

Introduction to the C-Language and Programming Environment

Part 2: The Language

C Language Tutorial
System Programming 251-0053
Winter Semester 2005/06

René Müller, IFW B45.2

Content

- Variables
- Input/Output
- Expressions
- Statements
 - Assignment, if-else, while, do-while, for, switch
- Subroutines
- Pointers
 - Function Pointers
- Arrays
- Enumeration Types
- Structs and Unions

General Remark



Caution: The C Language is case-sensitive.

`main` \neq `Main`

`myVariable` \neq `MyVariable`

Data Types in C

Name	Description	Size and Format
<code>short</code>	„short integer“	platform dependent
<code>int</code>	Integer	platform dependent
<code>long int</code> (<code>long</code>)	„long integer“	platform dependent
<code>long long int</code>	„very long integer“	platform dependent
<code>float</code>	single precision floating-point number	32 bit IEEE 754
<code>double</code>	double precision floating-point number	64 bit IEEE 754
<code>char</code>	single character	8 bit

Caution: The size of a an integer type is platform dependent. The size of a type can be determined at runtime with the `sizeof` operator. It returns the size in units of chars.

e.g.: `sizeof(long long int) ⇒ 8 (byte)`

Size of Data Types

Size of `int` types for different platforms (in bytes):

Type	Linux x86	MacOS PowerPC	Tru64 Unix DEC Alpha	SUN Solaris sparcv9 (64 bit)	LEGO Mindstorm RCX Brick Hitachi H8
<code>short</code>	2	2	2	2	2
<code>int</code>	4	4	4	4	2
<code>long int (long)</code>	4	4	8	8	4
<code>long long int</code>	8	8	8	8	4

Be aware of different sizes when porting C applications to different platforms.

Signed and Unsigned Types

In C you can specify for `int` types whether their value is to be interpreted as a `signed` (2-complement) or as an `unsigned` value.

Thus:

- a) `signed` for positive and negative values
- b) `unsigned` for positive values only

Type	Size (in bytes on x86)	Range
<code>signed char</code>	1	-128 .. +127
<code>unsigned char</code>	1	0 .. 255
<hr/>		
<code>signed short</code>	2	-32'768 .. +32'767
<code>unsigned short</code>	2	0 .. 65'535
<hr/>		
<code>signed int</code>	4	-2'147'483'648 .. + 2'147'483'647
<code>unsigned int</code>	4	0 .. 4'294'967'295
<hr/>		
<code>signed long long int</code>	8	-9'223'372'036'854'775'808 .. +9'223'372'036'854'775'807
<code>unsigned long long int</code>	8	0 .. 18'446'744'073'709'551'615

Unless `unsigned` is explicitly specified the compiler uses `signed`.

Declaration of Variables

Variables need to be **declared (with type)** before they can be used. Note that variables must be declared at the beginning of a code block (unless C99 standard is used).

	<i>Type</i>	<i>name of variable</i>
e.g.:	<code>int</code>	<code>a;</code>
	<code>double</code>	<code>f;</code>
	<code>boolean</code>	<code>error;</code>

Variable names must be valid identifiers, i.e., keywords such as `void`, `int`, `struct` must not be used as identifiers.

Literals

An **initial value** can be assigned when the variable is declared.

```
Example:  int anInt = 4;  
          double pi = 3.141592654;
```

