


A short introduction to the C Programming Language

and a little about integration of C with Python

Table of Contents

Introduction: Compiled vs Interpreted	3
Hello World	13
Data Types (briefly)	24
printf	34
scanf	41
For-loop	44
Special operators (i++,i--, i+=1)	49 
While-loop	56
If-then-else	59
Grouping	62
Integer division & Casting	63
Integer number representation	68
Numeric data types	73
Symbolic constants	81
Functions	85
Stdin, Stdout, Stderr	93
File copying, character, line, word counting	96
Bools	106
Arrays	107
Strings	110
Common mistakes	118
sizeof	121
Pointers	125
Pointers and arrays	131
Memory allocation	137
Structs	144
Fibonacci example	149
Speed comparison	155
Cython	169
What next?	188
GPU computing	189
PyOpenCL	192
Writing to and reading from files in C	196
Reading command line arguments	201
Contiguous memory allocation (1d, 2d, 3d)	203

The C-programming language

Overview

- Compiled versus Interpreted language
- Python, Matlab and C
- Hello World
- basic numerical data types
- printf
- for-loops
- while-loops
- if-then, grouping

- integer division truncation
- type casting
- stdin, stderr, stdout
- Character input and output
- Arrays
- functions
- strings (s 99)
- Char array: print longest line
- Gotchas
- Summary introduction

Introduction

- Purpose of this short course: provide *an introduction* to the C programming language
- Will roughly cover chapter 1 of K&R, “*The C Programming Language*”, 2nd edition (1988), and a bit more
- Assume knowledge of either Python or Matlab
- Show how compiled code can be linked with high level code
- Aim to provide sufficient fundamental understanding to enable self learning

Instruct the computer

We use different levels of programming languages. We distinguish three fundamentally different ones:

- machine code (assembler: very low level machine code)
- higher level languages (C, Fortran, typically compiled to machine code)
- very high level language (Python, Matlab, commonly interpreted)

Machine code

- instructions that the Central Processing Unit (CPU) of a computer can understand and execute
- Machine code is *low-level code*. Typical statements are
 - ▷ load a datum (typically a number) from a memory address
 - ▷ add another number to it, increment it by one, . . .
 - ▷ save the datum to (another) memory address
 - ▷ read next instruction from some memory address
- machine code depends on the CPU and the Operating System (OS)
 - ▷ UNIX machine code cannot be executed on Windows PC
- machine code is stored in not-human readable format (binary)
 - ▷ can use *assembler* to read/write machine code
 - ▷ very tedious, only for extreme problems reasonable (if speed must be maximised at all costs)

Interpreted language: Example Matlab

- Source file (the m-file) is *parsed* line by line by the *interpreter* while it is being executed
- The interpreter converts the code into *machine code* which can be executed by the CPU
- For m-files, MATLAB is the interpreter
- The interpretation process is
 - ▷ hard work for the computer,
 - ▷ reason why interpreted languages are slower than compiled languages:
 - ▷ statements in for-loops are (generally) being interpreted again and again
- Interpreted code can be executed on any OS and CPU as long as interpreter is available

Compiled language: Example C

- Source file (the c-file) is *compiled* once
- The resulting machine code is stored in a file
- To execute the program, the file containing the machine code is executed
 - ▷ the source file is not required for execution
 - ▷ file with machine code is often called “executable”
 - ▷ Under Windows: executable files end with .exe
 - ▷ Under UNIX/Linux: any filename can be executable
- Compiled programs execute much faster than interpreted programs
 - ▷ mainly because machine code is generated in advance (when we compile the source) but not at run-time

- Compiled programs are harder to write
 - ▷ no interactivity when developing code
 - ▷ have to declare types of variables



Compiled and Interpreted languages: Examples

- Compiled
 - ▷ C, C++
 - ▷ Fortran
 - ▷ Pascal
 - ▷ Lisp
 - ▷ Ada
- Interpreted languages
 - ▷ MATLAB , Octave, Maple (specialised languages for numeric/symbolic computation)
 - ▷ Perl
 - ▷ Tcl
 - ▷ Basic (and dialects)
 - ▷ Lisp
 - ▷ Python (byte code)



Comparison compiled and interpreted languages

- Compiled languages
 - + execute very fast — tool of choice if speed is primary concern
 - don't need source code to execute (there may be commercial interest in not revealing the source code)
 - more complicated to write, debug and read code
- Interpreted languages
 - + writing code is easier than for compiled languages (saves a lot of time)
 - execution of code is slower (can differ by factor 10-1000 from compiled languages)
 - cannot execute without source code



C, Matlab and Python

- Assume reader knows at least one programming language
- Assume has either learned Matlab or Python in the past
 - ▷ Opportunity to learn Python: [FEEG1001](#)
 - ▷ Opportunity to learn Matlab: [SESG1009](#) (past)
- We compare Matlab/Python code with C throughout this course.
- We briefly look at combining Python and C



Why C?

- Prototyping and a lot of simulation can be done in high-level (interpreted) language (such as Python and Matlab)
- Sometimes, we need very fast code (\rightarrow C, Fortran)
- Parallel execution (\rightarrow C, Fortran): MPI, OpenMP
- Unix, Linux and Mac OS X are written in C
- high portability (ANSI C 90)
- Many languages and tools provide a C-interface
- Fundamental concepts important: should have seen C at least once.



Hello World

It is tradition to write a “Hello World” program when learning a new language.

- Matlab:

matlab/hello.m

```
fprintf('Hello World\n')
```

- Python:

python/hello2.py

```
print("Hello World")
```



Hello World (in C)

Our first C program:

c/hello.c

```
1 #include <stdio.h>
2
3
4 int main( void ) {
5
6     printf("Hello World!\n");
7
8     return 0;
9 }
```

- Line 1 includes the STandard Input Output (stdio) *header* file
 - ▷ provides commands such as `printf`
 - ▷ we will usually need (at least) this header file



- In line 4, the “main program” starts
 - ▷ the main program is a function and returns an integer (therefore `int`)
 - ▷ the main function has to be called “main” (therefore `main`)
 - ▷ the main function does not take arguments (therefore `void`)
 - ▷ the main function body starts after the opening curly brace (`{`)
- Line 6 is a print statement (very similar to MATLAB’s `fprintf`) but use double quotes to enclose format string
- Line 8 returns 0 to the operating system to indicate that the program executed and finished correctly (a convention).
- Line 9 ends the main function with the closing curly brace (`}`)
- *Every* statement (such as `printf` and `return`) has to be followed by a semicolon (`;`)
- For now, use lines 1–4 and 8–9 as a template (*i.e.* just copy them) for your C programs.



Compiling `hello.c`

- Compilation depends on the operating system and environment
- If command line compiler available (for example on UNIX/Linux/Mac OS X/MinGW):
 - ▷ Compile `hello.c` into executable `hello` using compiler `gcc`
`gcc -o hello hello.c`
 - ▷ `cc` is the C-compiler
 - ▷ `-o hello` states that the output file should be called `hello`
 - ▷ `hello.c` is the name of the input file (typically the last argument)
 - ▷ many possible compilers available, including `gcc`, Portland Group Compiler, . . . but all have similar syntax (as shown above)
 - ▷ For `gcc`, we recommend switches `-ansi -pedantic -Wall`, i.e.:
`gcc -ansi -pedantic -Wall -o hello hello.c`
- A good C-compiler is GNU's `gcc` (see www.gnu.org for details). Can be used on UNIX/Linux, Windows and many other platforms.



- A simple and free graphical user interface and editor for GCC (on Windows) is Quincy → laboratories
- On Linux/Mac, use editor that supports C (Emacs, ...) and command prompt to call gcc.



Execution of hello

- On UNIX/Linux
 - ▷ `./hello`
- On Windows
 - ▷ Either find `hello.exe` and type `hello.exe` in MS-DOS prompt
 - ▷ or click on appropriate button to execute (→ labs)
- Output is
`Hello World!`

Hello World (in B)

The first ever “Hello World” program is believed to have been written for the programming language B:

b/hello.b

```
1 main( ) {
2     extrn a, b, c;
3     putchar(a); putchar(b); putchar(c); putchar('!*n');
4 }
5 a 'hell';
6 b 'o, w';
7 c 'orld';
```

Source: http://en.wikipedia.org/wiki/Hello_world_program (accessed 1/10/2010)

ooo

→ **lab 1**

Absence of restriction

Central design ideas of the C programming language were *economy of expression* and *absence of restriction*.

This requires some discipline from the programmer. How useful are these hello world programs?

c/hello3.c

```
1 #include <stdio.h>
2 main() {printf("Hello World!\n");}
```

c/hello4.c

```
1 #include <stdio.h>
2 int main(
3 int argc, char* argv[]
4 ) {printf("Hello World!\n")
5 ;return 0;}
```

Summary Hello World

- Write C-program in so-called source file (`hello.c`)
 - ▷ looks similar to m-file, py-file, etc
- Compile source to create executable file
 - ▷ Can choose name of executable, usually same root as source file but no extension or `.exe`, (*i.e.* `hello` or `hello.exe`)
 - ▷ executable is binary file with machine code
- Run program by executing executable
 - ▷ Don't need source code any more for execution of executable
 - ▷ Don't need C compiler for execution of executable

Execution comparison with Matlab

- Comparison with MATLAB

- ▷ Could have achieved the same Hello World greeting in MATLAB with

matlab/hello.m

```
fprintf('Hello World\n')
```

```
[fangohr@lyceum ~]$ matlab -nojvm -nodisplay -r hello2
```

```
< M A T L A B (R) >
```

```
Copyright 1984-2008 The MathWorks, Inc.
```

```
Version 7.7.0.471 (R2008b)
```

```
September 17, 2008
```

```
To get started, type one of these: helpwin, helpdesk, or demo.
```

```
For product information, visit www.mathworks.com.
```

```
Hello World
```

- ▷ Matlab executes the m-file given after the `-r` switch.

Execution comparison with Python

- Python Hello World program

python/hello2.py

```
print("Hello World")
```

- needs the Python interpreter to be executed:

```
[fangohr@lyceum ~]$ python hello.py  
Hello World
```

- Python executes the argument (hello.py) given to the interpreter program (python).



Data types

C knows 4 basic data types:

- `char` (character): letters, *i.e.* 'a', 'b', . . .
- `int` (integer): integers, *i.e.* 1, 10, -344, 10000, . . .
- `float` (floating point number): *i.e.* 1.1, -10.34, 1.41, . . .
- `double` (more accurate floating point number)

Chars and ints can be signed (default) or unsigned. ints can be short or long. The actual number of bits used for short int, int, long int, float, double, and long double is implementation dependent (see standard headers `<limits.h>` and `<float.h>`.)



Whenever we want to use a variable in C, we have to say (“declare”) in advance of which type it is!

For example:

```
int a;          /* declares an integer variable with name a */
double f;      /* ... a double float with name f */
int i=0;       /* declares the variable i to be of type int and
               initialises it with the value 0 */
```

Useful to combine *declaration* of variable (for example `int i;` with the *initialisation* of the variable (for example `int i=0;`. This avoids accidental use of `i` with an undefined (i.e. random) value.



Using variables in C

- C-Code

c/var1.c

```
1 #include <stdio.h>
2 int main( void ) {
3     int a, b, c;    /* define variables first:
4                     a, b and c are of type int */
5     a=10;          /* assign values */
6     b=20;
7     c = a + b;     /* do computation */
8     printf("The sum of %d and %d is %d.\n",a,b,c);
9     return 0;
10 }
```

- Output

The sum of 10 and 20 is 30.



- Comments

- ▷ text inclosed by `/*` and `*/` is ignored by the compiler
- ▷ The latest ANSI C (and C++) allows to use `//` to ignore the rest of the line (like MATLAB 's `%` and Python's `#`) but thsi is not part of the ANSI C90 standard (and the switch `-ansi` will prevent usage of `//`).

- Equivalent programs in Python and Matlab

python/var1.py

```
a = 10
b = 20
c = a + b
print("The sum of %d and %d is %d." % (a, b, c))
```

matlab/var1.m

```
a=int32(10); %int32() to get an integer: by default
b=int32(20); %numbers are doubles in Matlab
c=a+b;
fprintf('The sum of %d and %d is %d.\n',a,b,c)
```



When we forget to declare variables . . .

c/var1wrong1.c

```
1 #include <stdio.h>
2 int main( void ) {
3     a=10;          /* assign values */
4     b=20;
5     c = a + b;    /* do computation */
6     printf("The sum of %d and %d is %d.\n",a,b,c);
7     return 0;
8 }
```

- Output from `gcc -ansi -pedantic -Wall -o var1wrong1 var1wrong1.c`

```
var1wrong1.c: In function main :
var1wrong1.c:3: error: a undeclared (first use in this function)
var1wrong1.c:3: error: (Each undeclared identifier is reported only once
var1wrong1.c:3: error: for each function it appears in.)
var1wrong1.c:4: error: b undeclared (first use in this function)
var1wrong1.c:5: error: c undeclared (first use in this function)
```



When we forget a semicolon . . .

c/var1wrong2.c

```
●
1 #include <stdio.h>
2 int main( void ) {
3     int a, b, c;    /* define variables first:
4                     a, b and c are of type int */
5     a=10           /* assign values */
6     b=20;
7     c = a + b;    /* do computation */
8     printf("The sum of %d and %d is %d.\n",a,b,c);
9     return 0;
10 }
```

- Output from gcc -ansi -pedantic -Wall -o var1wrong2 var1wrong2.c

```
var1wrong2.c: In function main :
var1wrong2.c:6: error: expected ; before b
```



Summary Compilation

- Writing compiled code is not easy
 - ▷ have to get everything right, *i.e.*
 - ▷ declare variables before using them
 - ▷ stick to strict syntax (don't forget semicolon, parenthesis, braces, . . .)
 - ▷ otherwise the compiler will issue an error message and stop
- Compiler checks syntax of code at compile time



Documentation

- It is similarly important to document C-code as it is for any other code. For example

c/hellodoc.c

```
1  /* hello.c - a program printing "Hello World!" to the screen
2  Hans Fangohr, March 2003, University of Southampton, fangohr@soton.ac.uk
3
4  History: - inspired by standard C-greeting program
5  Purpose: - to provide students a small example of working C-code
6  Modifications:
7
8      Sept 2010: minor changes for SESG6025
9  */
10 #include <stdio.h>
11 int main( void ) {           // The latest ANSI C (>1999) allows one-line
12                               // comments starting with a double slash
13                               // but best not to rely on this.
14     printf("Hello World!\n");
15     return 0;
16 }
```




- Syntax for comments:
 - ▷ Enclose (long) comments between `/*` and `*/`
 - ▷ New feature: Use `//` to mark end of line as a comment
 - * This was introduced in C++, and
 - * later (around 1998) made its way back into ANSI C, but not part of ANSI C90 and gcc's switch `-ansi` prevents usage of `//`.
 - * Best avoided to maximise portability of our code.
- Suggestions for comments:
 - ▷ Do not explain what is visible anyway:
`s = sta + stu; /* computing the sum of sta and stu */`
instead provide the context or reason for the operation, for example
`s = sta + stu; /* total number s of STAff and STUdents */`
 - ▷ if a piece of code is self-explaining, there is no need to add comments.
- Indentation:
 - ▷ Commonly, deeper levels of functions/blocks of statements/etc are indented according to their depth.



