doesn't make sense.

For example, `int` is a primitive, though strings aren't (because they are meant to see parts of it, its characters).

Even such small types have structure, which have different trade off in speed, memory usage and precision.

### 2.2 Product types

It's the normal way to put multiple values of different types together.

As a structure, we only care when putting only a finite amount of types together, so that in C terms, it's a `struct`. (In some languages, it's possible to put an infinite amount of value together, those are the case for functions which return type depends of its input in some languages).

## 2.3 Sum types

Sum types are types which are the union of multiple other types. And they are instanciated by having a variable.

### 2.3.1 Enums

Enums are sum types which have some fixed known to be values. In C, they are litteraly `enums`.

### 2.3.2 Tagged unions

Tagged unions are the “true” sum types (other that they are the sum of finite amount of types), where you can have more complexe stuffs than predetermined list of values.

The idea is that you will have a “tag”, which will tell you what the structure of the rest is to be read as.

A very simple example is making an `int_or_float_t`.

```
1 enum int_or_float_e {
2     INT,
3     FLOAT,
4 };
5
6 struct int_or_float_s {
7     enum int_or_float_e tag;
8     union {float f; int i;} value;
9 };
```

And you know if you need to access `.value.i` or `.value.f` depending on `.tag`.

It's not as good as with other languages as it's not checked, but it is useful.

A particular one that is quite useful is `option`, due to its simplicity you can make it without unions nor enums.

```
1 struct option_s {
2     bool has_value;
3     T value;
4 }
```

Where `.value` can be left uninitialised or undefined of `.has_value` is `false`.

## 3 Container structures

A container structure is one which its goal is to put multiple values of the same type into one structure, so that you can access them, pass them around, etc.

### 3.1 Linked lists

The idea is that each item has a pointer to the next one, unless you're the last value.

Adding a new element is  $O(n)$ , accessing the  $i$ -th element is  $O(1 + i)$ , and also it's not very compact memory wise, and will cache miss a lot.

In general it's not such a great data structure, it is mostly useful when you want an expandable container where its elements are very expensive to copy.

Implementation:

```
1 typedef struct linked_list_s {
2     T value;
3     struct linked_list_s *next;
4 } *linked_list_t;
5
6 linked_list_t init_linked_list() {
7     return NULL;
8 }
9
10 // 0-based indexing
11 T get_ith(linked_list_t l, int i) {
12     if (l == NULL)
13         exit(1); // access out of bound, deal with it as you want
14     if (i == 0)
15         return l->value;
16     return get_ith(l->next, i-1);
17 }
18
19 /* In imperative form:
20 T get_ith(linked_list_t l, int i) {
21     while (l && i) {
22         l = l->value;
23         --i;
24     }
25     if (!l)
26         exit(1);
27     return l->value;
28 }
29 */
30
31 linked_list_t append(linked_list l, T val) {
32     if (!l) {
33         linked_list_t new = (linked_list_t) malloc(sizeof(linked_list_s));
34         new->value = val;
35         new->next = NULL;
36         return new;
37     } else {
38         l->next = append(l->next, val);
39         return l;
40     }
41 }
42
43 void free_list(linked_list_t l) {
44     if (l) {
45         free_list(l->next);
```

```

46     free(l);
47 }
48 }

```

Note that for appending, it is better to have a reference to the last element, so that it becomes an  $O(1)$  operation, but then it's not truly a linked list anymore. (also, yes I know that there are ways to make generic linked list, some aren't the best (`void *value; struct linked_list *next;`) and quite good one do some "advanced" C tricks (<https://felipec.github.io/good-taste/parts/1.html>), I don't cover it because it doesn't matter and it's not translatable directly into most other languages.)

## 3.2 Array

An array is a continuous memory blob.

There are some very good part about it, indexing is an  $O(1)$  operation, so it's the backbone of a lot of other structures.

Though, it has a huge problem if taken as is, you don't know if your indexing is valid unless you know its size.

There are two main way to have some size indication.

### 3.2.1 Sentinel valued array

The idea is that at the end of your array, you put a value that is supposed to be impossible to get, so once you get there, you know it's the end.

For example, common such values are `'\0'` for strings or `-1` for positive integer arrays.

Having this value is very useful for the implementation of some algorithms, such as doing a recursive decent parser.

Though, if you want to check your indexing, you have to go through all values to see if you are indexing after the sentinel value, making it an  $O(1 + \min(n, i))$  operation.

### 3.2.2 Fat pointers

The idea of a *fat pointer* is that it is a pointer, to which you couple with the size of the array underneath.

Basically, it is

```

1 struct fat_s {
2     T *p;
3     size_t size;
4 };

```

It doesn't require having an impossible value, has  $O(1)$  out of bound check and  $O(1)$  knowing the length.

A very good feature of it is that it can do  $O(1)$  sublististing, and that without any (heap) memory allocation (and not allocating is a good way to not have memory problems).

```

1 // 0 based indexing
2 struct fat_s sublist(struct fat_s of, size_t start, size_t end) {
3     if (end > of.size - 1)
4         exit(1); // out of bound sublisting
5     struct fat_s new;
6     new.size = end - start + 1;
7     new.p = of.p + start;
8     return new;
9 }

```

Be careful when freeing though, as you have indirect references to the same memory.

### 3.2.3 Vectors

Now, all fun and good, but adding a new element to an array seems *horrible*! You have to allocate a new bigger array, and copy all your previous elements to your new one, thus being an  $O(n)$  operation.

Thankfully some smart people made a quite good in general data structure to deal with it, the “vector”.

The idea is that you want to reallocate as little as possible.

We won’t be using the standard idea of the Landau’s  $O$  notation, we’ll consider the “amortized” time, aka. we’ll look at how it behaves when adding  $n$  element (from empty), not a single one.

Let’s analyse some ideas on how to grow the size of the array.

Adding one to the size: the operation is then  $\sum_{k=1}^n O(k) = O(n^2)$  (because  $\sum_{k=1}^n k = \frac{n(n+1)}{2} = \frac{n^2}{2} + \frac{n}{2} = O(n^2)$ ).

Now, what about magically increasing the array to the right size ? It is then  $\sum_{k=1}^n O(1) = O(n)$ .

Though, we don’t know the size, but at least we know that the operation is  $\Omega(n)$ , so that’s what we’ll aim for.

What about increasing by multiplying by a constant greater (or equal) than 2 ? We’ll call this constant  $c$ .