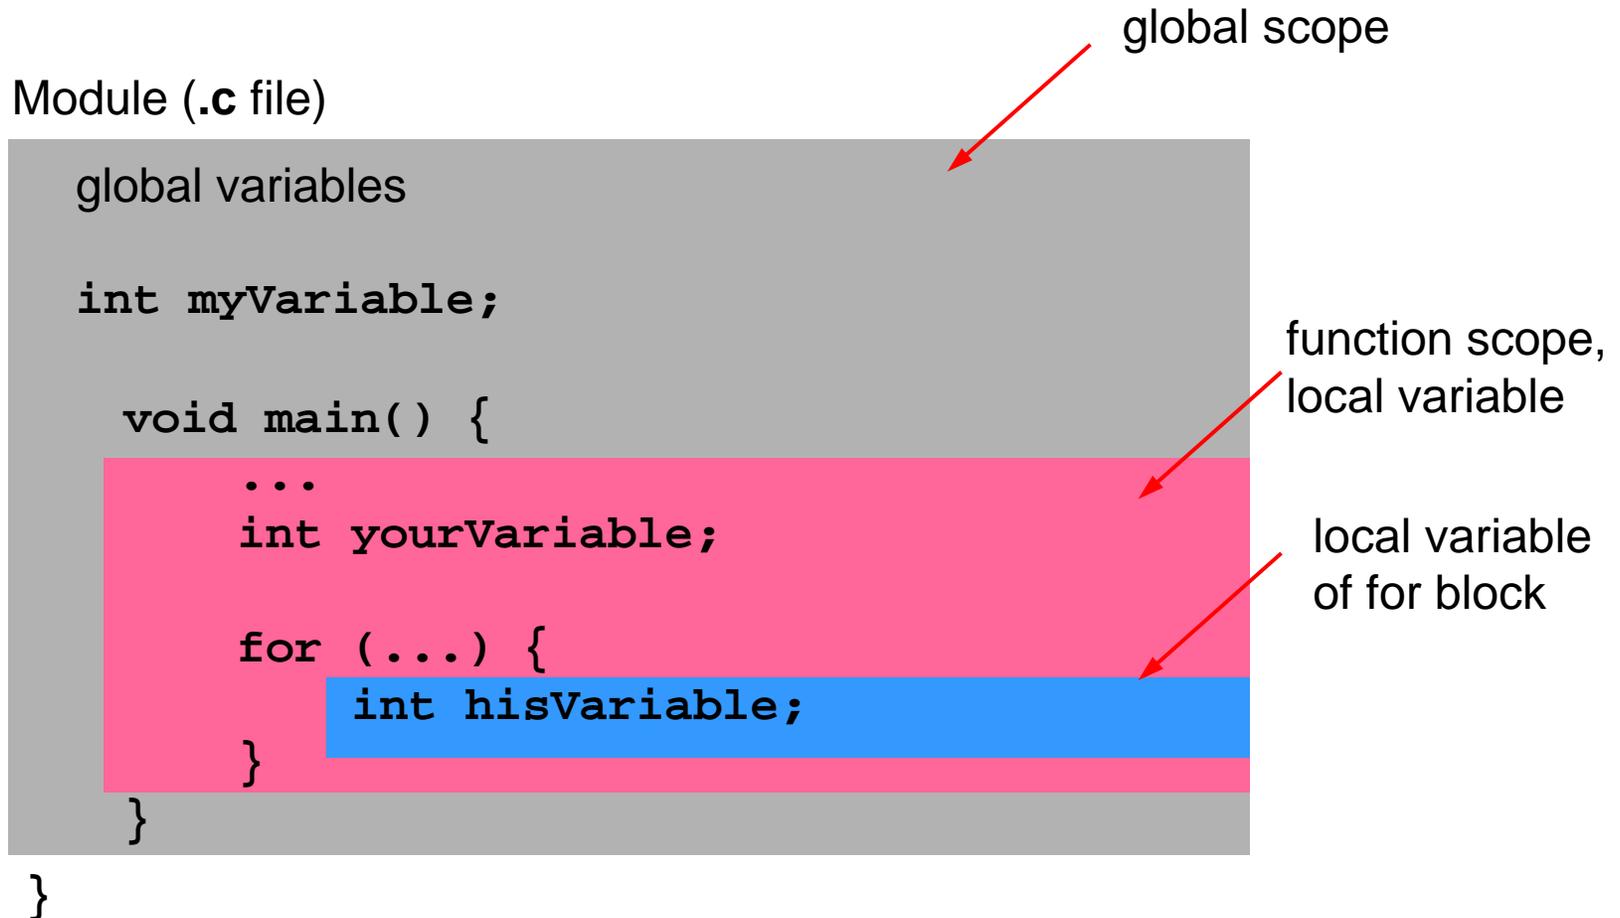
Some other literals with their data type:

Literal	Type	
178	int	
8864L	long	
37.266	double	
37.266d	double	
87.383f	float	
26.77e3	double	26.77×10^3
1.25e-1f	float	1.25×10^{-1}
'c'	char	

Note: 8864L is of type long and 8864 is of type int.