- ▷ However, C does not require this (as commands are grouped through the use of curly braces), so it is the programmers responsibility to indent code correctly (most editors will help).
- ▷ A lot of time has been wasted through not correctly indented code.



printf

- the `printf` function is used to PRINT Formatted values to the standard output (stdout).
- Same principle as in Matlab/Python: a string with format specifiers is followed by a list of values, such as

```
printf("%d bottles of milk at %4.2f each cost %5.2f in total",5,1.42,5*1.42);
```

with output

```
5 bottles of milk at 1.42 each cost 7.10 in total.
```

- Import specifiers are
 - ▷ `d` and `i` for *signed integers, short, chars* (int, char, short)
 - ▷ `ld` and `li` for *signed long integers* (long)



- ▷ `f` for *floating point representation of double* (double)
 - ▷ `F` for *floating point representation of float* (float)
 - ▷ `e` for *exponential notation of float and double* (double,float)
 - ▷ `E` for *exponential notation or floating point representation: the shorter one is chosen* (double,float)
 - ▷ `s` for *arrays of characters* (char *, char[])
 - ▷ `u` for *unsigned integers* (unsigned int, short, char)
- The `printf` function is part of the standard library and we must include `<stdio.h>` in the beginning of the file.
 - Related commands `fprintf` and `sprintf` to print into file stream and string, respectively.
 - Choosing the wrong format specifier results in **undefined behaviour**.

Example 1

c/printf1.c



```
1 #include <stdio.h>
2 int main(void){
3     int var1=-42;
4     unsigned int var2 =142;
5     float var3=3.14;
6     double var4 = 3.1428;
7     char var5[] = "Hello World";
8     printf("Printing int var1=%d\n",var1);
9     printf("Printing unsigned int var2=%u\n",var2);
10    printf("Printing float var3=%f\n",var3);
11    printf("Printing double var4=%f\n",var4);
12    printf("Printing double in scientific notation var4=%e\n",var4);
13    printf("Printing char[] var5=%s\n",var5);
14    return 0;
15 }
```

Output:

```
Printing int var1=-42
Printing unsigned int var2=142
Printing float var3=3.140000
Printing double var4=3.142800
Printing double in scientific notation var4=3.142800e+00
```



Printing char[] var5=Hello World

Example 2

c/printf2.c

```
1 #include <stdio.h>
2 int main(void){
3     double pi = 3.1415926535897931;
4     printf("as a standard float representation: pi=%f\n",pi);
5     printf("in exponential notation: pi=%e\n",pi);
6     printf("whichever of the above two is shorter: pi=%g\n",pi);
7     printf("Request 10 digits overall: pi=%10f\n",pi);
8     printf("With 3 postdecimal digits: pi=%10.3f\n",pi);
9     printf("12 postdecimal digits: pi=%.12f\n",pi);
10    return 0;
11 }
```

Output:

```
as a standard float representation: pi=3.141593
in exponential notation: pi=3.141593e+00
whichever of the above two is shorter: pi=3.14159
Request 10 digits overall: pi= 3.141593
```



With 3 postdecimal digits: $\pi = 3.142$
12 postdecimal digits: $\pi = 3.141592653590$



Another printf example

c/printf0.c

```
1 #include <stdio.h>
2 int main(void){
3     double price=1.42;
4     int n=5;
5     double total=0.;
6     total = n*price;
7     printf("%d bottles of milk at %4.2f "\
8           "each cost %5.2f in total.",
9           n,price,total);
10    return 0;
11 }
```

c/printf0.out

```
5 bottles of milk at 1.42 each cost 7.10 in total.
```




python/printf0.py

```
price = 1.42
n = 5
total = n * price
print("%d bottles of milk at %4.2f "\
      "each cost %5.2f in total." \
      % (n, price, total))
```

matlab/printf0.m

```
price=1.42;
n=5;
total=n*price;
fprintf('%d bottles of milk at %4.2f ',n,price);
fprintf('each cost %5.2f in total.\n',total);
```

scanf()



- This is a useful function for enabling user input to a program. You may not fully understand the use of it until future lectures but for now it can be used as a template to enable more interesting programs.
- The function `scanf()` is very similar to `printf()` and uses the same formatting rules but in the other direction.



- The basic form is as seen here:

c/scanf0.c

```
1 #include <stdio.h>
2
3 int main(void)
4 {
5     int answer;
6     scanf("%d",&answer);
7     printf("You entered %d.\n",answer);
8     return 0;
9 }
```

- the integer variable `answer` is declared prior to the use of `scanf`. The same `%d` format code is used within the quotation marks as for `printf` when using an integer. The variable that input is being assigned to is the second argument and needs the `&` before it.

- When using scanf to accept user input in the form of char values, it may be necessary to use the form:



c/scanf1.c

```
1 #include <stdio.h>
2
3 int main(void)
4 {
5     char answer;
6     scanf(" %c", &answer);
7     printf("You entered %c.\n", answer);
8     return 0;
9 }
```

the space before the % allows any white space still in the input buffer to be ignored before the most recent user input is retrieved.

The for-loop in C



- Example: printing numbers $i=0$ to $i=9$. Desired output

0 1 2 3 4 5 6 7 8 9

- In Matlab

matlab/forloop1.m

```
for i=0:9
    fprintf('%d ',i);
end
```

- In Python (using print in C-style)

python/forloop1.py

```
for i in range(10):
    print("%d" % i),
```



- In C:

c/forloop1.c

```
1 #include <stdio.h>
2 int main( void ) {
3     int i;
4     for (i=0; i<10; i = i + 1) {
5         printf("%d ",i);
6     }
7     return 0;
8 }
```

- `for`-command has 4 parts:

```
for ( A ; B ; C ) {
    D}
```

- ▷ A: initial assignment (executed once)
- ▷ B: logical expression, for-loop continues while this expression is true
- ▷ C: assignment to be executed *after* every loop
- ▷ D: block of statements to be executed in every loop (limited by enclosing curly braces)

- Improving on our programming style: make loop more general



▷ count from 0 to $N - 1$

c/forloop1a.c

```
1 #include <stdio.h>
2 int main( void ) {
3     int i;
4     int N;
5     N=10;
6     for (i=0; i<N; i=i+1) {
7         printf("%d\n",i);
8     }
9     return 0;
10 }
```

- ▷ can *initialise* variable when *declaring* it: combine lines 4 and 5 (nothing do to with for-loop):



c/forloop1c.c

```
1 #include <stdio.h>
2 int main( void ) {
3     int i;
4     int N = 10;
5     for (i=0; i<N; i=i+1) {
6         printf("%d ",i);
7     }
8     return 0;
9 }
```




- Use C-typical notation:

▷ `i = i + 1;` can be written as `i++;` (or `++i;`) using the “increment operator”

c/forloop1b.c

```
1 #include <stdio.h>
2 int main( void ) {
3     int i;
4     int N = 10;
5     for (i=0; i<N; i++) {
6         printf("%d ",i);
7     }
8     return 0;
9 }
```

ooo

→ lab 2



The increment operator ($i=i+1$)

- $i = i + 1$ can be written as $i++$ or $++i$

c/increment0.c

```
1 #include <stdio.h>
2 int main(void) {
3     int k=5,n=0;
4     n = k;
5     k = k + 1; /* could also write 'k++;' or '++k' */
6     printf("case 1: -> k=%d, n=%d\n", k, n);
7     k = 5; n=0; /* reset to original values */
8     k = k + 1; /* could also write 'k++;' or '++k' */
9     n = k;
10    printf("case 2: -> k=%d, n=%d\n", k, n);
11    return 0;}
```

Output:

case 1: -> k=6, n=5

case 2: -> k=6, n=6



- The `++` operator can be used *within* other statements (such as – in this example – assignments), and then it matters whether prefix notation (`++i`) or postfix notation (`i++`) is being used:

c/increment.c

```
1 #include <stdio.h>
2 int main(void) {
3     int k=5,n=0;
4     n = k++;      /* instead of: n=k; k=k+1 */
5     printf("postfix: n=k++ -> k=%d, n=%d\n", k, n);
6     k = 5; n=0;
7     n = ++k;     /* instead of: k=k+1; n=k */
8     printf("prefix : n=++k -> k=%d, n=%d\n", k, n);
9     return 0;
10 }
```

Output:

```
postfix: n=k++ -> k=6, n=5
prefix : n=++k -> k=6, n=6
```



Other concise expressions

- ▷ `i--` and `--i` are postfix and prefix decrements
- ▷ `i += c` \leftrightarrow `i = i + c`
- ▷ `i -= c` \leftrightarrow `i = i - c`
- ▷ `i *= c` \leftrightarrow `i = i * c`
- ▷ `i /= c` \leftrightarrow `i = i / c`

Legal but not soo nice (example for-loop)

c/forloop1d.c

```
●
1 #include <stdio.h>
2 int main( void ) {
3     int i;
4     int N = 10;
5     for (i=0; i<N; printf("%d ",i++)) {
6         ;
7     }
8     return 0;
9 }
```

c/forloop1e.c

```
1 #include <stdio.h>
2 int main(void){int i;int N=10; for(i=0;i
3 <N;printf("%d ",i++)){;}return 0;}
```



Tabulating \sqrt{x}

A more complex example: computing a table of square roots as shown below

```
sqrt(1.000000) = 1.000000
sqrt(1.500000) = 1.224745
sqrt(2.000000) = 1.414214
sqrt(2.500000) = 1.581139
sqrt(3.000000) = 1.732051
sqrt(3.500000) = 1.870829
sqrt(4.000000) = 2.000000
sqrt(4.500000) = 2.121320
sqrt(5.000000) = 2.236068
sqrt(5.500000) = 2.345208
sqrt(6.000000) = 2.449490
sqrt(6.500000) = 2.549510
sqrt(7.000000) = 2.645751
sqrt(7.500000) = 2.738613
sqrt(8.000000) = 2.828427
sqrt(8.500000) = 2.915476
sqrt(9.000000) = 3.000000
sqrt(9.500000) = 3.082207
sqrt(10.000000) = 3.162278
```

can be generated using



c/forloop2.c

```
1 #include <stdio.h>
2 #include <math.h>      /* include math-header for sqrt-function */
3
4 int main( void ) {
5     /* print table of N square roots sqrt(x) for x in [a,b] */
6
7     int N = 19;        /* total number of lines in table */
8     double a=1;       /* starting with x=a */
9     double b=10;      /* ending with x = b */
10    double x;          /* being used in for-loop */
11    int i;             /* iteration counter for for-loop */
12
13    for (i=0; i<N; i++) {
14        x = a + i*(b-a)/(N-1); /* compute x */
15
16        /* compute sqrt(x) and print result */
17        printf(" sqrt(%f) = %f\n", x, sqrt(x) );
18    }
19
20    return 0;
21 }
```

In Matlab/Python (can be done even shorter in both)

matlab/forloop2.m

```
%print table of N square roots sqrt(x) for x in [a,b]  
N=19;  
a=1;  
b=10;  
for i=0:N-1  
    x = a + i*(b-a)/(N-1.);  
    fprintf('sqrt(%f) = %f\n',x,sqrt(x));  
end
```

python/forloop2.py

```
"""print table of N square roots sqrt(x) for x in [a,b]"""  
from math import sqrt  
N = 19  
a = 1  
b = 10  
for i in range(N):  
    x = a + i * (b - a) / (N - 1.)  
    print("sqrt(%f) = %f" % (x, sqrt(x)))
```



The while loop in C



c/while.c

```
1 #include <stdio.h>
2 int main( void ) {
3     double x = 1;
4     while (x < 50) {
5         printf("In this iteration , x=%f\n",x);
6         x = x * 1.5;
7     }
8     return 0;
9 }
```

The `while`-loop has two parts:

```
while ( A ) {
    B
}
```

- A: logical expression, while-loop continues while this expression is true
- B: block of statements to be executed in every loop (limited by enclosing curly braces)

The for-loop and the while loop are closely related in C. In fact, this for-loop



```
for ( A; B; C ) { D; }
```

is equivalent to this while-loop:

```
A;  
while( B ) { D; C; }
```

Output:

```
In this iteration , x=1.000000  
In this iteration , x=1.500000  
In this iteration , x=2.250000  
In this iteration , x=3.375000  
In this iteration , x=5.062500  
In this iteration , x=7.593750  
In this iteration , x=11.390625  
In this iteration , x=17.085938  
In this iteration , x=25.628906  
In this iteration , x=38.443359
```



Python:

python/while.py

```
x = 1
while x < 50:
    print("In this iteration, x=%f" % x)
    x = x * 1.5
```

Matlab:

matlab/while.m

```
x=1;
while x < 50
    fprintf("In this iteration, x=%f\n", x);
    x = x * 1.5;
end
```

If-then statements



- MATLAB

matlab/ifthen1.m

```
a=10;  
b=20;  
if a==b  
    fprintf('a is the same as b\n')  
else  
    fprintf('a is not equal to b\n')  
end
```

- Python

python/ifthen1.py

```
a = 10  
b = 20  
if a == b:  
    print('a is the same as b')  
else:  
    print('a is not equal to b')
```



- C

c/ifthen1.c

```
1 #include <stdio.h>
2 int main( void ) {
3     int a=10;
4     int b=20;
5     if (a==b) {
6         printf("a is the same as b\n");
7     }
8     else {
9         printf("a is not equal to b\n");
10    }
11    return 0;
12 }
```

In C,

- the logical expression `(a==b)` has to be placed in parenthesis in `if`-statement (5)
- block of statements to be executed (6) if `a==b` is true has¹ to be enclosed by

¹This is not strictly true, but we strongly advise you to do this. If there are no curly braces, then only the first



curly braces (5, 7)

- block of statements for the `else`-part (16) has to be enclosed in curly braces (8, 10)

statement is considered to be a part of the `if`-statement. Using curly braces is the safe way.

