

Rules for teaching C as a first programming language

John C. Lusth

Revision Date: September 2, 2017

Introduction

This document concerns the teaching of C as a first language in a CS1 course. As it is often the desire to include the ideas of object encapsulation in a first course, C should be taught in a style that emphasizes encapsulation. Students who learn this style should have little trouble in moving to an object-oriented language later, especially if the advanced ideas (last section) are taught.

The remainder of this document assumes a familiarity with C and some object oriented language.

The basic idea

Instructors should always stress the object oriented nature of the program, emphasizing that each class gets its own module. In general, the public interface of a class *X* is placed in a file named *X.h*, often called a *header file*. The public interface consists of the structure that encapsulates the components of the class and the prototypes of the public functions (methods) that (typically) operate upon the structure. The public interface also includes external declarations of any components that are shared between objects of a class. The implementation of the class *X* is placed in a file named *X.c*. The implementation consists of the function definitions and variable declarations that were mentioned in the public interface along with definitions and declarations that are private (i.e., not visible outside the class).

An example class declaration

Here is the public interface for a *NODE* class, which can be used to make a linked list. The public interface is stored in a file named *node.h*:

```
#ifndef __NODE_INCLUDED__
#define __NODE_INCLUDED__

typedef struct node NODE; /* forward declaration of our structure */

extern char *INTEGER;
extern char *REAL;
extern char *STRING;

extern NODE *newNODEinteger(int value,NODE *next);
extern NODE *newNODEreal(double value,NODE *next);
extern NODE *newNODEstring(char *value,NODE *next);
extern char *getNODEtype(NODE *n);
extern int getNODEinteger(NODE *n);
extern double getNODEreal(NODE *n);
extern char *getNODEstring(NODE *n);
extern NODE *getNODEnext(NODE *n);
```

```
#endif
```

We can have three kinds of nodes, one that stores an integer value, one that stores a real value, and one that stores a string value.

We also declare external the three character constants to populate the *type* field, so we know what kind of node we have. Note that it is not necessary to explicitly declare the methods of the class external, but we do so, so that the instance variables and the method prototypes look consistent in the *.h* file.

The public interface exposes three node creating functions, named *newX*, depending on what kind of node you would like to create. It also exposes the access functions, named *getX*, to retrieve the various instance variables of the node. You may also have mutators, usually named *setX*; this example has no need of mutators, however.

The implementation, which should be stored in a file named *node.c*, looks like this:

```
#include <stdio.h>
#include <stdlib.h>
#include "node.h"

struct node
{
    char *type;
    int ival;
    double rval;    /* ival, rval, sval: one will hold the actual value */
    char *sval;
    NODE *next;
};

static NODE *newNODE(void);

/***** public interface *****/

char *INTEGER = "integer";
char *REAL = "real";
char *STRING = "string";

/* constructors */

NODE *
newNODEinteger(int v,NODE *n)
{
    NODE *p = newNODE();
    p->type = INTEGER;
    p->ival = v;
    p->next = n;
    return p;
}

NODE *
newNODEreal(double v,NODE *n)
{
    NODE *p = newNODE();
    p->type = REAL;
    p->rval = v;
    p->next = n;
    return p;
}
```

```

NODE *
newNODEstring(char *v,NODE *n)
    {
    NODE *p = newNODE();
    p->type = STRING;
    p->sval = v;
    p->next = n;
    return p;
    }

/* accessors */

char *getNodeEtype(NODE *n)    { return n->type; }
int   getNodeInteger(NODE *n)  { return n->ival; }
double getNodeEreal(NODE *n)   { return n->rval; }
char *getNodeEstring(NODE *n)  { return n->sval; }
NODE *getNodeEnext(NODE *n)    { return n->next; }

/***** private methods *****/

static NODE *
newNODE()
    {
    NODE *n = (NODE *) malloc(sizeof(NODE));
    if (n == 0) { fprintf(stderr,"out of memory"); exit(-1); }
    return n;
    }

```

We can see from the implementation that our node structure has separate fields for each of the types our node can store. We could use a union to save some space, but these days of lots of memory, it's hardly worth the trouble.