It is then  $\sum_{k=1}^n O(1) + \sum_{k=0}^{\lfloor \log_c(n) \rfloor} O(c^k) = O(n) + \sum_{k=0}^{\lfloor \log_c(n) \rfloor} O(c^k) = O(n)$

Because  $\sum_{k=0}^{\lfloor \log_c(n) \rfloor} c^k = \frac{1-c^{1+\lfloor \log_c(n) \rfloor}}{1-c} = O(c^{1+\lfloor \log_c(n) \rfloor}) = O(c^{\lfloor \log_c(n) \rfloor}) = O(n)$ .

Thus it is optimal in term of time complexity (and exercise: show that it is optimal in term of space complexity), so good enough, we’ll use that scheme of allocation.

We’ll keep in mind the actual size of the array and the length of it (different here, because you may over allocate).

Implementation-wise it goes as so (here  $c = 2$ , though literature suggest that  $c = 3$  is better):

```

1 struct vector_s {
2     T *data;
3     size_t len;
4     size_t ize;
5 }
6
7 void grow(struct vector_s *t) {

```



```

8  size_t new_size = t->size == 0 ? 1 : 2*t->size;
9  T *new_data = malloc(new_size * sizeof(T));
10 if (t->data) {
11     memcpy(new_data, t->data, t->len * sizeof(T));
12     free(t->data);
13 }
14 t->data = new_data;
15 t->size = new_size;
16 }
17
18 void add(struct vector_s *t, T el) {
19     if (t->len == t->size)
20         grow(t);
21     t->data[(t->len)++] = el;
22 }
23
24 T get(struct vector_s *t, int n) {
25     return t->data[n];
26 }
27
28 T pop(struct vector_s *t) {
29     return t->data[--(t->len)];
30 }
31
32 T top(struct vector_s *t) {
33     return t->data[(t->len) - 1];
34 }
35
36 struct vector_s init_vector() {
37     struct vector_s v;
38     v->size = 0;
39     v->len = 0;
40     v->data = NULL;
41     return v;
42 }
43
44 int free_vector(struct vector_s *l) {
45     if (l->data)
46         free(l->data);
47     l->size = 0;
48     l->len = 0;
49 }

```

It is a very common data structure in most of programming languages.

### 3.3 Stacks

A stack is a data structure last in, first out (LIFO), the most common way to implement it is to just use a vector underneath, it trivially accomodate all the stack operations so I won't explain further.

### 3.4 Queues

A queue is a first in, first out (FIFO) structure, a quite common way to implement it is to use 2 vectors underneath.

A basic implementation goes as follow

```
1 typedef struct queue_s {
2     size_t read_position;
3     struct vector_s *to_read;
4     struct vector_s *to_fill;
5 } *queue_t;
6
7 queue_t init_queue() {
8     queue_t q = malloc(sizeof(struct queue_s));
9     q->read_position = 0;
10    q->to_read = malloc(sizeof(struct vector_s));
11    q->to_fill = malloc(sizeof(struct vector_s));
12
13    *(q->to_read) = init_vector();
14    *(q->to_fill) = init_vector();
15
16    return q;
17 }
18
19 void push_queue(queue_t q, T element) {
20     add(q->to_fill, T);
21 }
22
23 T pop_queue(queue_t q) {
24     if (q->to_read->len == 0 && q->to_fill->len == 0)
25         exit(1); // no element to pop
26
27     if (q->read_position < q->to_read->len)
28         return get(q->to_read, q->read_position++);
29
30     // swap the reader and writer
31     q->read_position = 0;
32     swap(&q->to_fill, &q->to_read);
33
34     return pop_queue(q);
35 }
36
37 void free_queue(queue_t q) {
38     free_vector(q->to_read);
39     free_vector(q->to_fill);
40     free(q);
41 }
```

All operations being  $O(1)$  amortized.

### 3.5 Hashmaps

There are a lot of different implementation of hashmaps, though, they all come down from the same idea, the *hash*.

It's a function, taking T and returning a `size_t`.

The idea is that you first find in which congruency class induced by hash you element is, and then do your operations in the class.

Because memory isn't infinite, you store the number of class you have,  $n$ , and you'll  $\% n$  the class number given by  $n$ .

So a basic implementation of an hashmap would be:

```

1 typedef fat<list {
2     struct {
3         K key;
4         V val;
5     }
6 }> hash<K, V>;
7
8 void set(hash h, K key, V val) {
9     size_t class_num = hash(key) % h.size;
10    list {
11        K key;
12        V val;
13    } class = hash.p[class_num];
14    while (class->next != NULL) {
15        if (class->val.key == key) {
16            class->val.val = val;
17            return;
18        }
19    }
20    struct {
21        K key;
22        V val;
23        list *next;
24    } *new = malloc(sizeof(struct {
25        K key;
26        V val;
27    }));
28    new->key = key;
29    new->val = val;
30    new->next = NULL;
31    class->next = new;
32 }
33
34 void get(hash h, K key) {
35     size_t class_num = hash(key) % h.size;
36     list<
37         struct {
38             K key;
39             V val;
40         }
41     > class = hash.p[class_num];
42
43     while (class) {
44         if (class->key == key)
45             return class->value;
46         class = class->next;
47     }
48     exit(1); // asking for non existent key
49 }

```

### 3.5.1 Hashing

As we've seen, a crucial part of an hashmap is hashing, but how to devise a good hashing function ?

First of all, it depends on the patterns your input will have.

You don't really want your hashing to be predictable, you want it to seem random (while being deterministic) as it lowers the risk of intentional collision (people will have a harder time retro engineering conflicting keys).

And we want it to seem random under modulo of some number.

By the way it's a good idea to have this "modulo number" be prime, and for similar reasons of unpredictability, abusing prime numbers for hashing (but a different one).

For example, some hashing ideas:

```

1 size_t hash_string(char *s) {
2     size_t p = 1000000009;
3     size_t res = 0;
4     size_t r = 1;
5     while (*s) {
6         res += r * (*s);
7         s++;
8         r *= p;
9     }
10    return res;
11 }
12
13 size_t hash_size_t(size_t n) {
14     size_t p = 1000000009;
15     return pow(n, p); // you have to defined your pow, making it O(
16                        log(n)) plz
17 }
```

### 3.5.2 Transposition tables

Sometimes, you want to speedup computation of some redundant tasks, but it'll be too big to memoize it.

For that, you can use a transposition table, it is used a lot to make game bot using some sort of minimax search.