Grouping

The grouping of commands to form a block is done

- in Matlab: through *some keyword* (such as `for`, `if`, `else`, `while`) followed by the keyword `end`.
- in C: through enclosing curly braces `{` and `}`.
- in Python: through *some keyword* (such as `for`, `if`, `else`, `while`, `def`) followed by a colon (`:`) and indenting all commands in the block to the same level.

ooo

→ **lab 3**

Integer division truncates



Integer division truncates:

c/intdiv.c

```
1 #include <stdio.h>
2
3 int main( void ) {
4     int a=10; int b=3; double c;
5
6     c = b/a;
7
8     printf("a=%d, b=%d, c=%f\n", a, b, c);
9
10    return 0;
11 }
```

The output is:

a=10, b=3, c=0.000000



Solution: “cast” `a` or `b` into a double (line 6):

`c/intdiv1a.c`

```
1 #include <stdio.h>
2
3 int main( void ) {
4     int a=10; int b=3; double c;
5
6     c = (double) b/a;      /* OR c = b/(double) a; */
7
8     printf("a=%d, b=%d, c=%f\n", a, b, c);
9
10    return 0;
11 }
```

The output is

`a=10, b=3, c=0.300000`



Casting (changing type)

If required, we can change the type of a variable by *casting* it. This is done by the unary operator *(newtype) variable-of-old-type*, for example:

c/sqrt_int.c

```
1 #include <stdio.h>
2 #include <math.h>
3
4 /* main: sqrt expects a float or double input argument */
5 int main(void){
6     int k=9;                /* but have int */
7     printf("sqrt(%d)=%f\n", k,
8           sqrt( (double) k) ); /* cast int to double
9                               before calling sqrt */
10    return 0;
11 }
```

Output:

```
sqrt (9)=3.000000
```



Integer division truncates, another example

- Watch for intermediate integer calculations

c/intdiv2.c

```
1 #include <stdio.h>
2
3 int main( void ) { /* convert from fahrenheit to celsius */
4
5     double celsius;
6     double fahrenheit;
7     fahrenheit = 82;
8
9     celsius = (fahrenheit - 32) * (5/9);
10
11     printf("%f degree Fahrenheit are %f degree Celsius.\n",
12           fahrenheit, celsius);
13
14     return 0;
15 }
```

Output:

```
82.000000 degree Fahrenheit are 0.000000 degree Celsius.
```

- Reason: `(5/9)` is integer operation, thus it is truncated, thus `(5/9)` evaluates to zero! (line 9)



- Solution: modify 5 or 9 (or both) to be floating point numbers

- ▷ `celsius = (fahrenheit-32)*(5.0/9);`
- ▷ `celsius = (fahrenheit-32)*(5/9.0);`
- ▷ `celsius = (fahrenheit-32)*((double) 5/9);`
- ▷ `celsius = (fahrenheit-32)*(5/(double) 9);`
- ▷ `celsius = 5*(fahrenheit-32)/9;`

Numbers (in mathematics)

- Natural numbers:

$$\mathbb{N} = \{0, 1, 2, 3, \dots\}$$

- Integer numbers:

$$\mathbb{Z} = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$$

- Rational numbers : (all fractional numbers, fraction = ratio)

$$\mathbb{Q} = \left\{ \frac{m}{n} : m \in \mathbb{Z}, n \in \mathbb{Z}^* \right\}$$

- Real numbers (irrational numbers):

$$\mathbb{R} = \mathbb{Q} \cup \{ \text{“all other scalars”} \}$$

Examples for $x \in \mathbb{R}$: $x = \sqrt{2}$, $x = \pi$, $x = \exp(1) = e$

- (for completeness): Complex numbers:

$$\mathbb{C} = \{a + ib : a \in \mathbb{R}, b \in \mathbb{R}\} \quad \text{and} \quad i = \sqrt{-1}$$

- Formally:

$$\mathbb{N} \subset \mathbb{Z} \subset \mathbb{Q} \subset \mathbb{R}$$

Memory: bits and bytes



- Computers store information in “bits”
- a bit is either 1 or 0
 - ▷ RAM: low voltage or high voltage
 - ▷ Hard disk: magnetisation pointing up or down
- eight bits are one byte
- $1024 \approx 10^3$ byte are one kilo byte (kb)
- $1024^2 \approx 10^6$ byte are one mega byte (Mb)
- $1024^3 \approx 10^9$ byte are one giga byte (Gb)

Natural and integer numbers

How can we represent integer numbers in a computer using bits?

- Suppose we have one byte (i.e. eight bit):

natural number	bit-representation
0	00000000
1	00000001
2	00000010
3	00000011
4	00000100
5	00000101
⋮	⋮
254	11111110
255	11111111

- with 8 bit we can represent 256 natural numbers (for example from 0 to 255 as in the `unsigned char` type) – $2^8 = 256$.

- with 8 bit we can also represent 256 integer numbers (for example from -127 to +128 as in the (signed) `char` type)

(Numeric) data types



The numeric data types of C are (K&R, page 9)

- `char`: character - a single byte
- `short`: short integer
- `int`: integer (16 bits [between -32768 and +32767] or 32 bits [between -2147483648 and 2147483647])
- `long`: long integer
- `float`: single precision floating point (typically 32 bit with at least 6 significant digits and magnitude between 10^{-38} and 10^{38})
- `double` double-precision floating point (typically 2x32 bit with at least 15 significant digits and magnitude between 10^{-308} and 10^{308})



Guidelines:

1. Use `double` whenever you need floating point numbers
2. Do not use `float` (unless you know exactly what you are doing)
 - `float` is less accurate than `double`
 - with today's PCs (great CPU-power and memory) there is no need to use `floats` anymore: computation with floats is generally not slower than computation with doubles.
 - the main reason for using `floats` instead of `doubles` at all is (usually) to reduce memory requirements

Note: when people speak about “floats” they probably mean doubles – terminology varies from language to language. Python, for example, has only one data type “float” which is in C-terms a double.



Floating point numbers (“floats” and “double”)

- Any real number x (thus also integers) can be written as

$$x = a \cdot 10^b$$

where a is the mantissa and b the exponent.

- Examples:

x		a	b
123.456	=	1.23456	2
100000	=	1.0	6
0.0000024	=	2.4	-6

Can write a and b as integers:

123456	2
100000	6
240000	-6

- → we can use 2 integers to code one floating point number



$$x = a \cdot 10^b$$

- For example: x is an 8 byte float: these 64 bits are split as
 - ▷ 10 bit for the exponent b and
 - ▷ 54 bit for the mantissa a .

This results in

- ▷ largest possible float $\approx 10^{308}$
 - ▷ smallest possible float $\approx 10^{-308}$ (quality measure for b)
 - ▷ distance between 1.0 and next larger number $\approx 10^{-16}$ (quality measure for a)
- This is *in principle* how floating point numbers are represented.

Numbers in computer

- Mostly, we use *approximations* to the real number, i.e. $\pi \approx 3.14159265358979$.
- Often, this doesn't matter
- but sometimes it does:

```
>>sin(pi)
ans =
    1.2246e-16
```

- Even “simple floating point numbers” cannot be represented exactly (using Python to demonstrate but works the same in C):

```
>>> 0.0
0.0
```

```
>>> 0.1
0.100000000000000001
>>> 0.2
0.200000000000000001
>>> 0.3
0.299999999999999999
>>> 0.4
0.400000000000000002
>>> 0.5
0.5
```

- Be aware that
 - ▷ there is an upper limit to possible floats
 - ▷ there is a lower limit to possible floating point numbers
 - ▷ the next largest number to 1 is `1+eps`
 - ▷ see `float.h`
- This is an import issue in computation

- ▷ we'd like to know *how* inaccurate our results are,
- ▷ we look for *stable* numerical methods which don't amplify small errors.



Example program to determine eps

eps is the 'distance' between 1.0 and the next larger number (i.e. $1.0 + \text{eps}$) that can be represented:

c/eps.c

```
1 #include <stdio.h>
2 int main(void) {
3     double eps = 1.0;
4     while (1.0+eps>1.0) {
5         eps /= 2.0;
6     }
7     printf("eps is approximately %g.\n", eps);
8     return 0;
9 }
```

eps is approximately 1.11022e-16.

Symbolic constants



(K&R, Sec 1.4, page 14)

If we use constants in a program, it is good practice to define a *symbolic name* or *symbolic constant*:

```
#define name replacement-text
```

The pre-processor (runs silently before the actual compilation) will replace all occurrences of *name* in the file with *replacement-text*.

For example: allow to change type of int easily.



With a (32-bit) int integer variable:

c/define.c

```
1 #include <stdio.h>
2 #define MYINT int
3
4 MYINT square(MYINT x) {
5     return x*x;
6 }
7
8 int main(void) {
9     MYINT a=3;
10    printf("Myint a=%d uses %ld bytes. %d^2=%d\n",
11           a, sizeof(a), a, square(a));
12    return 0;
13 }
```

Output:

Myint a=3 uses 4 bytes. 3^2=9



With char integer variable (here 8 bit):

c/define2.c

```
1 #include <stdio.h>
2 #define MYINT char
3
4 MYINT square(MYINT x) {
5     return x*x;
6 }
7
8 int main(void) {
9     MYINT a=3;
10    printf("Myint a=%d uses %ld bytes. %d^2=%d\n",
11           a, sizeof(a), a, square(a));
12    return 0;
13 }
```

Output:

```
Myint a=3 uses 1 bytes. 3^2=9
```

The only change is in line 2.



We can also use `#define` to define global constants:

`c/symbolicconstant.c`

```
1 #include <stdio.h>
2
3 #define LOWER 0
4 #define UPPER 100
5 #define STEP 20
6
7 int main () /*print Fahrenheit-Celsius table */
8 {
9     int fahr;
10    for (fahr=LOWER; fahr <= UPPER; fahr += STEP) {
11        printf("%3d %6.1f\n", fahr, 5.0/9.0*(fahr-32));
12    }
13    return 0;
14 }
```

Output:

```
 0  -17.8
20   -6.7
40    4.4
60   15.6
80   26.7
100  37.8
```

Functions



c/hello2.c

```
1 #include <stdio.h>
2 int main( void ) {
3     printf("Hello World!\n");
4     return 0;
5 }
```

- C programs consist only of functions
- there has to be exactly one “main”-function with name `main` (this is executed first, and may call other functions)
- have to define what argument(s) each function takes (line 2)
- have to define what type these arguments are (different from Matlab and Python): in example above (line 2):
 - ▷ main takes no arguments (`void`) and
 - ▷ main returns an `int` (line 4)



Functions, example 1

c/func1.c

```
1 #include <stdio.h>
2
3 int square( int a ) {
4     return a*a;
5 }
6
7 int main( void ) {
8     int mynumber;           /* declare variable */
9     mynumber = 4;          /* initialise variable */
10
11     printf("mynumber is %d, and mynumber squared is %d.\n",
12           mynumber, square(mynumber) );
13
14     return 0;              /* return 0 to operating system */
15 }
```

Two functions: `main` (lines 7-15) and `square` (lines 3-5)

Output:

```
mynumber is 4, and mynumber squared is 16.
```



In summary, a function definition has this form (K&R, p25):

```
return-type function-name(parameter declaration, if any)  
{  
    declarations  
    statements  
}
```




Functions, prototypes K&R, 1.7, p24

c/powers.c

```
1 #include <stdio.h>
2
3 int power(int m, int n);          /* function prototype */
4
5 /* main(): test power function */
6 int main(void) {
7     int i;
8     for (i=0; i<6; ++i) {
9         printf("%2d %6d %6d\n", i, power(2,i), power(-3,i));
10    }
11    return 0;
12 }
13
14 /* power: raise base to the n-th power; n >= 0 */
15 int power(int base, int n) {
16     int i,p;
17     p = 1;
18     for (i=1; i<=n; ++i) {
19         p = p*base;
20     }
21     return p;
22 }
```



Output:

0	1	1
1	2	-3
2	4	9
3	8	-27
4	16	81
5	32	-243

- Functions definitions need to be given before the function can be called.
- Instead of the full function definition, we can provide a *function prototype* which provides
 - ▷ the name of the function (`powers`)
 - ▷ the type of the function parameters (`int, int`)
 - ▷ the type of the return value (`int`).

The actual function definition can follow then at any point in the file.

It is convention to have the `main` function as the first (or last) function definition in the file.