We can now make and manipulate nodes:

```

#include <stdio.h>
#include "node.h"

int
main()
    {
    NODE *n;

    n = newNODEinteger(3,0); /* zero is the null pointer */
    n = newNODEreal(5.5,n);
    n = newNODEstring("hello",n);

    while (n != 0)
        {
        char *t = getNodeEtype(n);
        if (t == INTEGER)
            printf("%d\n",getNodeInteger(n));
        else if (t == REAL)
            printf("%f\n",getNodeEreal(n));
        else if (t == STRING)
            printf("%s\n",getNodeEstring(n));
        }
    }

```

```

        n = getNODEnext(n);
    }

    return 0;
}

```

Here are links to this source code: [nodetest.c](#), [node.c](#), and [node.h](#).
 Compile *nodetest* with the command:

```
gcc -Wall -Wextra nodetest.c node.c -o nodetest
```

A more complicated example

The *X.h* file contains the forward definition of the typedef-ed structure representing class X as well as the prototypes for all the public methods. Consider a *STUDENT* class which holds the ID number of a student, a list of the classes he or she is currently taking, a list of past classes, and the student's cumulative gpa. The file named *student.h*, might look like:

```

#ifndef __STUDENT_INCLUDED__
#define __STUDENT_INCLUDED__

#include <stdio.h>
#include "class.h"

typedef struct student STUDENT;

extern STUDENT *newSTUDENT(int id);
extern void addSTUDENTclass(STUDENT *s, CLASS *c,int grade);
extern void setSTUDENTgrade(STUDENT *s, CLASS *c,int grade);
extern void setSTUDENTgpa(STUDENT *s, FILE *where);
extern void displaySTUDENT(STUDENT *s, FILE *where);

#endif

```

Note the typedef that gives a simple name to a *student* structure/object. Note also that *STUDENT* references another class named *CLASS*.

Public methods for a class follow a naming convention. For example, an insert method for objects of class *X* would be named *insertX*. Another convention is that the first argument for a class method is always the object.

Public versus private

Students should be well versed in the use of the keywords **static** and **extern** for controlling the visibility of definitions and declarations. Public components that are shared by all objects in a class are declared in the associated *.c* file. For example, suppose a shared component is the total number of students. This would be declared as a global in the *.c* file as:

```
int NumberOfStudents = 0;
```

if it is a public shared attribute. These public, shared components are also re-declared **extern** in the *.h* file, as in:

```
extern int NumberOfStudents;
```

If the component is to be shared, but made private, it would be declared static in the implementation file, as in:

```
static int NumberOfGoodStudents;
```

and would not appear in the header file.

Private methods in the implementation are placed in the `.c` file but are declared static.

Constructors

The constructor for a class X is named *newX*, by convention. It allocates space for the object and then initializes the components of the object. For example, if the student structure looked like:

```
struct student
{
    int id;
    CLASS **currentClasses;
    int currentClasseCount;
    CLASS **pastClasses;
    int pastClassCount;
    double gpa;
};
```

the constructor *newSTUDENT* might look like:

```
STUDENT *
newSTUDENT(int id)
{
    STUDENT *p;

    /* allocate the object */

    p = (STUDENT *) malloc (sizeof(STUDENT));

    if (p == 0) Error("Allocating a STUDENT: out of memory");

    /* initialize components */

    p->id = id;
    p->currentClasses = 0;
    p->currentClassCount = 0;
    p->pastClasses = 0;
    p->pastClassCount = 0;
    p->gpa = 0.0;

    ++NumberOfStudents;

    return p;
}
```

Multiple constructors require unique names. By convention, constructors for class X would all have the form *newXalpha*, where *alpha* is either empty or a description that distinguishes the constructor from the other constructors.

Encapsulating methods

It is easy enough for an object to “carry along” methods as function pointers, relieving the programmer from coming up with unique names for public methods. This is probably too much for a first semester course, however. Also, realize that the object must still be passed to these methods so that the object’s components are all accessible.

