

Implementation of algorithms and data structures

8. seminar

Jirka Fink

<https://ktiml.mff.cuni.cz/~fink/>

Department of Theoretical Computer Science and Mathematical Logic
Faculty of Mathematics and Physics
Charles University in Prague

Summer semestr 2021/22

Last change 13. dubna 2022

Licence: Creative Commons BY-NC-SA 4.0

Termination

- Last element points to NULL/nullptr/None (Null-terminated)
- Last element points to the first one (Circular)

Termination

- Last element points to NULL/nullptr/None (Null-terminated)
- Last element points to the first one (Circular)

How to store next/previous pointers

- Pointers are stored inside object (Intrusive)

```
1 class MyClass {  
2     string name; # Variables we need in the class  
3     node *next, *prev;  
4 };
```

Termination

- Last element points to NULL/nullptr/None (Null-terminated)
- Last element points to the first one (Circular)

How to store next/previous pointers

- Pointers are stored inside object (Intrusive)

```
1 class MyClass {  
2     string name; # Variables we need in the class  
3     node *next, *prev;  
4 };
```

- Pointers are stored outside object (Non-intrusive)

```
1 class MyClass {  
2     string name;  
3 };  
4 list<MyClass> my_list;
```

Termination

- Last element points to NULL/nullptr/None (Null-terminated)
- Last element points to the first one (Circular)

How to store next/previous pointers

- Pointers are stored inside object (Intrusive)

```
1 class MyClass {  
2     string name; # Variables we need in the class  
3     node *next, *prev;  
4 };
```

- Pointers are stored outside object (Non-intrusive)

```
1 class MyClass {  
2     string name;  
3 };  
4 list<MyClass> my_list;
```

Goal

We can easily add pointers into a class, but we do not want to repetitively write functions working with pointers.

How to design a general API?

Linked list: NULL-terminated lists

```
1 struct node {
2     node *next, *prev;
3 };
4 struct list {
5     node *first, *last;
6 };
7 void list_add_tail(list *l, node *n) {
8     n->next = NULL;
9     n->prev = l->last;
10    if(l->last)
11        l->last->next = n;
12    else
13        l->first = n;
14    l->last = n;
15 }
16 void list_remove(list *l, node *n) {
17     if(n->prev)
18         n->prev->next = n->next;
19     else
20         l->first = n->next;
21     if(n->next)
22         n->next->prev = n->prev;
23     else
24         l->last = n->prev;
25 }
```

```
1 struct node {
2     node *next, *prev;
3 };
4 struct list {
5     node head;
6 };
7
8 void list_init(list *l) {
9     l->head.next = l->head.prev = &l->head;
10 }
11
12 void list_insert_after(node *what, node *after) {
13     node *before = after->next;
14     what->next = before;
15     what->prev = after;
16     before->prev = what;
17     after->next = what;
18 }
19
20 void list_remove(node *n) {
21     node *before = n->prev;
22     node *after = n->next;
23     before->next = after;
24     after->prev = before;
25 }
```

```
1 template<class T>
2 class list {
3     struct node {
4         T data;
5         node *next, *prev;
6     };
7     node *first, *last;
8
9 public:
10    using iterator = node*;
11    iterator insert(iterator pos, const T& value);
12    iterator erase(iterator pos);
13    // ...
14 };
15
16 std::list<std::string> fruits {"orange", "apple", "raspberry"};
17
18 // To store one instance in multiple lists, create lists of pointers
19 std::list<MyClass*> list1, list2;
```


Intrusive list without templates

```
1 struct node { node *next, *prev; };
2 struct list { node head; };
3
4 node* list_get_first(list *l) {
5     return l->head.next;
6 }
7
8 struct company {
9     char *name;
10    list employees, consumers;
11 };
12
13 struct employee {
14     node n; // Assume that this node can be the first member variable
15     char *name;
16 }
17
18 void company_add_employee(company *c, employee *e) {
19     list_add_tail(&c->employees, &e->n);
20 }
21
22 employee* company_get_first_employee(company *c) {
23     return (employee*) list_get_first(&c->employees);
24 }
```

Intrusive list without templates

```
1 #define offsetof(TYPE, MEMBER) ((size_t)&((TYPE *)0)->MEMBER)
2
3 #define container_of(PTR, TYPE, MEMBER) \
4     ((TYPE *)((char *)PTR - offsetof(TYPE, MEMBER)))
5
6 struct company {
7     char *name;
8     list employees, consumers;
9 };
10
11 struct consumer {
12     char *name;
13     node n; // Assume that this node cannot be the first member variable
14 };
15
16 consumer* consumer_from_node(node *node_in_consumer) {
17     return container_of(node_in_consumer, consumer, n);
18 }
19
20 consumer* company_get_first_consumer(company *c) {
21     return consumer_from_node(list_get_first(&c->consumers));
22 }
```