Some terminology: function's parameters and arguments

c/snippet-only/powers_func.c

```
1 /* power: raise base to the n-th power; n >= 0 */
2 int power(int base, int n) {
3     int i,p;
4     p = 1;
5     for (i=1; i<=n; ++i) {
6         p = p*base;
7     }
8     return p;
9 }
```

- K&R refer to the as the function's *parameters* when they talk about the *type* of the variables "base" and "n".
- K&R like to refer to "base" and "n" as the function's *arguments* when talking about the actual values (that are passed to "base" and "n" in this case when the function is called).

Other names used: formal arguments (for parameters), actual argument(for arguments), or – less specific – input arguments, arguments, input.



Arguments – call by value

Arguments in C are passed *by value*: the called function is given the values of its arguments in *temporary variables* rather than the originals.

→ function can safely modify argument values without causing side effects

BUT, this is not true for (i) arrays and (ii) pointers.

Example:

c/passbyvalue.c

```
1 #include <stdio.h>
2
3 void f(int k){ /* f: print k dots in line*/
4     while (k>0) {
5         printf(".");
6         --k;
7     }
```



```
8   printf("\n");
9 } /* k is zero when function is left */
10
11 int main(void) {
12     int k = 10;
13     printf("k=%d\n",k);
14     f(k);
15     printf("k=%d\n",k);
16     return 0;
17 }
```

c/passbyvalue.out

```
1 k=10
2 .....
3 k=10
```

ooo

→ lab 4

stdin, stdout and stderr



- Unix(/Linux/Mac OS X) and (to some degrees) Windows have three default file descriptors:
 - ▷ 0: *standard input* (stdin)
 - ▷ 1: *standard output* (stdout)
 - ▷ 2: *standard error* (stderr)
- stdin is (by default) connected to the keyboard
- stdout and stderr are connected to the display
- Can use redirection (“>” and “<”) to change this



stdout and stderr redirection

- A command such as `ls` (Unix) or `dir` (Windows) will list the files in the current working directory (by sending the output to stdout).
- We can direct the stdout output (which has file descriptor 1) to a file:
`dir 1> filelist.txt`
- Error messages will be sent to stderr and do not go into this file. To capture those, we need to redirect stderr (with file descriptor 2):
`dir *.ccc 1>filelist.txt 2>filelist.error`
- Instead of `1>` can simply use `>`
- Common notation: `dir *.ccc >dir.stdout 2>dir.stderr`

Character Input and Output

(K&R, Sec 1.5, page 15)

- Command `c=getchar()` reads the *next input character* from the standard input text stream and returns it
- Command `putchar(c)` writes the `c` to the standard output text stream.
- EOF is a special macro representing End Of File (Linux: use CTRL+d on the keyboard to create this, Windows command: use CTRL+z (may have to be at beginning of new line, followed by RETURN)):

Often EOF = -1, but implementation dependent. Must be a value that is not a valid value for any possible character.

For this reason, `c` is of type `int` (not `char` as one may have expected).

EOF is defined in `<stdio.h>`.

File copying

(K&R, Sec 1.5.1, page 16)

c/cat_long.c

```
1 #include <stdio.h>
2
3 /* main: copy stdin to stdout */
4 int main(void) {
5     int c;
6     c=getchar();
7     while (c != EOF) { /* != means 'not equal' */
8         putchar(c);
9         c=getchar();
10    }
11    return 0;
12 }
```

Example usage (making copy of cat.c to copy_of_cat.c):

```
fangohr$ ./cat_long <cat_long.c >copy_of_cat_long.c
```

Experienced C programmers will shorten `cat_long.c` to:

`c/cat.c`

```
1 #include <stdio.h>
2
3 /* main: copy stdin to stdout */
4 int main(void) {
5     int c;
6     while ((c=getchar()) != EOF) {
7         putchar(c);
8     }
9     return 0;
10 }
```

Counting Characters in file

(K&R, Sec 1.5.2, page 17)

c/charcount.c

```
1 #include <stdio.h>
2
3 /* main: count number of characters in stdin */
4 int main(void) {
5     long nc=0;          /* Number of Characters */
6     while ( getchar() != EOF){
7         nc++;
8     }
9     printf("%ld\n", nc);
10    return 0;
11 }
```

```
$ ./charcount < charcount.c
```

```
215
```

```
$ wc -m charcount.c
```

```
215 charcount.c
```

Alternative implementation:

c/charcount2.c

```
1 #include <stdio.h>
2
3 /* main: count number of characters in stdin */
4 int main(void) {
5     double nc;          /* Number of Characters */
6     for (nc=0; getchar() != EOF; nc++){
7         ;              /* 'null statement' */
8     }
9     printf("%.0f\n", nc);
10    return 0;
11 }
```

- use for instead of while (just because we can)
- use double instead of long (can count higher)
- null statement (“;”) required as for-loop must have a body

Line counting

(K&R, Sec 1.5.3, page 19)

c/linecount.c

```
1 #include <stdio.h>
2
3 /* main: count number of lines in stdin */
4 int main(void) {
5     long nl=0;          /* Number of Lines */
6     int c;
7     while ( (c=getchar())!=EOF ) {
8         if (c=='\n')
9             nl++;
10    }
11    printf("%ld\n", nl);
12    return 0;
13 }
```

```
$ ./linecount < linecount.c
14
$ wc -l linecount.c
14 linecount.c
```

Word counting

(K&R, Sec 1.5.4, page 20)

c/wordcount.c

```
1 #include <stdio.h>
2
3 #define IN 1 /*inside a word*/
4 #define OUT 0 /*outside a word*/
5
6 /* main: count lines, words and characters in input */
7 int main(void) {
8     int c, nl, nw, nc, state;
9     state = OUT;
10    nl = nw = nc = 0;
11    while ( (c=getchar())!=EOF ) {
12        ++nc;
13        if (c=='\n')
14            ++nl;
```

```
15     if (c==' ' || c=='\n' || c == '\t')
16         state = OUT;
17     else if (state == OUT) {
18         state = IN;
19         ++nw;
20     }
21 }
22 printf("%d %d %d\n", nl, nw, nc);
23 return 0;
24
25 }
```

```
$ ./wordcount < wordcount.c
$ wc < wordcount.c
26  84 474
```


ASCII table

The American Standard Code for Information Interchange (ASCII, pronunciation: ass-kee) is a character-encoding scheme based on the ordering of the English alphabet. ASCII codes represent text in computers, communications equipment, and other devices that use text. (from Wikipedia, <http://en.wikipedia.org/wiki/ASCII>, 6 Nov 2011

C-program that prints ASCII numbers from 32 up to 126.

c/ascii-table.c

```
1 #include <stdio.h>
2 int main(void) {
3     int c;
4     for (c=32; c<127; c++) {
5         printf("%3d : %c\n",c,c);
6     }
7     return 0;
8 }
```

ASCII table output

32 :	52 : 4	72 : H	92 : \	112 : p
33 : !	53 : 5	73 : I	93 :]	113 : q
34 : "	54 : 6	74 : J	94 : ^	114 : r
35 : #	55 : 7	75 : K	95 : _	115 : s
36 : \$	56 : 8	76 : L	96 : `	116 : t
37 : %	57 : 9	77 : M	97 : a	117 : u
38 : &	58 : :	78 : N	98 : b	118 : v
39 : '	59 : ;	79 : O	99 : c	119 : w
40 : (60 : <	80 : P	100 : d	120 : x
41 :)	61 : =	81 : Q	101 : e	121 : y
42 : *	62 : >	82 : R	102 : f	122 : z
43 : +	63 : ?	83 : S	103 : g	123 : {
44 : ,	64 : @	84 : T	104 : h	124 :
45 : -	65 : A	85 : U	105 : i	125 : }
46 : .	66 : B	86 : V	106 : j	126 : ~
47 : /	67 : C	87 : W	107 : k	
48 : 0	68 : D	88 : X	108 : l	
49 : 1	69 : E	89 : Y	109 : m	
50 : 2	70 : F	90 : Z	110 : n	
51 : 3	71 : G	91 : [111 : o	



Bools

- False is represented by the value 0
- True is represented by the value 1
- Any non-zero value evaluates to True
- the `||` operator is a logical OR
- the `&&` operator is a logical AND
- AND has just slightly higher precedence than OR (K&R, p 21)
- Expressions connected by `||` and `&&` are evaluated left to right, evaluation stops as soon as truth or falsehood is known

Arrays

(K&R, Sec 1.6, page 22)

c/digitcount.c

```
1 #include <stdio.h>
2
3 /* main: count digits, white space, others */
4 int main(void) {
5     int c, i, nwhite, nother;
6     int ndigit[10];
7     nwhite = nother = 0;
8     for (i = 0; i < 10; ++i) {
9         ndigit[i] = 0;
10    }
11
12    while ( (c=getchar())!=EOF) {
13        if (c >= '0' && c <= '9') {
14            ++ndigit[c-'0'];
15        }
```

```

16     else {
17         if (c==' ' || c=='\n' || c=='\t')
18             ++nwhite;
19         else
20             ++nother;
21     }
22 }
23
24 printf("digits =");
25 for (i=0; i<10; ++i) {
26     printf(" %d", ndigit[i]);
27 }
28 printf(", white space = %d, other = %d\n",
29         nwhite, nother);
30
31 return 0;
32 }

```

```

$ ./digitcount < digitcount.c
digits = 10 3 0 0 0 0 0 0 0 1, white space = 168, other = 374

```



Arrays

- Can declare array variable (here `ndigit`) providing an array with space for n elements of a particular (and fixed) type: `int ndigit[10]` (here $n = 10$).
- Values can be accessed by indexing `ndigit[i]` using square brackets
- First element has index 0 and last element has index $n - 1$.
- Must remember the size of the array (`ndigit` cannot tell us)
- Must avoid read/write to elements outside the array (i.e. < 0 or $\geq n - 1$)
- Must initialise elements (may carry random values)



Strings (=character arrays)

- Strings (“chunks of text”) are represented through arrays of char
- Can declare char array s of length 10 as:

```
char s[10];
```

- s is actually a pointer to the beginning of the array.
- Convention in C: put special value `\0` in array to indicate end of string (called the *null character*, K&R, p30). For example “hello\n” is stored as an array of char with 7 entries:

```
-----  
| 0 | 1 | 2 | 3 | 4 | 5 | 6 |  
-----  
| h | e | l | l | o | \n | \0 |  
-----
```

c/helloarray0.c

```
1 #include <stdio.h>
2
3 /* demonstrate string termination I */
4 int main( void ) {
5     char a[10];
6     a[0]='h'; a[1]='e'; a[2]='l'; a[3]='l';
7     a[4]='o'; a[5]='\n'; a[6]='\0';
8     printf("%s",a);
9     return 0;
10 }
```

hello

Allocate array of appropriate size automatically from initial value:

c/helloarray.c

```
1 #include <stdio.h>
2
3 /* demonstrate string termination I */
```




```
4 int main( void ) {
5     char s []="hello\n";
6     int i;
7     printf ("%s",s);
8     for (i=0; i<7; i++)
9         printf ("s [%d]='%c '=%3d\n",i,s [i],s [i]);
10    return 0;
11 }
```

hello

s [0]='h '=104

s [1]='e '=101

s [2]='l '=108

s [3]='l '=108

s [4]='o '=111

s [5]='

'= 10

s [6]=' '= 0



Manual setting of \0:

c/stringtest.c

```
1 #include <stdio.h>
2
3 int main(void) {
4     char s[27];
5     int i;
6     for (i=0; i<26; i++) {/* put abc...xyz in s */
7         s[i] = 'A'+i;
8     }
9     s[26]=' \0'; /* terminate string */
10    printf("s=%s\n",s);
11    s[10]=' \0'; /* terminate string earlier*/
12    printf("s=%s\n",s);
13    return 0;
14 }
```

s=ABCDEFGHIJKLMNOPQRSTUVWXYZ

s=ABCDEFGHIJ

Character Arrays: print longest line

K&R, 1.9, page 28

Write a program that reads a set of text lines and prints the longest. Strategy:

```
while (there is another line)
    if (it is longer than the previous longest)
        save it
        save its length
print longest line
```

- clear strategy divides problem in subproblems
- write code to reflect strategy: use functions
 - ▷ `getline()` to fetch the next line of input and
 - ▷ `copy(to,from)` to save line

plus main program putting it all together.

c/longestline.c

```
1 #include <stdio.h>
2 #define MAXLINE 1000
3
4 /* function prototypes */
5 int get_line(char line[], int maxline);
6 void copy(char to[], char from[]);
7
8 /* main: print longest input line */
9 int main(void) {
10     int len;           /* current line length      */
11     int max;          /* maximum length seen so far */
12     char line[MAXLINE]; /* current input line        */
13     char longest[MAXLINE]; /* longest line saved here   */
14
15     max = 0;
16     while ((len = get_line(line, MAXLINE)) > 0) {
17         if (len > max) {
18             max = len;
19             copy(longest, line);
20         }
21     }
22     if (max > 0) { /* there was a line */
23         printf("%s", longest);
24     }
25     return 0;
```

```

26 }
27
28 /* get_line: read a line into s, return length */
29 int get_line(char s[], int lim) {
30     int c, i;
31
32     for (i=0; i<lim-1 && (c=getchar())!=EOF && c!='\n'; ++i) {
33         s[i]=c;
34     }
35     if (c == '\n') {
36         s[i] = c;
37         ++i;
38     }
39     s[i] = '\0';
40     return i;
41 }
42
43 /* copy: copy 'from' to 'to': assume to is big enough */
44 void copy(char to[], char from[]) {
45     int i=0;
46     while ((to[i]=from[i])!='\0') {
47         i++;
48     }
49 }

```

- `getline()` communicates through
 - ▷ (i) return value (length of line), 0 means EOF (because an empty line contains one character: `'\n'`).
 - ▷ (ii) parameter `char s[]`: `s` is pointer to location in memory that keeps array
- `copy()` communicates through parameters
- `char c[1000]` reserves space for 1000 characters.

ooo

→ **lab 5,6**



Common coding Gotchas in C

“Gotcha” from “Got you”

1. missing semicolon
2. missing closing curly brace
3. using uninitialised variables
4. unexpected integer division truncation
5. format specifier in printf-statement of wrong data type
6. missing parenthesis around logical expression in if-statement
(. . . there are more . . .)