The idea, is that you'll store the last element of the same hash, and erase the previous ones.

This is less costly memory-wise, and the (time) proximity makes it good enough as a *transposition table* for lots of games, thus its name.

The types are as so:

```

1 typedef fat<option<struct {K key; V val;}>> transposition_table<K,
2     V>;
3 option<V> get(transposition_table_t, K);
4 void set(transposition_table_t, K, V);
```

### 3.5.3 Perfect hashmaps

Now, there are cases where your hash function is *very* good, so good that you won't ever have a collision.

In that case, you don't need the lists shenanigans, neither do you need to store the keys.

So the type become:

```

1  typedef fat<V> perfect_hashmap<K, V>;
2
3  V get(perfect_hashmap<K, V> h, K key) {
4      return h.p[hash(key) % h.size];
5  }
6
7  void set(perfect_hashmap<K, V> h, K key, V val) {
8      h.p[hash(key) % h.size] = val;
9  }

```

### 3.5.4 Sets

Sets are often defined as `hashmap`.  
Basically:

```
1 typedef hashmap<K, void> set<K>;
```

You can make some optimisations because of `void` having only one value, and change the `get` to return a `bool`.  
And if your set is a perfect hashmap, you can use an `fat<bool>`.

### 3.6 Bitsets

Hashmaps are quite slow, so for some use cases, when the keys are *small* `size_t`, you can use an integer as a set.

The main use case is being able to do very efficient intersection and merging of sets.

The idea is to abuse the representation of an integer, we'll say that `size_t` are 64-bits.

We'll look at the bits of a `size_t`, and say that 0 means that the “element” isn't included and 1 if it is.

The element being a number from 0 to 63.

[illegible]

Note that you can only store `sizeof(size_t) * 8` elements, if you need more you can make some structs.

It is useful to translate your definition domain, for example, doing `c - 'a'` when dealing with latin letters strings.

The basic operations are as follow:

```

1 // add a new element e to s
2 s |= ((size_t) 1) << e;
3
4 // check if e is in s
5 0 != s & ((size_t) 1) << e;
6
7 // remove an element e to s
8 s &= ~(((size_t) 1) << e);
9
10 // remove an element e to s if it's here, add it otherwise

```

```

11 s ^= ((size_t) 1) << e;
12
13 // union of s and s2
14 s | s2;
15
16 // intersection of s and s2
17 s & s2;

```

### 3.7 Run length encodings

There are times where the same element is repeated multiple times, it is sometimes useful to know how time the number is repeated.

So instead of storing  $[v_0, \dots, v_n]$ , you store  $[(v_0, \text{how many time } v_0, \dots)]$

It's a bit clearer by seeing the algorithm:

```

1 typedef fat<struct {T val; size_t repeated;} rle<T>;
2
3 rle<T> init(fat<T> v) {
4     if (v.size == 0) {
5         rle = rle<T>.init(0);
6         return rle;
7     }
8     T last = v.p[0];
9     size_t num_diff = 1;
10    for (size_t i = 1; i < v.size; ++i) {
11        num_diff += last != v.p[i];
12        last = v.p[i];
13    }
14
15    rle<T> r = rle<T>.init(num_diff);
16    last = v.p[0];
17    size_t num = 1;
18    int j = 0;
19
20    for (size_t i = 1; i < v.size; ++i) {
21        if (last == v.p[i]) {
22            num_diff++;
23        } else {
24            rle.p[j].val = last;
25            rle.p[j].repeated = num;
26            num = 1;
27            j++;
28        }
29        last = v.p[i];
30    }
31 }

```

It is also sometimes useful to sort the vector first.

## 4 Graphs

I'll assume quite a lot of familiarity with graphs, if you don't know a term, wikipedia them.

We'll denote the nodes from 0 to  $n - 1$  (where  $n$  is the number of nodes), this choice leads to much simpler and efficient implementations.

## 4.1 General graphs

A graph to which you don't know many structural features, thus to which you don't optimize its structure much.

I'll assume that between two nodes, there can be only one edge, it is the case for most useful graphs, and the implementations are much better.

There are a lot of ways to represent them, but three of them are the most usefuls (never got a case where incidence matrix was useful).

### 4.1.1 Adjacency lists

You represent your graph with a length  $n$  fat pointer  $p$  of fat pointers of int (or vector of int if you want to be able to change your graph).

Basically, the node  $i$  will direct into all nodes of  $p.p[i]$ .

If your graph is weighted, you can have a tuples of nodes and weights instead of just nodes.

This is especially good for sparse graphs and for common graphs traversals (dfs and bfs).

For directed graphs, you either specify only out neighbors, or you use two different fat pointers to for in and out.

In "C", an example to type it would be:

```
1 typedef fat<struct {
2     fat<struct {
3         size_t node;
4         double weight;
5     }> in, out;
6 }> adjacency_list_t;
```

### 4.1.2 Adjacency matrices

We'll consider that non weighted graphs are weighted graphs with all edge weights put to the value 1.

Basically, we construct an  $n \times n$  matrix, and at index  $i, j$ , you put 0 if there are no edge from  $i$  to  $j$  and you put the weight of this edge otherwise.

It is possible that in some cases, 0 is a valid weight, so use  $-1$ , but sometimes it is a valid weight, in that case, use an `option<T>`.

So in "C", its type would be:

```
1 typedef fat<option<T>> adjacency_matrix_t;
```

It is useable only when  $n$  is relatively small (as it is  $O(n^2)$  memory wise).

The best trick to know about it is that if  $M$  is the adjacency matrix of  $G$ , a non weighted graph.

Then  $M_{i,j}^k$  is the number of path of length  $k$  going from  $i$  to  $j$ .

And the operation  $M^k$  is an  $O(n^2 \log k)$  operation, so it can be very efficient for really big  $k$ .

### 4.1.3 Actual graph definition

No, not incidence matrix, I never had to use it.

We know that a graph is a pair  $(V, E)$  where  $V$  are the vertices and  $E$  the edges. Because of our numbering scheme, we don't have to store  $V$  and only need  $|V|$ . And for the edges, we can store them as tuple (from, to, weight), where from and to are vertices.

Thus a graph is defined as so

```
1 struct graph_s {
2     size_t n;
3     fat<struct {
4         size_t from, to;
5         double weight;
6     }> edges;
7 };
```

It can be convenient to only keep .edges.

This is to use when you want to iterate through edges by some order, or when  $|E|$  is negligible compared to  $|V|$ .