Scope of Variables

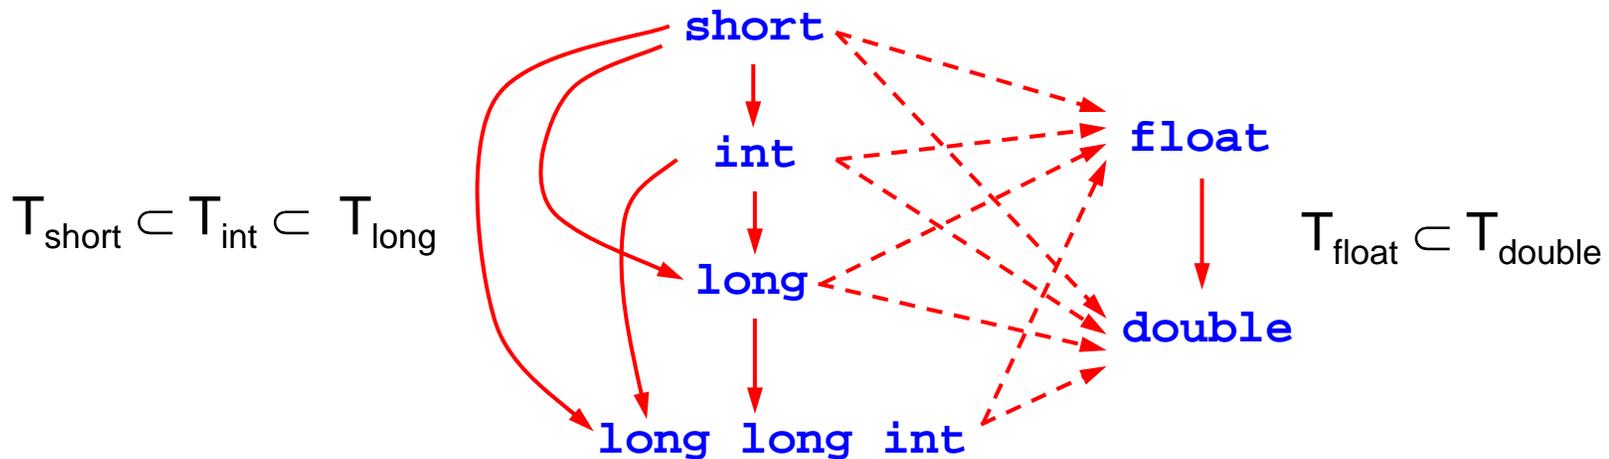
Variables are only valid in the scope they are defined (unless declared static).



Example: `hisVariable` may only be accessed inside of the for-block and `yourVariable` only inside `main`.

Type Conversion

Casting converts between different types. In general type conversion is done implicitly but it can be specified explicitly with the cast operator. In C the following conversions are done implicitly, i.e., this is done automatically.



These conversions guarantee that there is no loss of precision. (Except for conversions an integer to a floating-point type.) Conversions in the other directions may lead to some loss.

Example:

```
double aDouble = 3.14;
```

```
int anInt;
```

```
anInt = aDouble;
```

```
aDouble = anInt;
```

```
⇒ aDouble = 3.000000
```

„loss of precision“ ⚡

Explicit Casting of Types

Explicit type cast by the use of a [cast operator](#).

```
Casting:  valueInNewType = (NewType)valueInOldType;
```

Examples:

```
int anInt;  
double aDouble = 3.14159;
```

```
aDouble = (int)aDouble; // truncation of fraction  
printf("%f\n", aDouble); // --> 3
```

```
anInt = aDouble; // truncation of fraction  
aDouble = anInt/2; // integer division 3/2 = 1  
printf("%f\n", aDouble); // --> 1
```

```
aDouble = (double)anInt/2; // floating-point division 3/2  
printf("%f\n", aDouble); // --> 1.50000
```

```
aDouble = anInt/2.0; // floating-point division 3/2  
printf("%f\n", aDouble); // --> 1.50000
```

Constants

A constant is some place of storage that can hold a value (similar to a variable) but cannot be changed after its initialisation.