Finding errors(I): compiling fails

1. read the error message carefully:
 - it could point to the error (the ideal case)
 - it could point close to the point of error
2. check code for missing semicolons and non-matching curly-braces
3. “reduce” your program by commenting out parts of the code (using `/*` and `*/`) until the remaining program compiles.
4. Then
 - remove the comments (line by line) and
 - re-compile until compilation fails

That’s where the bug is hiding.



Finding errors(II): execution fails

Typical failure modes:

- program never stops (infinite loop)
- program computes the wrong answer
- execution aborts with error message (or computer crash)

General advice:

- Introduce printf-statements in the code that by reading these (while the program executes), you can deduce what happens.

(There are more sophisticated tools to de-bug programs, so-called “debuggers”.)

sizeof

- The “sizeof(t)” operator returns the number of bytes² used to store a variable of type “t”.
- The return value of sizeof is of type size_t (which must be an unsigned integer type).
- Need “sizeof” to allocate memory dynamically.
- Can use to study the platform we are working on using a program like

c/sizeof.c

```
1 #include <stdio.h>  
2
```

²strictly the number of chars which is the basic unit in C, but effectively all important architectures use char=byte today.

```

3 int main(void) {
4     printf("sizeof(char)    = %d\n", (int) sizeof(char));
5     printf("sizeof(short)   = %d\n", (int) sizeof(short int));
6     printf("sizeof(int)     = %d\n", (int) sizeof(int));
7     printf("sizeof(long)    = %d\n", (int) sizeof(long int));
8     printf("sizeof(float)   = %d\n", (int) sizeof(float));
9     printf("sizeof(double)  = %d\n", (int) sizeof(double));
10    printf("sizeof(double*) = %d\n", (int) sizeof(double*));
11    printf("sizeof(char*)   = %d\n", (int) sizeof(char*));
12    printf("sizeof(FILE*)   = %d\n", (int) sizeof(FILE*));
13
14    return 0;
15 }

```

Output on a 64-bit Mac OS X 10.6 on Intel Core 2 Duo (MacBook) and

Output on a 64-bit Ubuntu Linux 10.6 on Intel Intel i7 CPU Nehalem (PC)

c/sizeof.out64bit

```

1 sizeof(char)    = 1
2 sizeof(short)  = 2
3 sizeof(int)    = 4
4 sizeof(long)   = 8
5 sizeof(float)  = 4
6 sizeof(double) = 8

```

```
7 sizeof(double*) = 8
8 sizeof(char*)   = 8
9 sizeof(FILE*)   = 8
```

Output on a 32-bit Debian 5.0 on AMD64 CPU (PC)

c/sizeof.out32bit

```
1 sizeof(char) = 1
2 sizeof(short) = 2
3 sizeof(int) = 4
4 sizeof(long) = 4
5 sizeof(float) = 4
6 sizeof(double) = 8
7 sizeof(double*) = 4
8 sizeof(char*) = 4
9 sizeof(FILE*) = 4
```



Pointers

- Pointer is variable that contains address of a variable (K&R, chap 5, p 93).
- `&` is the *address operator*: for “char c;” and “p” being a pointer to char:

```
p = &c;
```

assigns the address of “c” to the pointer “p”. We say that “p” *points to* “c”.

- To *declare* a pointer of type “char”, we need to use

```
char *p;
```

- Outside declarations and outside function parameter lists, “*” is the *indirection* or *redirection operator*: when applied to a pointer, it accesses the object the pointer points to.



c/pointer1.c

```
1 #include <stdio.h>
2 int main(void) {
3     int x=1, y=2;
4     int *pi;           /* pi is pointer to int */
5     pi = &x;          /* pi now points to x */
6     y = *pi;          /* y is now 1 */
7     *pi = 0;          /* x is now 0 */
8     pi = &y;          /* pi points to y */
9     printf("x=%d, y=%d \n",x,y);
10    printf("address x=%p, address y=%p \n",
11           (void *) &x, (void *) &y );
12    printf("pi=%p, *pi=%d \n", (void *) pi, *pi);
13    return 0;}
```

x=0, y=1

address x=0x7fff5fbff02c, address y=0x7fff5fbff028

pi=0x7fff5fbff028, *pi=1

More pointer rules

- The declaration of a pointer “ip” to int “i”:

```
int *ip;
```

is intended as a mnemonic; it says that the expression “*ip” is an int.

The syntax of the declaration for a variable mimics the syntax of expressions in which the variable might appear (see function “swap”).

- If ip points to the integer i, then *ip can occur in any context where i could, for example *ip = *ip + 10, or y=*ip+1 or *ip+=1 which is equivalent to ++*ip and also (*ip)++.
- unary operators * and & bind more strongly than arithmetic operators
- unary operators associate from right to the left, i.e. *ip++ \Leftrightarrow *(ip++).

Example: swap, K&R, 5.2, p95

- C passes arguments to functions by value, i.e. by providing a local copy of the argument value.
- Thus, functions can not alter a variable in the calling function. Example:

c/pointer_swap1.c

```
1 #include <stdio.h>
2 void swap( int x, int y) { int tmp=x; x=y; y=tmp;}
3
4 int main(void) {
5     int a=1, b=2;
6     printf("Before swap a=%d, b=%d\n",a,b);
7     swap(a,b);
8     printf("After swap a=%d, b=%d\n",a,b);
9     return 0;}
```

Before swap a=1, b=2

After swap a=1, b=2

Solution: pass pointers of values to be swapped to function:

c/pointer_swap2.c

```
1 #include <stdio.h>
2 void swap(int *px, int *py)
3 {int tmp=*px; *px=*py; *py=tmp;}
4
5 int main(void) {
6     int a=1, b=2;
7     printf("Before swap a=%d, b=%d\n",a,b);
8     swap(&a,&b);
9     printf("After swap a=%d, b=%d\n",a,b);
10 return 0;}
```

Before swap a=1, b=2

After swap a=2, b=1

⇒ pointer arguments enable a function to access and change objects in the function that called it.

scanf() again

- We now know enough to understand exactly what is happening when using the scanf() function seen earlier. The argument that goes in to scanf after the format specifier is meant to be a pointer so that any input is assigned to the memory address of the appropriate variable.



Pointers and arrays K&R, 5.3, p97

- In C, there is a strong relationship between pointers and arrays
- Any operation that can be achieved using subscripting can also be done with pointers.
- The pointer version will generally be faster but, at least to the uninitiated, somewhat harder to understand.
- The declaration

```
int a[10];
```

defines an array `a` of size 10, that is, a block of 10 consecutive objects named `a[0]`, `a[1]`, ..., `a[9]`. Each object requires `sizeof(int)` bytes storage.

- With the declaration



```
int *pa;          /* pointer to int */
```

the assignment `pa = &a[0]` sets `pa` to point to `a[0]`.

- but since the array name `a` is a synonym for the location of the initial element, we can use `pa=a` instead of `pa = &a[0]`.
- Using so-called *pointer arithmetic*, one can write
 - ▷ `a+0=a` instead of `&a[0]`
 - ▷ `a+1` instead of `&a[1]`
 - ▷ `a+2` instead of `&a[2]`, i.e. generically
 - ▷ `a+i` instead of `&a[i]` to access *the address of element i*
- Using indirection operator on expressions above, we can access *the element i*:
 - ▷ `*(a+0)=*a` instead of `a[0]`
 - ▷ `*(a+1)` instead of `a[1]`
 - ▷ `*(a+2)` instead of `a[2]`, i.e. generically
 - ▷ `*(a+i)` instead of `a[i]` to access the value of element `i`

Array and pointer example

c/pointer_array.c

```
1 #include <stdio.h>
2 #include <string.h>
3 int main(void) {
4     char c[]="Hello World---";
5     int i=0;
6     for (i=0; i<strlen(c); i++) {/* use array      */
7         printf("%c", c[i]);           /* subscript  */
8     }
9     for (i=0; i<strlen(c); i++) {/* use pointer */
10        printf("%c", *(c+i));         /* arithmetic */
11    }
12 return 0;}
```

Hello World---Hello World---

strcpy example

c/snippet-only/strcpy1.c

```
1 /* strcpy: copy t to s; array subscript version */
2 void strcpy( char *s, char *t) {
3     int i=0;
4     while ((t[i]=s[i]) != '\0')
5         i++;
6 }
7
8 /* strcpy: copy t to s; pointer version */
9 void strcpy( char *s, char *t) {
10    while ((*t=*s) != '\0') {
11        s++;
12        t++;
13    }
14 }
15
16
```

```
17 /* strcpy: copy t to s; pointer version 2 */
18 void strcpy( char *s, char *t) {
19     while ((*t++=*s++) != '\0')
20         ;
21 }
```


Use of pointers

Use of pointers is required when

- allocating memory *dynamically*
- passing functions to functions (K&R, 5.11, p 118)
- Data structures (trees, lists, . . .)

show memory allocation example, trivial

c/mallocint.c

```
1 #include <stdio.h>
2 #include <stdlib.h> /* provides malloc */
3 int main(void) {
4     int *pi;
5     pi = (int *)malloc(sizeof(int));
6     if (pi == NULL) {
7         printf("ERROR: Out of memory\n");
8         return 1;
9     }
10    *pi = 5;
11    printf("%d\n", *pi);
12    free(pi);
13    return 0;
14 }
```

5

The `malloc` function (called in line 5):

- checks that enough memory is available (on heap)
- if not enough memory available, return address zero (`=NULL`) as an error code.
- if memory available, `malloc()` *allocates* or *reserves* the required amount of memory and
- returns a pointer to that reserved block (i.e. the pointer variable then contains the address of that reserved block).
- The return type of `malloc` is pointer to void (i.e. `void *`) so we need to *cast* the pointer to the right type (i.e. `int *`).

Lines 10 and 11 represent the main activity, i.e. use of the pointer `pi` to store an integer.

The `free` function in line 12 marks the reserved memory as released (“freed”).

Questions:

- *Is it really important to check that the pointer is zero after each allocation?*

Yes, should always be checked.

- *What happens if I forget to release a block of memory before the program terminates?*

The operating system will release the memory when the program terminates. Not releasing memory that is unused can lead to *memory leaks*, and is not considered good style either.

- What is the difference from using `malloc` to just declaring an int using `int i;`?

Memory allocation through `malloc` takes place *at run-time*, where as declarations such as `int i` are static and have to be known at compile time.

Memory allocation example for array

c/mallocint2.c

```
1 #include <stdio.h>
2 #include <stdlib.h> /* provides malloc */
3 #define N 10      /* define array length */
4 int main() {
5     int i;
6     int *a; /* array of int */
7     a = (int *)malloc(sizeof(int)*N);
8     if (a == NULL) {
9         printf("ERROR: Out of memory\n");
10        return 1;
11    }
12    else {
13        printf("Have allocated %d bytes for a.\n",
14              (int) sizeof(int)*N);
15    }
16    for (i=0; i<N; a[i++]=0) ; /* initialise */
17    a[4]=42; /* some activity */
18    for (i=0; i<N; printf("%d ",a[i++])) ;
19    free(a);
20    return 0;
21 }
```

Have allocated 40 bytes for a.

```
0 0 0 0 42 0 0 0 0 0
```

Slides starting on [203](#) show how to allocate memory for 2d and 3d arrays.

Typical use of pointers

- allocate memory using malloc and use pointer to point to that memory
- point the pointer to an address at which memory has been reserved already (and is typically in use by another variable), either by assigning an existing pointer to a pointer, or by using the address operator & to obtain the address of a variable.
- Initialise the pointer with NULL (==0). This is a signal (to the programmer) that the pointer does not point to a meaningful position.

Pointers are usually used if the “length” of data is not known at compile time. For example when

- reading a file and we don't know the length of the file

- running a simulation and we don't know the size of the simulation,
- need to use a vector, matrix, tensor where size will vary from run to run

ooo

→ **lab 7,8**

Structures

- It often makes sense to be able to group data together in to a single place so that it can be more easily understood and used.
- For instance, you may be dealing with shapes - all have a colour, a size, a position.
- You may be keeping records of people and each will have a name, a height, weight etc.
- It makes more sense and is often conceptually clearer to encapsulate this data.
- In C, we use the *struct* to do this.

struct declaration example

c/struct1.c

```
1 #include <stdio.h>
2
3 struct person
4 {
5     int age;
6     double height;
7     double weight;
8 };
9
10 int main(void) {
11     struct person BillyBob={34,1.93,85.0};
12     return 0;
13 }
```

struct usage example: name.member

c/struct2.c

```
1 #include <stdio.h>
2
3 struct person
4 {
5     int age;
6     double height;
7     double weight;
8 };
9
10 int main(void) {
11     struct person BillyBob={34,1.93,85.0};
12     printf("The age is %d.\n",BillyBob.age);
13     printf("The height is %.1f.\n",BillyBob.height);
14     printf("The weight is %.1f.\n",BillyBob.weight);
15     return 0;
16 }
```

Passing a struct to a function example

Using structs in functions is entirely the same as any other data type and this is the case with all C syntax - it is possible to have arrays of structs and pointers to structs etc.

c/struct3.c

```
1 #include <stdio.h>
2
3 struct person
4 {
5     int age;
6     double height;
7     double weight;
8 };
9
10 void printDetails(struct person A)
11 {
```

```
12     printf("The age is %d.\n",A.age);
13     printf("The height is %.1f.\n",A.height);
14     printf("The weight is %.1f.\n",A.weight);
15 }
16
17 int main(void)
18 {
19     struct person BillyBob={34,1.93,85.0};
20     printDetails(BillyBob);
21     return 0;
22 }
```



Example: computation of Fibonacci numbers

In Python:

python/fibrec.py

```
def fib(n):
    if n == 0:
        return 0
    elif n == 1:
        return 1
    else:
        return fib(n - 1) + fib(n - 2)

import time
starttime = time.time()
print fib(36)
print time.time() - starttime, "seconds"
```

f(32) took 2.1s (on Lyceum) and f(36) took 15.3s

In Matlab:

matlab/fibrec.m

```
function fibrec()  
  
tic  
disp(fib(36))  
toc  
  
function res=fib(n)  
  
if n==0  
    res=0;  
else  
    if n==1  
        res=1;  
    else  
        res = fib(n-1)+fib(n-2);  
    end  
end  
end
```

Lyceum: Matlab 7.7 (2008)

fib(32) took 72 sec (34 times slower than Python),

fib(36) took 462 sec (30 times slower than Python)



c/fibrec.c

```
1 #include <stdio.h>
2 #include <time.h>
3 /* Compute Fibonacci number n
4    (recursively */
5 int fib(int n) {
6     if (n==0)
7         return 0;
8     else
9         if (n==1)
10            return 1;
11        else
12            return fib(n-1) + fib(n-2);
13 }
14
15 int main(void) {
16     time_t starttime;
17     starttime = time(NULL);
18     printf("fib(36)=%d\n",fib(36));
19     printf("Time taken: %4.2g\n", (double) (time(NULL)-starttime));
20     return 0;
21 }
```

fib(36) took 0.33 sec.