Here is the *STUDENT* structure rewritten in such a style:

```
#include "student.h"

struct student
{
    /* components */
    int id;
    CLASS **currentClasses;
    int currentClassCount;
    CLASS **pastClasses;
    int pastClassCount;
    double gpa;

    /* pointer to methods */
    void (*addClass)(STUDENT *s,Class *c,int grade);
    void (*setGrad)(STUDENT *s,Class *c,int grade);
    void (*setGPA)(STUDENT *s,FILE *);
    void (*display)(STUDENT *s,FILE *);
}

extern STUDENT *newSTUDENT(int id);
```

The constructor for *STUDENT* sets the function pointers in the methods section of the typedef:

```
STUDENT *
newSTUDENT(int id)
{
    STUDENT *p;

    /* allocate the object */

    p = (STUDENT *) malloc (sizeof(STUDENT));

    if (p == 0) Error("Allocating a STUDENT: out of memory");

    /* initialize components */

    p->id = id;
    p->currentClasses = 0;
    p->currentClassCount = 0;
    p->pastClasses = 0;
    p->pastClassCount = 0;
    p->gpa = 0.0;

    /* set the methods */

    p->addClass = addClass;
    p->addGrade = addGrade;
```

```

    p->setGPA = setGPA;
    p->display = display;

    ++NumberOfStudents;

    return p;
}

```

The methods *addClass*, *addGrade*, *setGPA*, and *display* are defined with the static keyword (making them private) in *student.h*.

Here is an example call to the *display* method:

```

STUDENT *s;

s = newSTUDENT(123456789);
/* add past classes */
...
s->setGPA(s,"transcript");
s->display(s,stdout);

```

Even when encapsulating methods, one still needs to pass the object itself to the method.

Inheritance

As inheritance is not allowed, code reuse is accomplished via wrapper functions that dispatch to clients. Instead of a class *X* inheriting from class *Y*, an object of type *X* instead has a component which holds a pointer to an object of type *Y*. Suppose inserting into an *X* object is really an insertion into the *Y* object. If so, the solution is to define a function *insertX* that simply calls *insertY* on object component *Y*.

Advanced ideas

In the above scenario, all of the components of an object are private, but the object itself is exposed and can be manipulated outside of the public interface. If the object itself must be private, then the constructors are reformulated to return an index into a private array of objects (similar to the *open* constructor for file descriptor objects). The object is now an integer rather than a pointer to a structure. The *newSTUDENT* constructor, in this case, might look like:

```

int
newSTUDENT(int id)
{
    STUDENT *p;

    ++NumberOfStudents;

    p = (STUDENT *) malloc (sizeof(STUDENT));

    if (p == 0) Error("out of memory");

    p->id = id;
    p->currentClasses = 0;
    p->currentClassCount = 0;
    p->pastClasses = 0;
    p->pastClassCount = 0;
    p->gpa = 0.0;
}

```

```
return add(p);
}
```

The *add* function resizes the array holding the allocated *STUDENT* objects to make room for the new object and returns the index corresponding to where the new object was placed:

```
static int
add(STUDENT *p)
{
    /* increase the number of students */

    ++NumberOfStudents;

    /* reallocate where the object is stored */

    Students = (STUDENT **) realloc(Students * sizeof(STUDENT *)
        * NumberOfStudents);

    if (Students == 0) Error("out of memory");

    /* store the student object and return its index */

    Students[NumberOfStudents - 1] = p;
    return NumberOfStudents - 1;
}
```

The *Student* array is also declared static in the *student.c* file:

```
static STUDENT **Students = 0;
```

Of course, the *add* routine should use array size doubling if a large amount of objects are to be allocated.

The downside to privatizing the components is that type checking on the first argument to class methods (the object argument) is weakened since the first argument is now an integer index. Also, there is no protection against arbitrarily generating Student indices. One can mitigate this latter concern by having the constructor return long integer keys that map to actual internal indices. A red-black tree can store the mapping between keys and actual indices.