## 4.2 Rooted trees

If the tree isn't rooted in an obvious way, not specialising it and using an adjacency list is quite common.

### 4.2.1 General rooted trees

One of the most general way to implement them is as so:

```
1 struct rooted_tree_s {
2     struct rooted_tree_s *parent; // itself or NULL if it is the root
3     , depends on what you want
4     fat<struct rooted_tree_s> childs;
5 };
```

It's a bit cumbersome to construct, but it allows for decently easy tree related traversals.

For binary rooted tree with a left, right distinction, it's as so:

```
1 struct binary_rooted_tree_s {
2     struct binary_rooted_tree_s *parent; // itself or NULL if it is
3     the root, depends on what you want
4     struct binary_rooted_tree_s *left, *right;
5 };
```

Sometimes you don't specify the parent.



### 4.2.2 Parent oriented trees

Not an actual name.

We'll use the convention that the parent of the root is itself.

What you may think about is doing

```
1 typedef fat<struct parent_oriented_tree_s {
2     struct parent_oriented_tree_s *parent;
3 }> parent_oriented_tree_t;
```

But we'll once again use the the vertex numerotation is  $0..n - 1$ , we don't need to point to a struct, we can give out the vertice number directly.

And that for all vertices from  $0..n - 1$ .

So a fat pointer does the job perfectly

```
1 typedef fat<size_t> parent_oriented_tree_t;
```

Where the  $i$ th element is the parent of  $i$ .

It's very good for accessing the parent, doesn't use much space, and is quite efficient for what it's designed for, but horrible for a lot of common operations, so definitely not a general purpose data structure.

### 4.2.3 Complete binary trees

Due to its highly regular structure, there are a useful trick to represent complete binary tree which doesn't require to store for each nodes its parents, left and right side.

Instead we can store all of them in a fat pointer, and do some indexing tricks.

We'll put the root at the 0-th index, and we'll put the left node of the  $i$ -th index at index  $2i + 1$  and the right at index  $2i + 2$ .

To see why it'd work (why it uses the indexes from 0 to  $n - 1$ , without overlap), it's easier to reason in base 1.

So our rule become root at index 1 and left and right at  $2i$  and  $2i+1$  respectively.

The no overlap is obvious because composing some  $\times 2$  and some  $\times 2 + 1$  and applying it to 1 is akin to using Hörner polynomial decomposition, it creates a unique representation of an integer strictly greater than 0.

And we see that the left part have smaller index than the right one, thus the indexes are 1 to  $n$ .

Now to translate back to our 0 based index, we put everything one case left.

For practical reasons, we uses the indices for the methods.

So an implementation of such would be:

```
1 typedef fat<T> perfect<T>;
2
3 T get(perfect<T> p, size_t i) {
4     return p.p[i];
5 }
6
7 size_t parent(size_t i) {
8     return (i-1) / 2; // note that "-1/2" is 0 in C
9 }
10
```

```

11 size_t left(size_t i) {
12     return 2*i + 1;
13 }
14
15 size_t right(size_t i) {
16     return 2*i+2;
17 }

```

Note that it can be used for non perfect ordered binary trees as well, by using an `option<T>` instead of `T`, though the memory usage is polynomial with the tree height so it's not advisable for very badly balanced trees.

It can also be used for ternary, even  $n$ -ary perfect trees like so:

```

1 typedef fat<T> perfect<T, n>; // n is the number of possible
   children
2
3 T get(perfect<T, n> p, size_t i) {
4     return p.p[i];
5 }
6
7 size_t parent(size_t i) {
8     return (i-1) / n; // same n as before
9 }
10
11 // the num-th child of i, with 0 based indexing
12 size_t nth(size_t i, size_t num) {
13     return n*i + num + 1;
14 }

```

#### 4.2.4 Heaps

A heap is a complete binary tree where each element has a value less than its children's.

So, of course, we'll use `perfect<T>`, but with a `vector<T>` instead of a `fat<T>` to be able to add and remove elements.

The main goal is being able to find the minimum, adding an element and removing an element efficiently.

Finding the min is very easy:

```

1 void min(heap<T> h) {
2     return h.p[0];
3 }

```

Now, removing and adding elements is a bit harder, there are different ways to implement those but I'll introduce one which in my opinion, is easy to understand and to extend, the "swift up", "swift down" approach.

The "swift up" and when you have an almost heap, it is an heap for indexes from 0 to  $i - 1$  included, and you want to make it an heap from 0 to  $i$  by bubbleing up the element.

They'll both return the new index of the previously  $i$ -th element.

```

1 // n <= h.size
2 size_t swift_up(heap<T> h, size_t i, size_t n) {
3     while (1) {

```

```

4     size_t p = parent(i);
5     size_t l = left(p);
6     size_t r = right(p);
7
8     size_t smallest = p;
9
10    if (l < n && get(h, l) < get(h, p))
11        smallest = l;
12
13    if (r < n && get(h, r) < get(h, p))
14        smallest = r;
15
16    if (p == smallest)
17        return i;
18
19    swap(&h.p[p], h.p[smallest]);
20
21    i = p;
22 }
23 }

```

And swift down is when you have an almost heap, an heap from index  $i + 1$  to  $n - 1$  and you want to make it an heap from  $i$  to  $n - 1$ .

```

1 // n <= h.size
2 size_t swift_down(heap<T> h, size_t i, size_t n) {
3     while (1) {
4         size_t l = left(i);
5         size_t r = right(i);
6
7         size_t smallest = i;
8
9         if (l < n && get(h, l) < get(h, p))
10             smallest = l;
11
12         if (r < n && get(h, r) < get(h, p))
13             smallest = r;
14
15         if (p == smallest)
16             return i;
17
18         swap(&h.p[p], h.p[smallest]);
19
20         i = smallest;
21     }
22 }

```

It is very similar and that mean you can easily copy-paste your swift\_up and modify some value very quickly.