Intrusive NULL-terminated list without memory magic

```
1 template <class T>
2 struct node {
3     T *next, *prev;
4 };
5
6 template <class T, node<T> T::*M> // Member pointer
7 class list {
8     T *first = nullptr, *last = nullptr;
9
10 public:
11     T* begin() {
12         return first;
13     }
14
15     void add_head(T *n) {
16         if(first)
17             (first->*M).prev = n;
18         (n->*M).next = first;
19         first = n;
20     }
21 }
```

```
1 struct node {
2     node *next = nullptr, *prev = nullptr;
3 };
4
5 template <class T, node T::*M>
6 class list {
7     node head;
8
9     T* container_of(node *n) {
10    const size_t offset = (size_t)(&(((T*)nullptr)->*M));
11    return (T*)((char *)n - offset);
12    }
13
14 public:
15     T* begin() {
16     return head.next != &head ? container_of(head.next) : nullptr;
17     }
18 }
```

```
1 class matrix_element {
2     node row, column;
3 };
4
5 class matrix_row {
6     node n;
7     list<matrix_element, &matrix_element::row> row;
8 };
9
10 class matrix_column {
11     node n;
12     list<matrix_element, &matrix_element::column> column;
13 };
14
15 class matrix {
16     list<matrix_row, &matrix_row::n> rows;
17     list<matrix_column, &matrix_column::n> columns;
18 };
```

Example in C

```
1 struct tree_node{
2     list children;
3     node siblings; // Node for the list of children
4     struct tree_node *parent;
5 };
6
7 struct tree {
8     tree_node head;
9 };
```

Example in C++

```
1 class tree {
2     struct tree_node {
3         tree_node *parent;
4         node<tree_node> siblings;
5         list<tree_node, &tree_node::siblings> children;
6     };
7     tree_node head;
8 };
```

Circular vs. NULL-terminated

- If it is possible to use a circular linked list, it is better than NULL-terminated

Circular vs. NULL-terminated

- If it is possible to use a circular linked list, it is better than NULL-terminated

Non-intrusive `list<MyClass*>`

- Element can be in multiple lists
- Iterator cannot be obtained from a pointer to `MyClass` (e.g. to remove it from a list)
- Requires more memory and worse cache-efficiency

Circular vs. NULL-terminated

- If it is possible to use a circular linked list, it is better than NULL-terminated

Non-intrusive `list<MyClass*>`

- Element can be in multiple lists
- Iterator cannot be obtained from a pointer to `MyClass` (e.g. to remove it from a list)
- Requires more memory and worse cache-efficiency

Non-intrusive `list<MyClass>`

- Element is in one list only
- Some languages does not support obtaining an iterator from an instance
- Unsupported by some languages

Circular vs. NULL-terminated

- If it is possible to use a circular linked list, it is better than NULL-terminated

Non-intrusive `list<MyClass*>`

- Element can be in multiple lists
- Iterator cannot be obtained from a pointer to `MyClass` (e.g. to remove it from a list)
- Requires more memory and worse cache-efficiency

Non-intrusive `list<MyClass>`

- Element is in one list only
- Some languages does not support obtaining an iterator from an instance
- Unsupported by some languages

Intrusive

- Element can be in as many lists as it contains next pointers
- May require memory magic which is not available in some languages
- Appropriate for advanced algorithms which encapsulate lists
- Inappropriate between modules of complex programs

Do not use linked lists for everything

Array

- Significantly smaller overhead
- Better usage of cache, especially for small elements
- Sequential access does not require conditional jumps

Do not use linked lists for everything

Array

- Significantly smaller overhead
- Better usage of cache, especially for small elements
- Sequential access does not require conditional jumps

Hashing

- Expected complexity $O(1)$ for operations find, insert and delete
- Requires a good hashing function

Do not use linked lists for everything

Array

- Significantly smaller overhead
- Better usage of cache, especially for small elements
- Sequential access does not require conditional jumps

Hashing

- Expected complexity $O(1)$ for operations find, insert and delete
- Requires a good hashing function

Search trees

- Worst-case complexity $O(\log n)$
- Requires fast comparison of elements
- Support interval queries, finding successor element, etc.