In C a constant is declared as a variable with the preceding keyword `const`.

```
...  
const double pi = 3.14159254;  
const int maxval = 127;
```

When trying to assign a new value to a constant after its initialisation the compiler issues (only) a warning message.

```
...  
pi = 5;
```

`gcc` → „warning assignment of read-only variable `pi`“

Output from C

For writing output to the console the function `printf(..)` can be used. It is part of the C library. The prototype of the function is defined in header file `stdio.h`. Thus the use of the `#include <stdio.h>` directive at the beginning is recommended.

```
printf("Hello World!");
```

`printf` requires at least one argument. The first argument must be a string (to be precise, a pointer to the first element of a zero-terminated character array).

Strings must be quoted by “ ” (double quote). Strings may also contain ANSI-C escape sequences.

<code>\n</code>	new line
<code>\t</code>	horizontal tab stop
<code>\"</code>	double quote “
<code>\\</code>	back slash \

e.g.: `printf("Hello World\n\t\"This is a quoted string.\");`

Ausgabe:

Hello World

“This is a quoted string.”

Formatted Output

printf can also be used to print values of different types in various formats (formatted output). The format is specified in the string (the format string) using the special character preceded by %.

Format Character

d, i	int; decimal number
o	int; positive number in octal format
x, X	int; positive number in hexadecimal format
u	int; positive number
c	char; single character
s	A string (char *) until terminator character '\0' is reached
f	double; [-] m.dddddd
e, E	double; [-] m.dddddde±xx (scientific notation)
p	pointer (memory address)
%	prints character '%'

The variables or expressions to be printed are then added as additional arguments after the format string:

```
printf("String %....", arg1, arg2, ... argn)
```

Example – Output from C

Example:

```
#include <stdio.h>
main()
{
    const double pi = 3.141592654;
    int anInt = 42;
    double big = 1.24e23;

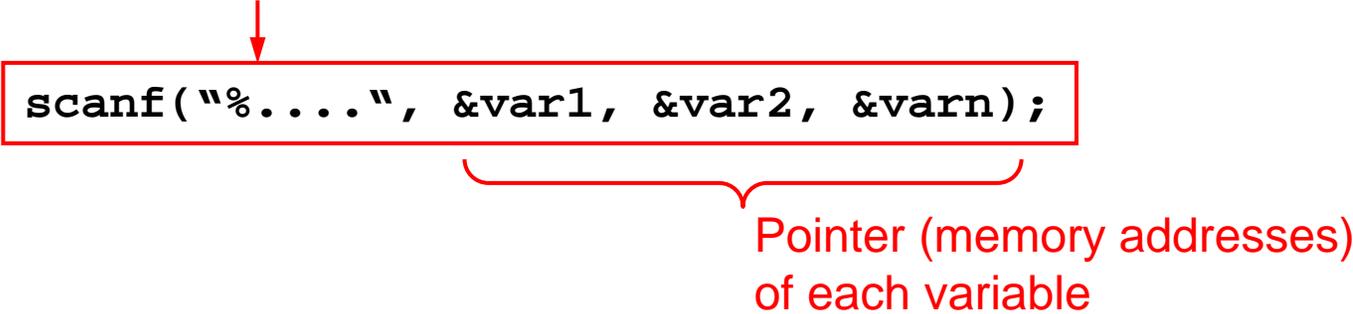
    printf("pi      = %f\n", pi);
    printf("anInt  = %d\n", anInt);
    printf("big    = %e\n", big);
    printf("anInt is at address: %p\n, &anInt);
}
```