To add an element, you can put it last, and swift it up.

```

1 void add(heap<T> h, T el) {
2     size_t n = h.size;
3     add_vector(&h, el);
4     swift_up(h, n, h.size);
5 }

```

To delete an element, you put the last element of the vector at its place, pop the vector, then heap the whole again.

```

1 void pop(heap<T> h, size_t i) {
2     h.p[i] = h.p[h.size - 1];
3     pop(h);
4     swift_down(h, swift_up(h, i, h.size), h.size);
5 }

```

Note that it can be used to create an easy  $O(n \log n)$  in place sorting algorithm, heap sort:

```

1 void heapsort(heap<T> p) {
2     heap<T> h = p;
3     size_t n = p.size;
4
5     // make it an actual heap
6     for (int i = 0; i < n; ++i)
7         swift_up(h, i, h.size);
8
9     // sort it in reverse
10    for (int i = n-1; i; --i) {
11        swap(&h.p[i], &h.p[0]);
12
13        swift_down(h, 0, i);
14    }
15
16    // reverse it
17    for (int i = 0; i < n/2; ++i)
18        swap(&h.p[i], &h.p[n-i-1]);
19 }

```

There are actually a lot of different kind of heaps, the other “important” one is the Fibonacci’s heap (used for better time complexity for Prim’s and Dijkstra’s algorithms), but its implementation is quite complicated.

## 5 Priority Queues

Priority queues are data structures which represent a queue where the first element out is the smallest one.

An easy way to implement it is by using an heap, different kind of heap makes different cases where the queue is good at.

## 6 Disjoint set unions

Imagine, you want to making a civilisation-like game.

People are in guilds (of just themselves if they don’t have friends), only one.

And they can decide to merge guilds.

And quite often, players will ask if they are in the same guild as someone else because otherwise they’ll start a war.

This is the kind of problems this datastructure will solve, the two main operations are `same(dsu, person1, person2)` and `merge(dsu, person1, person2)` (with `merge` returning if it indeed had to do something to merge).

I'll present multiple partial solutions as it's better for learning in my opinion.

Let  $n$  be the number of persons.

One of the easiest way to modilise it is with a bipartite  $n + n$  graph, where on the left side there are the persons and on the right the guilds, and a connection if that person is in that guild, we'll say that by default, persons are matched with guild of the same "level".

Making for  $O(1)$  same and  $O(n)$  merging.

```

1 typedef fat<size_t> dsu;
2
3 dsu init_dsu(size_t n) {
4     dsu d = init_fat<size_t>(n);
5     while (--n)
6         d.p[n] = n;
7     return d;
8 }
9
10 bool same(dsu d, size_t a, size_t b) {
11     return d.p[a] == d.p[b];
12 }
13
14 bool merge(dsu d, size_t a, size_t b) {
15     if same(d, a, b)
16         return false;
17     size_t guild_of_b = d.p[b];
18     for (int i = 0; i < d.size; ++i)
19         if (d.p[i] == guild_of_b)
20             d.p[i] = d.p[a];
21     return true;
22 }

```

You may think that iterating through everything is wasteful and you'd want to iterate only the one which are on the same guild as b.

It is indeed better in average, but it has the same mean and worstcase complexity, so we won't go that route.

Now, there are an obvious bijection between the guilds and the persons, so let's consider that a connect mean "has the same guild as this person", we want a parent oriented tree/graph !

We'll create a function `find` which find the root of the tree, the "representent" of the guild.

```

1 size_t find(dsu d, size_t i) {
2     while (i != d.p[i])
3         i = d.p[i];
4     return i;
5 }
6
7 bool same(dsu d, size_t a, size_t b) {
8     return find(d, a) == find(d, b);
9 }
10

```

```

11 bool merge(dsu d, size_t a, size_t b) {
12     size_t fa = find(d, a);
13     size_t fb = find(d, b);
14     if (fa == fb)
15         return false;
16     d.p[fb] = fa;
17     return true;
18 }

```

And if we look at the time complexity, it's now  $O(n)$  for all operations... that's not better, but we have some tricks up our sleeves.

We can remark that the operation `find` is  $O(n)$  because we allow our tree to be very deep, unnecessary deep.

In a perfect world, you'd want it to be a complete tree or at least complete until the leaf level, that make it  $O(\log n)$  and all operations will become  $O(\log n)$  as a consequence.

To have this kind of results, there are two common heuristics, the “weight” and “depth” ones.

The weight one is storing how many nodes are attached to it, and depth one is storing the height of the tree of nodes attached to it.

They give out the same time complexity ( $O(\log n)$  for `find` and thus for all operations), but I prefer the “weight” one because it is very slightly faster and will play nicely for other operations, though the “depth” one is more often used.

```

1  struct dsu {
2      size_t n;
3      size_t *succ; // array of size n
4      size_t *w; // array of size n
5  }
6
7  struct dsu init_dsu(size_t n) {
8      struct dsu d = {.n = n, .succ = malloc(n*sizeof(size_t)), .w =
9          malloc(n*sizeof(size_t))};
10     for (int i = 0; i < n; ++i) {
11         dsu.w[i] = 1;
12         dsu.succ[i] = i;
13     }
14 }
15
16 size_t find(struct dsu d, size_t i) {
17     while (i != d.succ[i])
18         i = d.succ[i];
19     return i;
20 }
21
22 bool same(struct dsu d, size_t a, size_t b) {
23     return find(d, a) == find(d, b);
24 }
25
26 bool merge(struct dsu d, size_t a, size_t b) {
27     size_t fa = find(d, a);
28     size_t fb = find(d, b);
29
30     if (fa == fb)
31         return false;

```

```

31
32     if (d.w[fb] > d.w[fa])
33         swap(&fa, &fb);
34
35     d.succ[fb] = fa;
36     d.w[fa] += d.w[fb];
37     return true;
38 }

```

Now, suppose you call `find(d, 3)` a lot of time without merging in between, it feels a bit wasteful to do an  $O(\log n)$  operation each time when you could memoize it.

And that is the idea behind the “path compression” optimisation, when you are finding the representent of a node, you put it as its successor directly, making it much more efficient in amortized time.

```

1 size_t _find(struct dsu d, size_t i) {
2     while (i != d.succ[i])
3         i = d.succ[i];
4     return i;
5 }
6
7 size_t find(struct dsu d, size_t i) {
8     size_t repr = _find(d, i);
9     while (i != repr) {
10         size_t next = d.succ[i];
11         d.succ[i] = repr;
12         i = next;
13     }
14     return repr;
15 }

```

Note on how we don’t care about the weights of nodes, also it have a quite nice non tail-call recursive way to write it, but it’s not tail call recursive so I prefer this way as it’s constant in memory usage.

All operations are now  $O(\alpha(n))$  amortized, where  $\alpha$  is the inverse function of the Ackermann’s, so it grows *very* slowly, and for all reasonable inputs it’s less than 4.