In summary, we find the relative speed *for this example* (all on Lyceum for `fib(36)`)

- Python: 15.3s → factor 45 slower than C
- Matlab: 462s → factor 1400 slower than C

See also slides [155](#) to [168](#) for a more detailed study on execution speed.

Fibonacci faster

c/fibarray.c

```
1 #include <stdio.h>
2 #define N 36
3 long fib(int n) {
4     long fibs[N+1];
5     long i;
6     fibs[0]=0; fibs[1]=1;
7     for (i=2; i<=N; i++) {
8         fibs[i]=fibs[i-1]+fibs[i-2];
9     }
10    return fibs[N];
11 }
12
13 int main(void) {
14     printf("fib(%d)=%ld\n",N,fib(N));
15     return 0;
16 }
```

```
fib(36)=14930352
```

Nice, but requires use of array. In program above, there is a maximum size hard coded (due to the array).

We can get around this with *dynamic memory allocation* (→ laboratory session).

A speed comparison example

Example content: The upper boundary of a two-dimensional geometric shape is described by $y(x) = \sqrt{1 - x^2}$ for $-1 \leq x \leq 1$, and the lower boundary by $y = 0$. If mirrored along the x-axis, this shape would be a circle with radius 1 and area π .

We use numerical integration of $y(x)$ to approximate $\pi/2$. (There are of course other ways of determining π .)

We use the (relatively) simple composite trapezoidal rule with $n = 10,000,000$ subdivisions to integrate $f(x) = y(x) = \sqrt{1 - x^2}$ from $a = -1$ to $b = 1$. (There are more efficient ways of computing such an integral numerically, but here we are more interested in the computational speed than the algorithmic sophistication.)

We will solve the same problem in Matlab, Python and C. They all produce the same answer (for the same number of subdivisions).

Python, conceptually clear

python/pi.py

```
import math
def f(x):
    return math.sqrt(1. - x * x)

def pi(n):
    s = 0.5 * (f(-1) + f(1))
    h = 2. / n
    for i in range(1, n):
        x = -1 + i * h
        s += f(x)
    return s * h * 2

import time
starttime = time.time()
b = pi(10000000)
stoptime = time.time()
print("%.15g, %g" % (b, math.pi - b))
print math.pi
print("Needed %g seconds" % (stoptime - starttime))
```

Output:

3.14159265348482, 1.04977e-10

3.14159265359

Needed 10.8633 seconds

Execution time:

Lyceum: 14.5s

MacBook Pro: 11s

Python, function $f(x)$ inline

Reducing the overhead of calling the function can make code faster.

python/pi_inlined.py

```
import math
def pi(n):
    s = 0
    h = 2. / n
    for i in range(1, n):
        x = -1 + i * h
        s += math.sqrt(1. - x * x)
    return s * h * 2

import time
starttime = time.time()
b = pi(10000000)
stoptime = time.time()
print("%.15g, %g" % (b, math.pi - b))
print("%.15g" % math.pi)
print("Needed %g seconds" % (stoptime - starttime))
```

Output:

3.14159265348482, 1.04977e-10
3.14159265358979
Needed 6.681 seconds

Lyceum: 8.2s

MacBook Pro: 6.6s

Python, using numpy

Delegate some of the operations to numpy (which uses compiled code).

python/pi_array.py

```
import time
import math
import numpy

def pi(n):
    x = 1. - numpy.linspace(0, 2, n + 1)
    y = numpy.sum(numpy.sqrt(1 - numpy.power(x, 2)))
    s = numpy.sum(y)
    return s * 4. / n

starttime = time.time()
b = pi(10000000)
stoptime = time.time()
print("%.15g, %g" % (b, math.pi - b))
print math.pi
print("Needed %g seconds" % (stoptime - starttime))
```

Output:

3.14159265348482, 1.04976e-10

3.14159265359

Needed 1.14756 seconds

MacBook Pro: 1.14s

Lyceum: 1.36s

Mixing C and Python



A common strategy is to replace those (and only those) parts of a program which consume most of the CPU time by fast (i.e. compiled) code. There are a number of Python tools supporting the mixing of C and Python (including Cython, Pyrex, Swig, Boost [in particular for C++], Weave). Here we use the `weave.inline` tool that comes with `scipy`:

`python/pi_weaveinline.py`

```
import time, math
from scipy import weave          # this is a Python program
def pi(n):
    my_C_code = """
        double s=0;          /* this is C code */
        double h=2./n;      /* in a string */
        double x;
        long i;
        for (i=1; i<n; i++) {
            x = -1+i*h;
            s += sqrt(1.-x*x);
        }
        return_val = s*h*2; /* last line C code */
    """
```

```
    return weave.inline(my_C_code, ['n'], compiler='gcc')

if __name__ == "__main__": # main program (Python)
    starttime = time.time()
    b = pi(10000000)
    stoptime = time.time()
    print("%.15g, %g" % (b, math.pi - b))
    print math.pi
    print("Needed %g seconds" % (stoptime - starttime))
```

Output on Lyceum:

```
[fangohr@lyceum ~/tmp]$ python pi_weaveinline.py
<weave: compiling>
creating /tmp/fangohr/python24_intermediate/compiler_c9100a4644ba0ca4d2d2a1b35e817f54
3.14159265348482, 1.04977e-10
3.14159265359
Needed 1.57684 seconds
[fangohr@lyceum ~/tmp]$ python pi_weaveinline.py
3.14159265348482, 1.04977e-10
3.14159265359
Needed 0.205193 seconds
```

Note that this is slow the first time we execute it (as the c-code string is compiled into a C program) but very fast the second time. (MacBook Pro: 0.17s)

C

[Code missing] → lab exercise

Measurements:

Lyceum: 0.42s (0.20s if using -O3)

MacBook Pro: 0.27s (0.16 if using -O3)

Matlab

matlab/pi.m

```
function pi()  
  
tic;  
    disp(piEuler(10000000));  
toc;  
  
function res=f(x);  
    res=sqrt(1.-x*x);  
  
function s=piEuler(n)  
s = 0.5*(f(-1)+f(1));  
h=2./n;  
for i=1:n-1  
    x = -1+i*h;  
    s = s+f(x);  
end  
s = 2*s*h;
```

Lyceum: 64s

MacBoo Pro: (Unknown)

Matlab, exploiting arrays

matlab/pi_array.m

```
tic
n=10000000;
x=-1:2./(n+1):1;
y=sqrt(1-x.^2);
s=sum(y);
disp(s*4./n);
toc
```

Lyceum: 0.21s

MacBoo Pro: (Unknown)

Speed example, summary

Use Lyceum numbers:

Language	Seconds	Factor
=====		
C (-O3)	0.2	1.00
Python weave	0.21	1.05
Matlab array	0.21	1.05
C	0.4	2.00
Python array	1.3	6.50
Python inline	8.2	41.00
Python plain	14.5	72.50
Matlab plain	64.0	320.00

Need interpretation warning on next slide.

Which one is fastest?

Actual performance will depend on

- Problem to solve, and implementation
- Hardware
- Version of OS, Compilers, Interpreters
- Other activity on the system

Example on previous slide shows that Matlab can be very fast if used well.

Not all performance possibilities have been exploited.

Cython — a modern approach to effective simulation code development

From <http://www.cython.org>: Cython is a language that makes writing C extensions for the Python language as easy as Python itself.

Cython gives you the combined power of Python and C to let you

- write Python code that calls back and forth from and to C or C++ code natively at any point.
- easily tune readable Python code into plain C performance by adding static type declarations.
- integrate natively with existing code and data from legacy, low-level or high-performance libraries and applications.

- interact efficiently with large data sets, e.g. using multi-dimensional NumPy arrays.
- use combined source code level debugging to find bugs in your Python, Cython and C code.
- quickly build your applications within the large, mature and widely used CPython ecosystem.

Another summary: *Cython is Python with C data types.*

Making Python code fast using Cython

- Central idea is to *automatically* translate Python code into C code (through the *Cython compiler*, and to compile the resulting C code using a (normal) C-compiler).
- This provides a modest performance gain over executing the Python code with the Python interpreter.
- We can improve the performance significantly if we can tell the Cython compiler which data types Python objects will be. For numeric code, this is often straightforward.

This is best demonstrated through an example:

- We stick to our previous example `python/pi.py`, and gather the following files in the direction `bench` for benchmarking performances:

bench/pi.py

```
import math

def f(x):
    return math.sqrt(1. - x * x)

def pi(n):
    s = 0.5 * (f(-1) + f(1))
    h = 2. / n
    for i in xrange(1, n):
        x = -1 + i * h
        s += f(x)
    return s * h * 2
```

- We have replaced `range` with `xrange` for clarity (`xrange` does not actually create the list as it is a generator, and thus needs much less memory).

The command to translate the python file `pi.py` into C is:

```
cython pi.py
```

- For clarity of the following steps, we copy our file `bench/pi.py` and save it under the new name `bench/cy_simple.pyx`, and execute

```
cython cy_simple.pyx
```

Cython files (i.e. the files that we will compile to C code using the Cython compiler) have by convention the extension `pyx`.

In this particular case, the code is both valid Python and valid Cython.

- The resulting file `cy_simple.c` is long (> 2000 lines) and difficult to read (show only beginning):

`bench/cy_simple.c`

```
1  /* Generated by Cython 0.19.1 on Wed Sep 25 15:53:09 2013 */
2
3  #define PY_SSIZE_T_CLEAN
4  #ifndef CYTHON_USE_PYLONG_INTERNALS
5  #ifdef PYLONG_BITS_IN_DIGIT
6  #define CYTHON_USE_PYLONG_INTERNALS 0
7  #else
8  #include "pyconfig.h"
9  #ifdef PYLONG_BITS_IN_DIGIT
10 #define CYTHON_USE_PYLONG_INTERNALS 1
```

```

11 #else
12 #define CYTHON_USE_PYLONG_INTERNALS 0
13 #endif
14 #endif
15 #endif
16 #include "Python.h"
17 #ifndef Py_PYTHON_H
18     #error Python headers needed to compile C extensions, please install development
19 #elif PY_VERSION_HEX < 0x02040000
20     #error Cython requires Python 2.4+.
21 #else
22 #include <stddef.h> /* For offsetof */
23 #ifndef offsetof
24 #define offsetof(type, member) ( (size_t) & ((type*)0) -> member )
25 #endif
26 #if !defined(WIN32) && !defined(MS_WINDOWS)
27     #ifndef __stdcall
28         #define __stdcall
29     #endif
30     #ifndef __cdecl

```

- To actually compile this into executable code which we can use from python, it is easiest to use the `pyximport` module as in the following example:

bench/call_cy_simple.py

```
import time

import pyximport      # this allows to import pyx modules as
pyximport.install()  # if they were normal python modules
                    # (cython will compile them automatically
                    # at that point)

import cy_simple     # access to compiled cython module

print cy_simple.pi(10000) # call pi from compiled code
print cy_simple.__file__  # print the module's file

# Compare cython with plain python version
import pi
print pi.__file__

n = 10 ** 7
t_start = time.time()
value_python = pi.pi(n)
print "Native Python needs %f seconds" % (time.time() - t_start)
t_start = time.time()
value_cython = cy_simple.pi(n)
print "(Simple) Cython needs %f seconds" % (time.time() - t_start)
if value_cython == value_python:
    print("Both compute the same value: %.16f" % value_python)
```


which produces this output (on MacBook Air 1.7GHz Inter Core i5, Sept 2012):

```
3.14158932743
/Users/fangohr/.pyxbld/lib.macosx-10.5-i386-2.7/cy_simple.so
Native Python needs 6.832224 seconds
(Simple) Cython needs 5.180885 seconds
Both compute the same value: 3.1415926534848162
```

Discussion

- the cython version is actually compiled code (because the `cy_simple` module is a shared object file with name ending in `.so`).
- both the cython and the python version compute exactly the same number
- the cython version is slightly faster.

Can we do better?

Annotating the c-code that Cython produced

Cython provides the `-a` switch which stands for Annotate. It produces html output that shows

- the python source on white or yellow background
- line numbers can be double clicked and will reveal the C-code required to represent each line
- the background color represents the closeness of the resulting code to Python (yellow) and C (white).

We can use this to identify which parts of the code are slow (more yellow).