Do not use linked lists for everything

Array

- Significantly smaller overhead
- Better usage of cache, especially for small elements
- Sequential access does not require conditional jumps

Hashing

- Expected complexity $O(1)$ for operations find, insert and delete
- Requires a good hashing function

Search trees

- Worst-case complexity $O(\log n)$
- Requires fast comparison of elements
- Support interval queries, finding successor element, etc.

Many other data structures

- Heaps
- Union-find
- ...

Popis

Exception handling is the process of responding to the occurrence of exceptions – anomalous or exceptional conditions requiring special processing – during the execution of a program.

Popis

Exception handling is the process of responding to the occurrence of exceptions – anomalous or exceptional conditions requiring special processing – during the execution of a program.

Examples of exceptions

- Index out of range of an array
- Retrieve the first element of an empty list
- Searched element does not exist
- Square root of a negative number
- Determine a circle from three colinear points
- File access without sufficient permission
- Error in reading data
- Invalid user input
- Insufficient amount of memory

Methods of handling exceptions

- Ignore, leads to an undefined behavior
`vector::operator[], list::pop_front`
- Throw an exception, call `longjmp`
`vector::at`
- Call a function
`set_terminate`, signal handling
- Set a status variable, e.g. `errno`
`fopen`, `scanf`
- Return an invalid value, e.g. `NULL`, a dummy object
`find`
- Return a valid value

```
double sqrt(double x) { if(x < 0) return 0; ... }
```
- Return a pair of a status and a result

```
pair<bool, int> find(x) { return exist(x) ? make_pair(true, search(x)) : make_pair(false, 0); }
```
- Use `goto` to an error label
- Terminate the program, e.g. using functions `abort` or `assert`

Return a valid value

- A calling function cannot recognize an error
- The error is propagated which makes it harder to find

Return a valid value

- A calling function cannot recognize an error
- The error is propagated which makes it harder to find

Set a status variable

- Hard to manage
- Requires detailed documentation
- Data race conditions

Return a valid value

- A calling function cannot recognize an error
- The error is propagated which makes it harder to find

Set a status variable

- Hard to manage
- Requires detailed documentation
- Data race conditions

Use goto to an error label

- Can be used only inside a function
- In rare situations, it may simplify a complex if-else statements

Return a valid value

- A calling function cannot recognize an error
- The error is propagated which makes it harder to find

Set a status variable

- Hard to manage
- Requires detailed documentation
- Data race conditions

Use goto to an error label

- Can be used only inside a function
- In rare situations, it may simplify a complex if-else statements

Call a function

Only for asynchronous or external events, e.g. signal handling.

Throw an exception

- Safe
- Checking conditions delays computation
- Requires catching exceptions
- Requires documentation of used exceptions

Throw an exception

- Safe
- Checking conditions delays computation
- Requires catching exceptions
- Requires documentation of used exceptions

Assert in a debug mode, ignore in a release mode

- Fast (in release mode)
- May lead to an undefined behavior
- Requires documentation of conditions on arguments

Return an invalid value

Prefer when

- the calling function is expected to handle
- the exception is a natural output (e.g. no element satisfies a given condition)

If an invalid values does not exist, return a pair with a status.

Return an invalid value

Prefer when

- the calling function is expected to handle
- the exception is a natural output (e.g. no element satisfies a given condition)

If an invalid values does not exist, return a pair with a status.

Throw an exception

Prefer when

- the exception may jump through many functions
- the situation is very rare
- many different types of errors may occur

Assert in a debug mode, ignore in a release mode

- Fast
- Easy to debug
- Cause undefined behavior in release mode

Assert in a debug mode, ignore in a release mode

- Fast
- Easy to debug
- Cause undefined behavior in release mode

A general rule

- Write documentation
- Always follow a project policy!

Example of an issue

Consider a function `push_back` inserting into a `vector<MyClass>` when its array is full. Reallocation moves all element into a new memory but move constructor throws an exception. What the function `push_back` should do?

Example of an issue

Consider a function `push_back` inserting into a `vector<MyClass>` when its array is full. Reallocation moves all element into a new memory but move constructor throws an exception. What the function `push_back` should do?

Exception safety

- No-throw guarantee: Operations are guaranteed to succeed and satisfy all requirements even in exceptional situations. If an exception occurs, it will be handled internally and not observed by clients.
- Strong exception safety: Operations can fail, but failed operations are guaranteed to have no side effects, so all data retains their original values.
- Basic exception safety: Partial execution of failed operations can cause side effects, but all invariants are preserved and there are no resource leaks (including memory leaks). Any stored data will contain valid values, even if they differ from what they were before the exception.
- No exception safety: No guarantees are made.