Note though that there are not only just one easy to backtrack operation done when doing a `find`, so to use it for backtracking-heavy algorithms (like `negamax`), you may decide to not implement the path compression.

Now, let’s add a functionality to our game, a player can decide to quit its guild and join its own.

An idea is that if nobody is connected to them, they can just take themselves as parent, otherwise, if they are their own guild representent, they need to take its ownership to someone else, and use the path compression to put everyone else to the other representent to be alone.

Though this require to have an accurate representation of weights for non root nodes.

```

1 void flatten(struct dsu d) {
2     for (int i = 0; i < d.size; ++i)
3         find(d, i);

```

```

4 }
5
6 size_t find(struct dsu d, size_t i) {
7     size_t repr = _find(d, i);
8     size_t prev_weight = d.w[i];
9     while (i != repr) {
10         size_t next = d.succ[i];
11         d.succ[i] = repr;
12         i = next;
13         if (i != repr) {
14             size_t tmp = d.w[i];
15             d.w[i] -= prev_weight;
16             prev_weight = tmp;
17         }
18     }
19     return repr;
20 }
21
22 // O(n)
23 bool disconnect(struct dsu d, size_t i) {
24     if (d.w[i] == 1) {
25         if (d.succ[i] == i)
26             return false;
27         --d.w[d.succ[i]];
28         d.succ[i] = i;
29         return;
30     }
31
32     flatten(d);
33
34     if (d.w[i] == 1) {
35         if (d.succ[i] == i)
36             return false;
37         --d.w[d.succ[i]];
38         d.succ[i] = i;
39         return;
40     }
41
42     // i is the representent of a guild with >= 2 members
43     for (int k = 0; k < d.n; ++k) {
44         if (i == k)
45             continue;
46         if (d.succ[k] == i) {
47             d.succ[i] = k;
48             d.succ[k] = k;
49             d.w[k] = d.w[i]--;
50             break;
51         }
52     }
53     flatten(d);
54     --d.w[d.succ[i]];
55     d.succ[i] = i;
56 }

```



## 7 Range queries

Imagine you want a binary operators  $a \top b$  and an array  $v$ , and you want to be able to quickly know  $\bigtop_{k=i}^j v_k$  for some  $i \leq j$  not given (those will be the queries). You could store all the queries ahead of time, having a time complexity of  $O(n^2)$  or you can use some of the followings data structures

### 7.1 Prefix sums

If  $\top$  is commutative, divisible and that the values wouldn't change, then we can use a prefix "sum" (the term is as so because you usually use it for addition):

```
1 typedef fat<T> prefix<T>;
2
3 prefix<T> init(fat<T> f) {
4     prefix<T> = copy(f);
5     if (p.size <= 1)
6         return p;
7
8     T s = p.p[0];
9
10    for (int i = 1; i < p.size; ++i)
11        p.p[i] = s = op(s, p.p[i]);
12
13    return p;
14 }
15
16 T query(prefix<T> p, size_t i, size_t j) {
17     if (i > j)
18         exit(1); // plz handle errors correctly, can't return 0 because
19                 // op isn't unitful
20     if (!i)
21         return p.p[j];
22     return div(p.p[j], p.p[i-1]);
23 }
```

With  $O(1)$  query,  $O(n)$  initialisation, and  $O(n)$  updating (just initialise it again).

### 7.2 Segment trees

We'll now suppose that  $\top$  is associative and unitful (if it isn't, we can make it so by using an `option<T>` and consider the `none` value as being the unit), so we'll consider  $\top$  to be unitful.

What you'd do is having a complete binary tree with last layer be your vector (and unit element to fill it up), and the value of a non-leaf node  $n$  is  $\text{node\_left}(n) \top \text{node\_right}(n)$ .

I'll consider that the unit element of  $\top$  is called `unit`.

In code it becomes:

```
1 typedef struct {
2     size_t depth;
3     T *p; // of size 2depth+1 - 1
4 }
```

```

4 } segment_tree<T>;
5
6 segment_tree<T> init(fat<T> f) {
7     size_t n = f.size;
8     size_t d = 0;
9     size_t size = 1;
10    size_t pow2 = 2;
11
12    while (size < n) {
13        d++;
14        size += pow2;
15        if (size < n)
16            pow2 *= 2;
17    }
18
19    segment_tree<T> st;
20    st.depth = depth;
21    st.p = malloc(size * sizeof(T));
22    for (int i = size - pow; i < size - pow + n; ++i)
23        st.p[i] = f.p[i - size + pow];
24    for (int i = size - pow + n; i < size; ++i)
25        st.p[i] = unit;
26    for (int i = size - pow - 1; i >= 0; --i)
27        st.p[i] = op(st.p[left(i)], st.p[right(i)]);
28 }
29
30 T _query(T *p, size_t i, size_t diameter, size_t l, size_t r) {
31     if (r - l + 1 == diam)
32         return p[i];
33     if (diameter == 1)
34         return unit;
35
36     size_t mid = diameter / 2;
37     if (r < mid)
38         return _query(p, left(i), mid, l, r);
39     if (l >= mid)
40         return _query(p, right(i), mid, l - mid, r - mid);
41     return op(_query(p, left(i), mid, l, mid), _query(p, right(i),
42         mid, 0, r - mid));
43 }
44
45 T query(segment_tree<T> st, size_t l, size_t r) {
46     size_t diam = 1ULL << st.depth;
47     if (l < 0 || r >= diam)
48         exit(1);
49     return _query(st.p, 0, 1ULL << st.depth, size_t l, size_t r);
50 }
51
52 void update(segment_tree<T> st, size_t i, T val) {
53     size_t diam = 1ULL << st.depth;
54     size_t size = (1ULL << (st.depth + 1)) - 1;
55     i = size - diam + i;
56     st.p[i] = val;
57     while (i != 0) {
58         i = parent(i);
59         st.p[i] = op(st.p[left(i)], st.p[right(i)]);
60     }
61 }

```

60 }

With  $O(\log n)$  query,  $O(n)$  initialisation,  $O(\log n)$  updating.

### 7.3 Fenwick/BIT trees

We will now suppose that  $\top$  forms a group.

There are actually two common types of Fenwick's (or Binary Interval Tree) trees. (yes, I repeat myself with the "tree tree"), the 1-based index and the 0-based index.

The 1-based index is more common, but I find that letting the index 0 unused is quite sad, so we'll speak about the 0-based index version.