Output:

```
pi      = 3.141592
anInt  = 42
big    = 1.240000e+23
anInt is at address: 0xffbefaf4
```

Reading From Standard Input

scanf can read formatted inputs from the standard input (e.g., what the user types) The first argument of **scanf** is also a format string. Next to the format string the addresses of the variables that should store the input are listed.



```
scanf("%....", &var1, &var2, &varn);
```

Pointer (memory addresses)
of each variable

Format character

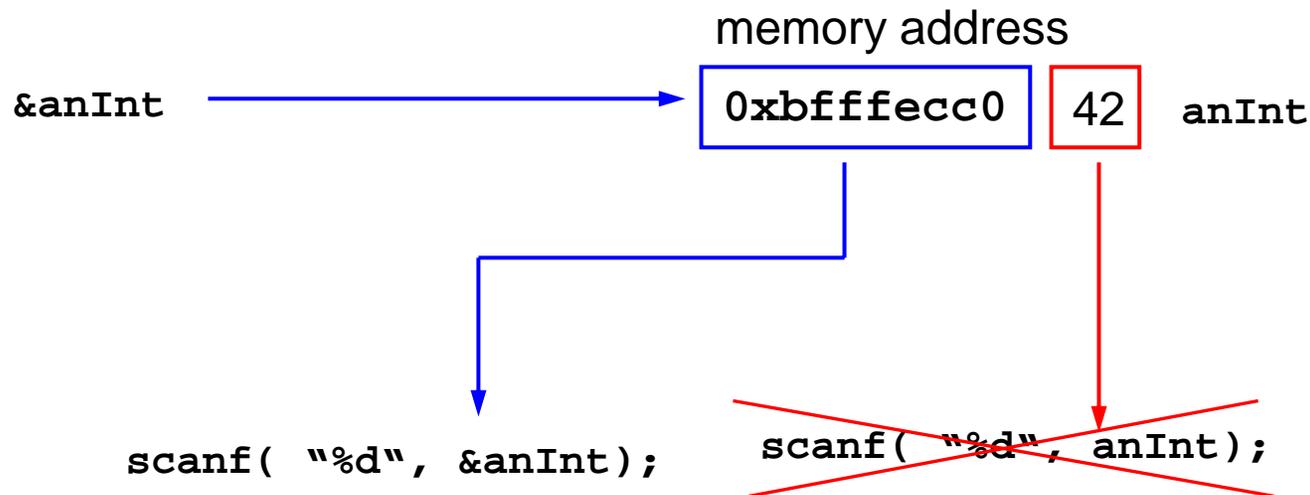
d	int; decimal number
i	int; positive decimal number or octal number (when a leading 0 is found) or hexadecimal number (when prefix 0x is found)
o	interprets input in octal format
x	interprets input in hexadecimal format
u	positive decimal number
c	a single character
s	char*; a string
e,f,g	floating-point number (type float)

Pointers in scanf

(Pointers are discussed later)

`scanf` has to store the value in a variable (memory location). Thus the call `scanf("%d", anInt)` cannot work, since `anInt` specifies the current value of the variable in the call.

`scanf` however requires the address of `anInt`. By using the `&` operator the address of `anInt` is determined.



Reading from Standard Input – Example

```
#include <stdio.h>
main()
{
    int age;
    float size;

    printf("your age: ");
    scanf("%d", &age);
    printf("your size in meter: ");
    scanf("%f", &size);

    printf("You are %d years old and %f m.\n",
           age, size);
}
```

Output:

```
your age:                ← 42
your size in meter:      ← 1.76
You are 42 years old and 1.760000 m.
```

Statements: Assignment

Assignments in C are specified by the **Operator =** .

E.g.:

y = x+1; \Leftrightarrow **y ← x + 1**

Assignments are only valid if:

- left-hand side and right-hand side have same type **or**
- right-hand side type is a subtype of left-hand side type ($T_{LHS} \subset T_{RHS}$)

C also allows multiple-assignment statements:

int c = a = b = x+y;

Variables a, b and c all have the same value x+y.