For a given pyx file, for example `cy_simple.pyx`, the command to produce the annotated html output is

```
fangohr$ cython -a cy_simple.pyx -o HTML/cy_simple.html
```

Example:

Generated by Cython 0.16 on Sat Sep 22 15:18:11 2012

Raw output: [cy_simple.html](#)

```
1: import math
2: def f(x):
3:     return math.sqrt(1.-x*x)
4:
5: def pi(n):
6:     s = 0.5*(f(-1)+f(1))
7:     h=2./n
8:     for i in xrange(1,n):
9:         x = -1+i*h
10:         s += f(x)
11:     return s*h*2
```

Step 2: declaring types of variables

- The Cython compiler can not know what object types the variables in the code will be. (Python is a dynamically typed language.)
- The compiled code that is flexible enough to deal with any object type, is fairly slow: running such generic code is essentially what the python interpreter has to do.
- To make the cython-compiled c code faster, we have to restrict the object types to particularly (c) types.

Cython provides keywords that we can use in `pyx` to fix the data types.

We modify `bench/cy_simple.pyx` to become `bench/cy_opt.pyx` and include lines 7 to 9 to fix the type of `s`, `h`, and `x` to be `double`, and variable `i` to be of type `int`:

bench/cy_opt.pyx

```
1 import math
2
3 cdef double f(double x):
4     return math.sqrt(1.0-x*x)
5
6 def pi(int n):
7     cdef:
8         double s, h, x
9         int i
10    s = 0.5*(f(-1)+f(1))
11    h=2./n
12    for i in xrange(1,n):
13        x = -1+i*h
14        s += f(x)
15    return s*h*2
```

We annotate the code again:

Generated by Cython 0.16 on Sat Sep 22 15:45:08 2012

Raw output: [cy_opt.html](#)

```
1: import math
2:
3: cdef double f(double x):
4:     return math.sqrt(1.0-x*x)
5:
6: def pi(int n):
7:     cdef:
8:         double s, h, x
9:         int i
10:    s = 0.5*(f(-1)+f(1))
11:    h=2./n
12:    for i in xrange(1,n):
13:        x = -1+i*h
14:        s += f(x)
15:    return s*h*2
```

and see that

- lines 10, 12, 13, 14 are truly white, and will thus execute at native C speed.
- In line 11, the code is still checking whether ZeroDivision might take place. Cython has a switch to disable this, however, the line is only executed once.

New performance data:

basic python:	6.8 seconds	->	4.0
naive cython:	5.0 seconds	->	2.9
typed cython:	1.7 seconds	->	1.0

A substantial amount of CPU time goes into the calculation of square roots. At the moment, we are using the `sqrt` function from the `mathematics` module. While this is very likely written in C as well, we need to interface from Cython's C code to Python's `math` module to then calculate a single square root using the (compiled C) code in the `math` module.

It would be faster if we could use the `sqrt` function from the C libraries immediately from Cython's C code:

bench/cy_cmath.pyx

```
1 cdef extern from "math.h":
2     double sqrt(double x)
3
4 cdef double f(double x):
5     return sqrt(1.0-x*x)
6
7 def pi(int n):
8     cdef:
9         double s, h, x
10        int i
11        s = 0.5*(f(-1)+f(1))
12        h=2./n
13        for i in xrange(1,n):
14            x = -1+i*h
15            s += f(x)
16        return s*h*2
```


Generated by Cython 0.16 on Sat Sep 22 15:45:08 2012

Raw output: [cy_cmath.html](#)

```
1: cdef extern from "math.h":
2:     double sqrt(double x)
3:
4: cdef inline double f(double x):
5:     return sqrt(1.0-x*x)
6:
7: def pi(int n):
8:     cdef:
9:         double s, h, x
10:        int i
11:        s = 0.5*(f(-1)+f(1))
12:        h=2./n
13:        for i in xrange(1,n):
14:            x = -1+i*h
15:            s += f(x)
16:        return s*h*2
```

The performance figures are now

basic python:	6.841452 seconds	->	43.04261	(pi.py}
naive cython:	5.201624 seconds	->	32.72572	(cy_simple.pyx}
typed cython:	1.827127 seconds	->	11.49527	(cy_opt.pyx}
cmath cython:	0.158946 seconds	->	1.0	(cy_cmath.pyx}

With relatively little effort (adding a few cdef static type statements, and interfacing of sqrt from C library), we have made our code 40 times faster using Cython.



Overview performance

Finally, we also include the performance of native C code (compiled with different optimisation flags) and the `scipy.weave` code into the following table:³

Method	Time	Factor
=====	=====	=====
basic python:	6.887 seconds	-> 43.88
xrange python:	6.710 seconds	-> 42.75
naive cython:	5.133 seconds	-> 32.70
typed cython:	1.845 seconds	-> 11.75
C (no opt flags):	0.249 seconds	-> 1.58
C (-O3):	0.158 seconds	-> 1.00
C (-O2):	0.158 seconds	-> 1.00
cmath cython:	0.157 seconds	-> 1.00
weave python:	0.157 seconds	-> 1.00

³The script `bench/benchmark.py` produces this table. These numbers are produced on Mac Book Air, Sept 2012, Mac OS X 10.8.1, 1.7 GHz Intel Core i5.

- Optimised Cython code can run as fast as optimised C code.
- (Note: the order of the last four lines can vary from run to run – they are essentially of the same speed.)



How can we write even faster code?

- With compiled C code, we have pretty much done as much as we can to make our code fast.

(Not entirely true: SESG6028 will teach more details on serial optimisation.)

- To gain further speed-up, we need to parallelise using
 - ▷ multiple cores per CPU
 - ▷ multiple CPUs per board
 - ▷ multiple machines connected
 - ▷ GPU cards (that host hundreds of cores)
 - ▷ emerging hybrid architectures



GPUs

- GPUs are programmed in specialised languages. Well known are CUDA (Nvidia-specific) and OpenCL (generic).
- GPUs have many computational cores but not a lot of GPU RAM (typically 2-4GB max at the moment)
- cores are fast, but communication to CPU RAM is very slow
- if problem fits into GPU RAM, speed ups of factor 10 to 100 can be achieved

GPU computing example: OpenCL code

Timings taken on a Ubuntu Linux machine with a GeForce GTX 560 Ti GPU with 256 cores, and Intel Core2 CPU 6420 @ 2.13GHz 2, using double floats on GPU ($n = 10^7$):

xrange python:	5.6759	seconds	->	897.104
basic python:	5.6001	seconds	->	885.130
naive cython:	4.3711	seconds	->	690.881
typed cython:	1.7268	seconds	->	272.923
C (no opt flags):	0.4600	seconds	->	72.705
C (-O2):	0.3700	seconds	->	58.480
cmath cython:	0.2744	seconds	->	43.364
weave python:	0.2685	seconds	->	42.435
C (-O3):	0.2600	seconds	->	41.094
PyOpenCL inc transfer:	0.0072	seconds	->	1.137
PyOpenCL exc transfer:	0.0063	seconds	->	1.000

- we can gain another factor ≈ 50 if we run the calculation on the GPU.
- The card used here (GeForce GTX 560 Ti GPU) costs about £150 (summer 2012).
- If we use single precision floats on the GPU (instead of double), the execution time is reduced by another factor ≈ 2 .

Acknowledgements: Justin Diviney has written and provided the OpenCL code (bench/OCL.py) used in these timings, and helped with the Cython and benchmarking programme.

PyOpenCL example

- There is a very convenient tool called PyOpenCL that helps with writing OpenCL code: only the parts ("kernels") of the calculation that need be fast are written in OpenCL; the surrounding glue is in Python. For example:

bench/pyopencl-example.py

```
import numpy
import pyopencl as cl

# create host side vectors
A_host = numpy.array([2, 4, 6, 5, 0, 8, 3], dtype = numpy.float32)
B_host = numpy.array([1, 1, 0, 5, 6, 3, 1], dtype = numpy.float32)
C_host = numpy.array([0, 0, 0, 0, 0, 0, 0], dtype = numpy.float32)

# OpenCL C source code for parallel addition of two vectors A and B
kernel = """
__kernel void vector_add(__global float* A,
                        __global float* B,
                        __global float* C)
{
```

```

    // each kernel instance has a different global id
    int i = get_global_id(0);
    C[i] = A[i] + B[i];
}
"""

# OpenCL routines and objects
context = cl.create_some_context()
queue = cl.CommandQueue(context)
program = cl.Program(context, kernel).build()

# create device side vectors and copy values from host to device memory
A_dev = cl.Buffer(context, cl.mem_flags.COPY_HOST_PTR, hostbuf = A_host)
B_dev = cl.Buffer(context, cl.mem_flags.COPY_HOST_PTR, hostbuf = B_host)
C_dev = cl.Buffer(context, cl.mem_flags.COPY_HOST_PTR, hostbuf = C_host)

# run the kernel (see string sourceCode above)
program.vector_add(queue, A_host.shape, None, A_dev, B_dev, C_dev)

# enqueue data transfer from device to host memory
cl.enqueue_read_buffer(queue, C_dev, C_host).wait()

print(C_host)

# Maxim Skripnik, Hans Fangohr, Southampton 2012

```

when run produces the output

```
fangohr@osiris:$ python pyopencl-example.py  
[ 3.    5.    6.   10.    6.   11.    4.]
```

PyOpenCL homepage: <http://mathematician.de/software/pyopencl>

Tidying up a few loose ends in C programming

- writing to files
- reading from files
- interchangeable nature of `*a` and `a[]`
- parsing command line arguments

Writing to files

c/write_to_file.c

```
1 #include <stdio.h>
2 #include <time.h>
3 int main(void) {
4     FILE *f;          /* pointer to file */
5     if ( (f=fopen("myfile.txt","w"))==NULL) {
6         printf("Cannot open 'myfile.txt' for writing");
7         return -1;
8     }
9     fprintf(f,"We can now print to the file f using Fprintf\n");
10    fprintf(f,"For example, a number %d and %d and %d, and\n",1,2,42);
11
12    if (fclose(f) != 0 ) {
13        fprintf(f,"File could not be closed.\n");
14        return -1;
15    }
16    return 0;
17 }
```

- Line 5: attempt to open a file for Writing
- if successful, `fopen` returns a pointer different from `NULL`

- line 7: return -1 to OS to signal failure if we couldn't open file
- line 12: `fclose` returns 0 if successful

Note: There are other ways of writing data to a stream, including `fputc` to write a character.

Reading from files

c/read_from_file.c

```
1 #include<stdio.h>
2 #include<stdlib.h>
3 int main(void) {
4     FILE *f;                /* pointer to file */
5     char c;
6     if ( (f=fopen("myfile.txt","r"))==NULL) {
7         printf("Cannot open 'myfile.txt' for reading");
8         return -1;
9     }
10    /* read file content char by char and print to stdout */
11    while ( (c=fgetc(f)) != EOF) {
12        printf("%c",c);
13    }
14    if (fclose(f) != 0 ) {
15        printf("File could not be closed.\n");
16        return -1;
17    }
18    return 0;
19 }
```

Note: There are other ways of reading data, including `fscanf`.

a* and a[]

Can use *a and a[] interchangeably in function parameter definitions:

c/array-and-pointer.c

```
1 #include <stdio.h>
2
3 void print_array1( int a[]) {
4     int i;
5     for (i=0; i<5; i++) {
6         printf("%d ",a[i]);
7     }
8     printf("\n");
9 }
10
11 void print_array2( int *a) {
12     int i;
13     for (i=0; i<5; i++) {
14         printf("%d ",a[i]);
15     }
16     printf("\n");
17 }
18
19 int main(void) {
20     int a[5];
```



```
21  int i;  
22  for (i=0;i<5;i++) a[i]=i*i;  
23  print_array1(a);  
24  print_array2(a);  
25  return 0;  
26 }
```

Output:

```
0 1 4 9 16  
0 1 4 9 16
```

Reading command line arguments

The `main()` function in general takes two input arguments: an `int argc` which provides the total number of arguments, and `char *argv[]`, an array of pointers to strings (i.e. to an array of chars).

c/args.c

```
1 #include <stdio.h>
2 int main(int argc, char *argv[]) {
3     int i;
4     for (i=0; i<argc; i++) {
5         printf ("Argument %d = '%s'\n",
6                 i, argv[i]);
7     }
8     return 0;
9 }
```

```
$ ./args This is a test
Argument 0 = './args'
Argument 1 = 'This'
```

```
Argument 2 = 'is '  
Argument 3 = 'a '  
Argument 4 = 'test '
```

Dynamic memory allocation for 1d, 2d and 3d arrays

The following code has kindly been provided by Dr Ian Bush, NAG <http://www.nag.co.uk/about/ibush.asp> to demonstrate how to allocate memory for 1d, 2d and 3d arrays so that it is arranged *contiguously* in memory.

The following test program demonstrates how to use the routines in `c/array_alloc.c`.