The idea is that you'll use a `fat<T>` of size  $n$ , and the index  $i$  is sum of values at indices from  $i \& (i + 1)$  (changing all leading 1 in the based 2 representation into 0s) to  $i$  included.

And thus, when updating a value of index  $i$ , you have to modify all indices  $j$  with the property:  $j \& (j + 1) \leq i \leq j$ .

To find, you can start with  $i$  and flip the last 0 into an 1 for each iteration, which is the operation  $j | (j + 1)$  and you stop once  $j$  will out of bound.

With those information actually writing the data structure becomes doable:

```
1 typedef fat<T> fenwick<T>;
2
3 T op_to(fenwick<T> fw, size_t to) {
4     T res = unit;
5     while (to >= 0) {
6         res = op(fw.p[to], res);
7         to = (to & (to + 1)) - 1;
8     }
9     return res;
10 }
11
12 T query(fenwick<T> fw, size_t from, size_t to) {
13     return div(op_to(fw, to), op_to(fw, from - 1));
14 }
15
16 void op_update(fenwick<T> fw, size_t i, T val) {
17     while (i < fw.size) {
18         fw.p[i] = op(fw.p[i], val);
19         i = i | (i+1);
20     }
21 }
22
23 void update(fenwick<T> fw, size_t i, T val) {
24     T cur = query(fw, i-1, i);
25     op_update(fw, i, div(val, cur));
26 }
27
28 fenwick<T> init(fat<T> f) {
29     fenwick<T> fw = fat<T>.init(f.size);
30     for (size_t i = 0; i < f.size; ++i)
31         fw[i] = unit;
32     for (size_t i = 0; i < f.size; ++i)
33         op_update(fw, i, f[i]);
```

```

34     return fw;
35 }

```

Though this initialisation is  $O(n \log n)$ , not the best if  $\mathbb{T}$  is also abelian, as you can do it in  $O(n)$  as so:

```

1 fenwick<T> init(fat<T> f) {
2     fenwick<T> fw = fat<T>.init(f.size);
3     for (size_t i = 0; i < f.size; ++i)
4         fw[i] = unit;
5     for (size_t i = 0; i < f.size; ++i) {
6         fw[i] = op(fw[i], f[i]);
7         size_t next = i | (i+1);
8         if (next < f.size)
9             fw[next] = op(fw[next], fw[i]);
10    }
11    return fw;
12 }

```

Giving similar time complexity of a segment tree but with better constants.

## 8 Memory management

### 8.1 Safe guards

I'll call by "safe guards" data structures which doesn't *do* anything new, but prevent doing *bad things*, such as freeing memory still in use, or having race conditions.

The second kind are built-in the CPU, but it still can be useful to know how it'd work.

#### 8.1.1 Rcs

Sometimes, it's very hard to track the lifetimes of your objects, especially when you're starting to pass around views and/or allocating in a multi-threaded environment.

Wouldn't it be so simple to just call free and if you don't have any other view of the memory, it'll free and otherwise do "nothing" ?

Well, that's what reference counting (Rc) is for.

The idea is that you count how many references to that memory you have, and when you free, you remove a reference, and if that reference count goes to 0, you actually free your memory.

```

1 typedef struct {T v; size_t count} *Rc<T>;
2
3 Rc<T> init_rc(T val) {
4     Rc r = malloc(sizeof(*(Rc)0)); // some compilers may really not
5     like it
6     r->v = val;
7     r->count = 1;
8     return r;
9 }

```

```

10 // to call each time you're creating a new reference
11 Rc<T> copy(Rc<T> r) {
12     r->count++;
13     return r;
14 }
15
16 // to call each time a reference is destroyed
17 Rc<T> free_rc(Rc<T> r) {
18     r->count--;
19     if (r->count == 0) {
20         free(r->val);
21         free(r);
22     }
23 }

```

Note that Rc is by itself a pointer instead of a struct to be able to modify the count accross references.

Be careful to not do circular referencing of Rcs, you won't be able to free them out.

But it's impossible to do so in languages without mutation, so Rc-based garbage collection is a viable strat for those.

### 8.1.2 Mutexes

There are a lot of data structures for thread safety, though the most common one is the mutex.

The idea is that along your T, you put a variable to lock the reading/writing when you need to access it yourself.

You could use a bool for the locking variable, but it itself is prone to race condition, so you have to look what your language's thread safe locking type is.

In C it is pthread\_mutex\_t.

In code you'd do:

```

1 typedef struct {T val; pthread_mutex_t mutex;} mutex<T>;
2
3 mutex<T> init(T val) {
4     mutex<T> m;
5     m.val = val;
6     pthread_mutex_create(&m.mutex, NULL);
7     return m;
8 }
9
10 void free_mutex(mutex<T> m) {
11     pthread_mutex_destroy(&m.mutex);
12 }
13
14 void lock(mutex<T> m) {pthread_mutex_lock(m.mutex);}
15 void unlock(mutex<T> m) {pthread_mutex_unlock(m.mutex);}

```

## 8.2 Allocators

Have you ever wondered where your memory comes from and what were the choices ?

Well, here we go.

Because a lot of our allocators will take others as an argument, I'll define a small allocator interface

```
1 typedef struct allocator_s {
2     void (*free_memory)(struct allocator_s *self, void *addr, size_t
      size); // yeah, no check if that works
3     void *(*alloc_memory)(struct allocator_s *self, size_t size);
4 } allocator;
```

And we'll demand that the first two fields of any allocators we'll use have those two fields in this order at first before any other potential field.

### 8.2.1 Page allocators

The user side and kernel side is a bit different as users have the already made allocator from the kernel.

We'll call this allocator a `page_allocator`.

It's not really a data structure per se, it's merely a general purpose allocator (kernel code) backed by a fix buffer allocator (the computer memory) (those term will be explained later).

Though I decided to explain it first as it is how you can create most of your own allocators.

Your OS gives your two several (virtual) memory spaces (so don't forget to thank it), there are two which will interest us, the stack and the heap.

To get stack space, you can use `strpbrk`, it is useful and you can use to allocate, but we'll look at the heap implementations (as they have less constraints in term of scoping).

To get a valid pointer to some heap memory (without `malloc`), the OS way (in the C stdlib) is the `mmap` method and the `munmap` to free.

The type signature:

```
1 void *mmap(void addr[.length], size_t length, int prot, int flags,
      int fd, off_t offset);
2 int munmap(void addr[.length], size_t length);
```

You can use `mmap` to map file memory, the `fd` argument, but here we'll use `NULL` to indicate that we're using a new heap.