Arithmetic Expressions

Binary arithmetic operators in C:

+	Addition	$x = 3+5 \Rightarrow 8$
-	Subtraction	$x = 5-3 \Rightarrow 2$
*	Multiplication	$x = 5*3 \Rightarrow 15$
/	Division	$x = 14/3 \Rightarrow 4$ (integer division if both operands int)
%	Modulo	$x = 14\%3 \Rightarrow 2$

Increment und Decrement Operators

In C there is a short hand for $x = x + 1$: statement $x++$ and $x--$ for $x = x - 1$.

Also allowed are $++x$ and $--x$. However note difference between pre- and postfix form:

```
int x = 3;  
int y;
```

```
y = x++;  
printf("%d\n", x); ⇒ 4  
printf("%d\n", y); ⇒ 3
```

```
int x = 3;  
int y;
```

```
y = ++x;  
printf("%d\n", x); ⇒ 4  
printf("%d\n", y); ⇒ 4
```

Differences:

$x++$ returns the current value of x and then increments x

$++x$ increments first and then returns new value of x

Compressed Assignment Operators

Assignments with expression of the form

x = x operand y

can be written in „compressed“ form:

	short form	expanded form
+=	x += y;	x = x + y;
-=	x -= y;	x = x - y;
*=	x *= y;	x = x * y;
/=	x /= y;	x = x / y;
%=	x %= y;	x = x % y;

Boolean Expressions

Note! There is no explicit Boolean type in traditional C (it was introduced in C99 standard). Instead C uses integer values for Booleans, with the following meaning:

false	Value 0 (zero)
true	Any value except 0

Note! C uses „=“ for assignments and for comparisons the equality operator „==“. It returns 1 if operands are equal and 0 otherwise.

e.g.:

```
printf( "%d\n", 1==2);    ⇒ 0  
printf( "%d\n", 1==1);   ⇒ 1
```

Boolean Expressions (2)

Following **comparison operators** lead to expressions of type Boolean:

Operator	Meaning	Example
<code>==</code>	equal	<code>x == 3</code>
<code>!=</code>	not equal	<code>x != y</code>
<code>></code>	greater	<code>4 > 3</code>
<code><</code>	less	<code>x < 3</code>
<code>>=</code>	greater or equal	<code>x >= y</code>
<code><=</code>	less or equal	<code>x <= y</code>

Logical Operators

Boolean expressions can be „logically“ linked.

Operators: AND, OR, XOR, Negation

Operands		AND	OR	XOR	NOT
		$a \wedge b$	$a \vee b$	$a \oplus b$	$\neg a$
a	b	<code>a && b</code>	<code>a b</code>	<code>a ^ b</code>	<code>!a</code>
0	0	0	0	0	1
0	1	0	1	1	1
1	0	0	1	1	0
1	1	1	1	0	0

(bitwise)



Precedence: `!` \Rightarrow `^` \Rightarrow `&&` \Rightarrow `||`

Otherwise use parentheses!

Lazy Evaluation of Boolean Expressions

In C Boolean sub-expressions are only evaluated until the final result can be determined.

Example:

`((expression_1) && (expression_2))`

`expression_2` is only evaluated if `expression_1` is `true`. Otherwise `false` is determined immediately without evaluating the second expression.

`((expression_1) || (expression_2))`

`expression_2` is only evaluated if `expression_1` is `false`. Otherwise `true` is returned without evaluating the second expression.

Illustration:

```
int x = 0 // or 1;
int result = (x == 0) && (++x);
printf("result: %d, x=%d\n", result, x);
```

`x was 0` \Rightarrow `result: 1, x=1`

`x was 1` \Rightarrow `result: 0, x=1`

If-Statement (I)

The syntax of if statement in C is as follows:

```
if (boolean_expression) statement;
```

`statement` is executed if `boolean_expression` evaluates to true, i.e., `!= 0`.