First, we compile `array_alloc.c` into an object file (`array_alloc.o`)

```
gcc -ansi -pedantic -Wall -c -o array_alloc.o array_alloc.c
```

and then we compile the executable `array_alloc_demo` using this object file `array_alloc.o`

```
gcc -ansi -pedantic -Wall \  
    -o array_alloc_demo array_alloc.o array_alloc_demo.c
```

c/array_alloc_demo.c

```
1  /* Simple example program to show use of the array_alloc routines */
2  /* Ian Bush, 2010 for SESG6025 */
3
4  #include <stdio.h>
5  #include <stdlib.h>
6
7  #include "array_alloc.h"
8
9  #define NROWS 3
10 #define NCOLS 5
11
12 int main( void ) {
13
14     float **array_float;
15     char **array_char;
16
17     double ***t;
18
19     int i, j, k;
20
21     /* First we will deal with a two dimensional array of floats */
22
23     /* Allocate a NROWS by NCOLS array */
24     array_float = alloc_2d_float( NROWS, NCOLS );
25
```

```

26  /* Check the allocation worked */
27  if( array_float == NULL )
28      /* It failed! Exit the program */
29      exit( EXIT_FAILURE );
30
31  /* Zero the array - this shows how to access the elements */
32  for( j = 0; j < NROWS; j++ )
33      for( i = 0; i < NCOLS; i++ )
34          array_float[ j ][ i ] = 0.0;
35
36  /* Print out the adress of each element of array
37     Note how the 4 byte floats are contiguous in memory */
38  printf("Addresses of the elements of the %i by %i two dimensional array:\n",
39         NROWS, NCOLS );
40  for( j = 0; j < NROWS; j++ )
41      for( i = 0; i < NCOLS; i++ )
42          printf( "a[ %i ][ %i ] is at %lu\n", j, i,
43                 ( unsigned long ) &( array_float[ j ][ i ] ) );
44  printf( "Note that the %i byte floats are contiguous in memory\n",
45         (int) sizeof( **array_float ) );
46  printf( "\n" );
47
48  /* Now do exactly the same, but this time with chars */
49
50  /* Allocate a NROWS by NCOLS array */
51  array_char = alloc_2d_char( NROWS, NCOLS );

```

```

52
53  /* Check the allocation worked */
54  if( array_char == NULL )
55      /* It failed! Exit the program */
56      exit( EXIT_FAILURE );
57
58  /* Zero the array - this shows how to access the elements */
59  for( j = 0; j < NROWS; j++ )
60      for( i = 0; i < NCOLS; i++ )
61          array_char[ j ][ i ] = 0.0;
62
63  /* Print out the adress of each element of array
64     Note how the 1 byte chars are contiguous in memory */
65  printf("Addresses of the elements of the %i by %i two dimensional array:\n",
66         NROWS, NCOLS );
67  for( j = 0; j < NROWS; j++ )
68      for( i = 0; i < NCOLS; i++ )
69          printf( "a[ %i ][ %i ] is at %lu\n", j, i,
70                 ( unsigned long ) &( array_char[ j ][ i ] ) );
71  printf( "Note that the %i byte chars are contiguous in memory\n",
72         (int) sizeof( **array_char ) );
73  printf( "\n" );
74
75
76  /* Finally a 3 dimensional example */
77  t = alloc_3d_double( 2, 3, 4 );

```

```

78  if( t == NULL )
79      exit( EXIT_FAILURE );
80  printf( "Addresses of the elements ");
81  printf( "of the %i by %i by %i three dimensional array of doubles:\n",
82          2, 3, 4 );
83  for( k = 0; k < 2; k++ )
84      for( j = 0; j < 3; j++ )
85          for( i = 0; i < 4; i++ )
86              printf( "a[ %i ][ %i ][ %i ] is at %lu\n", k, j, i,
87                      ( unsigned long ) &( t[ k ][ j ][ i ] ) );
88
89  return EXIT_SUCCESS;
90  }

```

The output from the program reads:

Addresses of the elements of the 3 by 5 two dimensional array:

```

a[ 0 ][ 0 ] is at 4296016032
a[ 0 ][ 1 ] is at 4296016036
a[ 0 ][ 2 ] is at 4296016040
a[ 0 ][ 3 ] is at 4296016044
a[ 0 ][ 4 ] is at 4296016048
a[ 1 ][ 0 ] is at 4296016052
a[ 1 ][ 1 ] is at 4296016056
a[ 1 ][ 2 ] is at 4296016060
a[ 1 ][ 3 ] is at 4296016064

```



```
a[ 1 ] [ 4 ] is at 4296016068
a[ 2 ] [ 0 ] is at 4296016072
a[ 2 ] [ 1 ] is at 4296016076
a[ 2 ] [ 2 ] is at 4296016080
a[ 2 ] [ 3 ] is at 4296016084
a[ 2 ] [ 4 ] is at 4296016088
```

Note that the 4 byte floats are contiguous in memory

Addresses of the elements of the 3 by 5 two dimensional array:

```
a[ 0 ] [ 0 ] is at 4296016128
a[ 0 ] [ 1 ] is at 4296016129
a[ 0 ] [ 2 ] is at 4296016130
a[ 0 ] [ 3 ] is at 4296016131
a[ 0 ] [ 4 ] is at 4296016132
a[ 1 ] [ 0 ] is at 4296016133
a[ 1 ] [ 1 ] is at 4296016134
a[ 1 ] [ 2 ] is at 4296016135
a[ 1 ] [ 3 ] is at 4296016136
a[ 1 ] [ 4 ] is at 4296016137
a[ 2 ] [ 0 ] is at 4296016138
a[ 2 ] [ 1 ] is at 4296016139
a[ 2 ] [ 2 ] is at 4296016140
a[ 2 ] [ 3 ] is at 4296016141
a[ 2 ] [ 4 ] is at 4296016142
```

Note that the 1 byte chars are contiguous in memory

Addresses of the elements of the 2 by 3 by 4 three dimensional array of doubles:

```
a[ 0 ][ 0 ][ 0 ] is at 4296016144
a[ 0 ][ 0 ][ 1 ] is at 4296016152
a[ 0 ][ 0 ][ 2 ] is at 4296016160
a[ 0 ][ 0 ][ 3 ] is at 4296016168
a[ 0 ][ 1 ][ 0 ] is at 4296016176
a[ 0 ][ 1 ][ 1 ] is at 4296016184
a[ 0 ][ 1 ][ 2 ] is at 4296016192
a[ 0 ][ 1 ][ 3 ] is at 4296016200
a[ 0 ][ 2 ][ 0 ] is at 4296016208
a[ 0 ][ 2 ][ 1 ] is at 4296016216
a[ 0 ][ 2 ][ 2 ] is at 4296016224
a[ 0 ][ 2 ][ 3 ] is at 4296016232
a[ 1 ][ 0 ][ 0 ] is at 4296016240
a[ 1 ][ 0 ][ 1 ] is at 4296016248
a[ 1 ][ 0 ][ 2 ] is at 4296016256
a[ 1 ][ 0 ][ 3 ] is at 4296016264
a[ 1 ][ 1 ][ 0 ] is at 4296016272
a[ 1 ][ 1 ][ 1 ] is at 4296016280
a[ 1 ][ 1 ][ 2 ] is at 4296016288
a[ 1 ][ 1 ][ 3 ] is at 4296016296
a[ 1 ][ 2 ][ 0 ] is at 4296016304
a[ 1 ][ 2 ][ 1 ] is at 4296016312
a[ 1 ][ 2 ][ 2 ] is at 4296016320
a[ 1 ][ 2 ][ 3 ] is at 4296016328
```

Array allocation tools header file

c/array_alloc.h

```
1
2
3 /* 1d routines */
4 char *alloc_1d_char ( int ndim1 );
5 int *alloc_1d_int ( int ndim1 );
6 float *alloc_1d_float ( int ndim1 );
7 double *alloc_1d_double( int ndim1 );
8
9 /* 2d routines */
10 char **alloc_2d_char ( int ndim1, int ndim2 );
11 int **alloc_2d_int ( int ndim1, int ndim2 );
12 float **alloc_2d_float ( int ndim1, int ndim2 );
13 double **alloc_2d_double( int ndim1, int ndim2 );
14
15 /* 3d routines */
16 char ***alloc_3d_char ( int ndim1, int ndim2, int ndim3 );
17 int ***alloc_3d_int ( int ndim1, int ndim2, int ndim3 );
18 float ***alloc_3d_float ( int ndim1, int ndim2, int ndim3 );
19 double ***alloc_3d_double( int ndim1, int ndim2, int ndim3 );
```

Array allocation tools implementation

c/array_alloc.c

```
1
2 /*
3  Simple wrappers for allocating arrays so that the elements are
4  contiguous in memory
5
6  The following data types are supported: char, int, float, double
7
8  The following dimensionality of arrays is supported: 1, 2, 3
9
10 The naming convention for routines is
11
12 alloc_<dimensionality>_<type>
13
14 where <dimensionality> is of the form nd, where n is the number of
15 dimensions, and <type> is the data type
16
17 Each function returns a pointer with the appropriate degree of
18 indirection for the dimensionality required. On error NULL is
19 returned.
20
21 The arguments to the functions are the dimensions in natural order.
```

```
22
23   Ian Bush, 2010 for SESG6025, University of Southampton,
24   very minor modifications by Hans Fangohr.
25
26  */
27
28  #include <stdlib.h>
29
30  char *alloc_1d_char( int ndim1 ) {
31     return malloc( ndim1 * sizeof( char ) );
32 }
33
34  int *alloc_1d_int( int ndim1 ) {
35     return malloc( ndim1 * sizeof( int ) );
36 }
37
38  float *alloc_1d_float( int ndim1 ) {
39     return malloc( ndim1 * sizeof( float ) );
40 }
41
42  double *alloc_1d_double( int ndim1 ) {
43     return malloc( ndim1 * sizeof( double ) );
44 }
45
46  char **alloc_2d_char ( int ndim1, int ndim2 ) {
47     char **array2 = malloc( ndim1 * sizeof( char * ) );
```

```

48  int i;
49  if( array2 != NULL ){
50      array2[0] = malloc( ndim1 * ndim2 * sizeof( char ) );
51      if( array2[ 0 ] != NULL ) {
52          for(i = 1; i < ndim1; i++)
53              array2[i] = array2[0] + i * ndim2;
54      }
55      else {
56          free( array2 );
57          array2 = NULL;
58      }
59  }
60  return array2;
61 }
62
63 int **alloc_2d_int( int ndim1, int ndim2 ) {
64     int **array2 = malloc( ndim1 * sizeof( int * ) );
65     int i;
66     if( array2 != NULL ){
67         array2[0] = malloc( ndim1 * ndim2 * sizeof( int ) );
68         if( array2[ 0 ] != NULL ) {
69             for(i = 1; i < ndim1; i++)
70                 array2[i] = array2[0] + i * ndim2;
71         }
72         else {
73             free( array2 );

```

```

74     array2 = NULL;
75     }
76     }
77     return array2;
78 }
79
80 float **alloc_2d_float( int ndim1, int ndim2 ) {
81     float **array2 = malloc( ndim1 * sizeof( float * ) );
82     int i;
83     if( array2 != NULL ){
84         array2[0] = malloc( ndim1 * ndim2 * sizeof( float ) );
85         if( array2[ 0 ] != NULL ) {
86             for(i = 1; i < ndim1; i++)
87                 array2[i] = array2[0] + i * ndim2;
88         }
89         else {
90             free( array2 );
91             array2 = NULL;
92         }
93     }
94     return array2;
95 }
96
97 double **alloc_2d_double( int ndim1, int ndim2 ) {
98     double **array2 = malloc( ndim1 * sizeof( double * ) );
99     int i;

```

```

100  if( array2 != NULL ){
101      array2[0] = malloc( ndim1 * ndim2 * sizeof( double ) );
102      if( array2[ 0 ] != NULL ) {
103          for(i = 1; i < ndim1; i++)
104              array2[i] = array2[0] + i * ndim2;
105      }
106      else {
107          free( array2 );
108          array2 = NULL;
109      }
110  }
111  return array2;
112 }
113
114 char ***alloc_3d_char ( int ndim1, int ndim2, int ndim3 ) {
115     char *space = malloc( ndim1 * ndim2 * ndim3 * sizeof( char ) );
116     char ***array3 = malloc( ndim1 * sizeof( char ** ) );
117     int i, j;
118
119     if( space == NULL || array3 == NULL )
120         return NULL;
121
122     for( j = 0; j < ndim1; j++ ) {
123         array3[ j ] = malloc( ndim2 * sizeof( char * ) );
124         if( array3[ j ] == NULL )
125             return NULL;

```



```

126     for( i = 0; i < ndim2; i++ )
127         array3[ j ][ i ] = space + j * ( ndim3 * ndim2 ) + i * ndim3;
128     }
129     return array3;
130 }
131
132 int     ***alloc_3d_int    ( int ndim1, int ndim2, int ndim3 ) {
133     int     *space = malloc( ndim1 * ndim2 * ndim3 * sizeof( int ) );
134     int     ***array3 = malloc( ndim1 * sizeof( int ** ) );
135     int     i, j;
136
137     if( space == NULL || array3 == NULL )
138         return NULL;
139
140     for( j = 0; j < ndim1; j++ ) {
141         array3[ j ] = malloc( ndim2 * sizeof( int * ) );
142         if( array3[ j ] == NULL )
143             return NULL;
144         for( i = 0; i < ndim2; i++ )
145             array3[ j ][ i ] = space + j * ( ndim3 * ndim2 ) + i * ndim3;
146     }
147     return array3;
148 }
149
150 float   ***alloc_3d_float ( int ndim1, int ndim2, int ndim3 ) {
151     float   *space = malloc( ndim1 * ndim2 * ndim3 * sizeof( float ) );

```

```

152 float ***array3 = malloc( ndim1 * sizeof( float ** ) );
153 int i, j;
154
155 if( space == NULL || array3 == NULL )
156     return NULL;
157
158 for( j = 0; j < ndim1; j++ ) {
159     array3[ j ] = malloc( ndim2 * sizeof( float * ) );
160     if( array3[ j ] == NULL )
161         return NULL;
162     for( i = 0; i < ndim2; i++ )
163         array3[ j ][ i ] = space + j * ( ndim3 * ndim2 ) + i * ndim3;
164 }
165 return array3;
166 }
167
168 double ***alloc_3d_double( int ndim1, int ndim2, int ndim3 ) {
169     double *space = malloc( ndim1 * ndim2 * ndim3 * sizeof( double ) );
170     double ***array3 = malloc( ndim1 * sizeof( double** ) );
171     int i, j;
172
173     if( space == NULL || array3 == NULL )
174         return NULL;
175
176     for( j = 0; j < ndim1; j++ ) {
177         array3[ j ] = malloc( ndim2 * sizeof( double* ) );

```

```
178     if( array3[ j ] == NULL )
179         return NULL;
180     for( i = 0; i < ndim2; i++ )
181         array3[ j ][ i ] = space + j * ( ndim3 * ndim2 ) + i * ndim3;
182     }
183     return array3;
184 }
```

```
changeset: 171:903387372b50
tag:       tip
user:     Hans Fangohr [MBA] <fangohr@soton.ac.uk>
date:     Wed Sep 25 16:10:06 2013 +0100
summary:  Updated table of contents with new page numbers (toc.pdf)
```