For speed reasons, it'll usually give you out a memory address multiple of the OS page size, and it'll give you multiple of your page size for acceptable read/write location.

In most systems, it is 4096, so we'll make it extra clear in our implementation.

```
1 #define PAGE_SIZE 4096
2
3 typedef allocator page_allocator; // doesn't need any new field
4
5 size_t to_page_size(size_t n) {
6     return PAGE_SIZE*((n/PAGE_SIZE) + (n % PAGE_SIZE != 0));
7 }
8
9 void free_memory_page(allocator *self, void *addr, size_t size) {
```

```

10 size = to_page_size(size);
11 return NULL;
12 munmap(addr, size);
13 }
14
15 void *alloc_memory_page(allocator *self, void *addr, size_t size) {
16     size = to_page_size(size);
17
18     // it's not defined what mmap should do for those values
19     // so I'd just return nonsense
20     if (size <= 0)
21         return mmap(NULL, size, PROT_READ | PROT_WRITE, MAP_ANONYMOUS |
22                     MAP_PRIVATE, -1, 0);
23 }
24
25 page_allocator init_page_allocator(void) {
26     page_allocator page;
27     page.free_memory = &free_memory_page;
28     page.alloc_memory = &alloc_memory_page;
29     return page;
}

```

Read the man pages for the flags and prot of mmap.

### 8.2.2 Fix buffer allocators

A huge problem of page allocators is that they do a system call for each allocation, and also they waste a lot of memory (for example, if you need 3 bytes, it'll give you 4096 instead).

What we can do is to allocate a bigger memory space, and then manage in that space what we'll allocate in there.

So here is the FBA, you have a memory address, you cannot free by the way (sad time) and you keep track of what space you used or you haven't used with an offset, what's left of the offset is used, and what's right isn't.

```

1 typedef struct {
2     // layout requirements
3     void (*free_memory)(struct allocator_s *self, void *addr, size_t
4                         size);
5     void (*alloc_memory)(struct allocator_s *self, size_t size);
6
7     char *buffer;
8     size_t size;
9     size_t head;
10 } FBA;
11
12 void free_memory_FBA(allocator *self, void *addr, size_t size) {
13     return; // yeah, no op
14 }
15
16 void *alloc_memory_FBA(allocator *self, size_t size) {
17     char *current = self->buffer + self->head;
18     self->buffer.head += size;
19     // you can check if self->head > self->size, in that case it'll
20     // buffer overflow
21     return current;
}

```

```

20 }
21
22 FBA init_FBA(allocator *altor, size_t size) {
23     FBA fba;
24     fba.free_memory = &free_memory_FBA;
25     fba.alloc_memory = &alloc_memory_FBA;
26     fba.buffer = altor.alloc_memory(size);
27     fba.size = size;
28     fba.head = 0;
29     return fba;
30 }
31
32 void free_fba(allocator *altor, FBA *fba) {
33     altor.free_memory(fba->buffer);
34 }

```

### 8.2.3 Arena allocators

ARENA = List<FBA> One problem with the previous way is that if we don't have enough memory to allocate our new object, we're just screwed, and that's bad, a way to circumvent that is to allocate a new FBA if we don't have enough space anymore.

And let's put those in a linked list (putting those in a vector wouldn't work because you can't copy memory addresses, only their values, that's one of the cases where using a vector isn't possible while using a linked list is).

In code:

```

1 typedef struct arena_node_s {
2     char *buffer;
3     size_t size;
4     size_t head;
5     arena_node_s *next;
6 } arena_node;
7
8 typedef {
9     // layout requirements
10    void (*free_memory)(struct allocator_s *self, void *addr, size_t
        size);
11    void (*alloc_memory)(struct allocator_s *self, size_t size);
12
13    allocator *altor;
14    arena_node *node;
15 } arena;
16
17 arena_node *init_arena_node(allocator *altor, size_t size) {
18     arena_node *n = (arena_node *) altor->alloc_memory(sizeof(
        arena_node));
19     n->buffer = (char *) altor->alloc_memory(size);
20     n->size = size;
21     n->head = 0;
22     n->next = NULL;
23     return n;
24 }
25
26 void free_memory_arena(void *addr, size_t size) {

```



```

27     return;
28 }
29
30 void *alloc_memory_node_arena(allocator *altor, arena_node *node,
    size_t size) {
31     if (node->head + size <= node->size) {
32         void *rv = node->head+node->buffer;
33         node->head += size;
34         return rv;
35     }
36     if (!node->next)
37         node->next = init_arena_node(max(size, node->size));
38     return alloc_memory_node_arena(altor, node->next, size);
39 }
40
41 void *alloc_memory_arena(allocator *altor, size_t size) {
42     arena *a = (arena *) (void *) altor;
43     return alloc_memory_node_arena(a->altor, &a->node, size);
44 }
45
46 arena *init_arena(allocator *altor, size_t size) {
47     arena a;
48
49     a->free_memory = &free_memory_arena;
50     a->alloc_memory = &alloc_memory_arena;
51
52     a.allocator = altor;
53     a.node = init_arena_node(allocator *altor, size_t size);
54
55     return a;
56 }
57
58 void arena_node_free(allocator *altor, arena_node *node) {
59     if (node->next)
60         arena_node_free(altor, node->next);
61     altor->free_memory(node->buffer);
62     altor->free_memory(node);
63 }
64
65 void arena_free(arena *a) {
66     arena_node_free(a->altor, a->node);
67 }

```

A good thing about it is that you can allocate all the memory you want, and then deallocate everything at once, without having to care the slightest about lifetimes.

Though like before, we cannot really free stuffs.

It's also quite common to create a function `reset` which won't deallocate your arena, but will invalidate all its memory, so that you can use it back again without having to allocate again.

#### 8.2.4 General purpose allocators

A lot of allocators can be used as “general purpose”.

The most common ones are the “slabs” and “buddy” allocators.

They are a bit too complexe for me to explain them clearly and make a good implementation of them (implementations are multiple hundreds of lines long), so I'll let you search online to learn how to do so.

## 9 After words

Thanks for reading, I hope you understand stuffs, and now feel much more competant for writing your own, personal data structures when you'll have a need to.

My goal in this guide was to show a bit the thought process on how to come up with those, though for brevety reasons, I didn't go in depth into how to discover all of them.

If you liked it, try reading others of my guides/articles/blogs/papers on my site:  
<https://juliapbn.pages.dev>