The `else-Block` is optional.

```
if (boolean_expression) statement1; else statement2;
```

`boolean_expression != 0` \Rightarrow `Statement1` is executed

`boolean_expression == 0` \Rightarrow `Statement2` is executed

Statements and Statement-Blocks

C allows grouping of a sequence of statements into a block. Then this **statement-block** is regarded as a **single statement**, e.g., it can be used as then- or else-block.

```
statement_sequence ::= { statement ";" }
statement_block    ::= "{" statement_sequence "}"
statement          ::= statement_block | ...
```

Example:

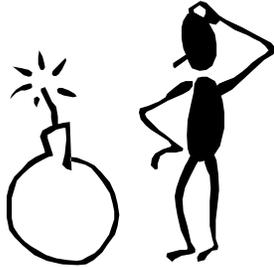
```
if (a>b)
{
    max = a;
    z++;
}
else
    max = b;
```

} statement block in then-block

} single statement in else-block

If-Statement Puzzle

Frequent Bug



```
#include <stdio.h>
main()
{
    int number;
    printf("enter number: ");
    scanf("%d", &number);

    if (number = 0)
        printf("equal 0\n");
    else
        printf("not equal 0\n");
}
```

What is the program output? What is wrong?

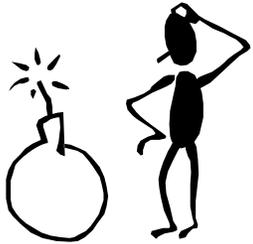
Bug: `number = 0` is an assignment, no comparison!

But it is also valid as a Boolean expression, since C allows multiple assignments `x = (number=0)`. Since `x` is `==0` is false, the else branch is taken, regardless of the value of `number`. But instead `number == 0` was meant.

Warning: The gcc compiler does not warn you here!
(Some other compilers do).

The „Dangling-Else“-Problem

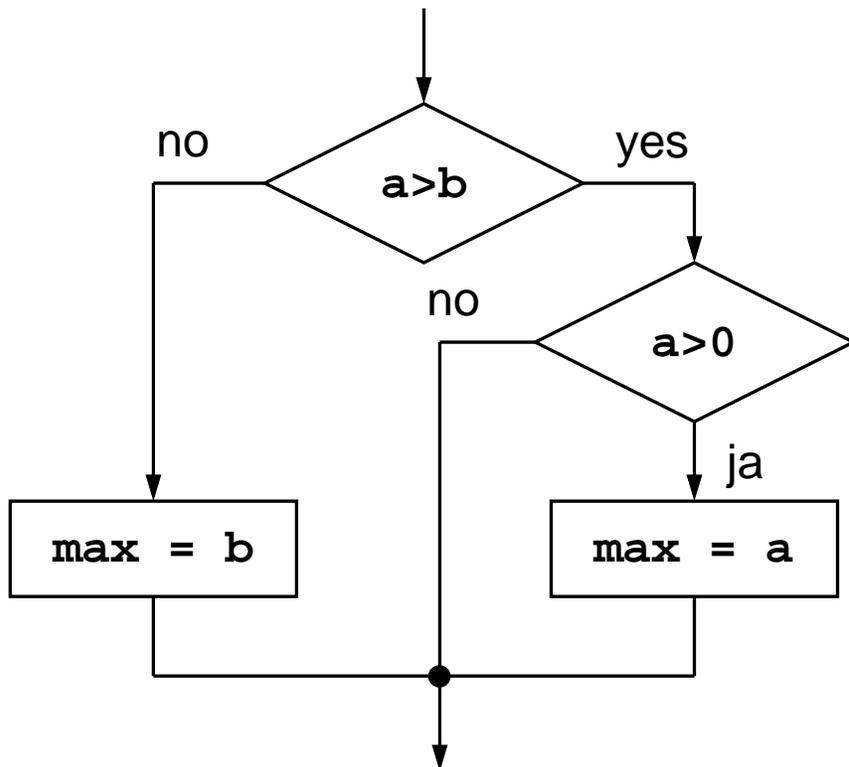
Consider the following code:



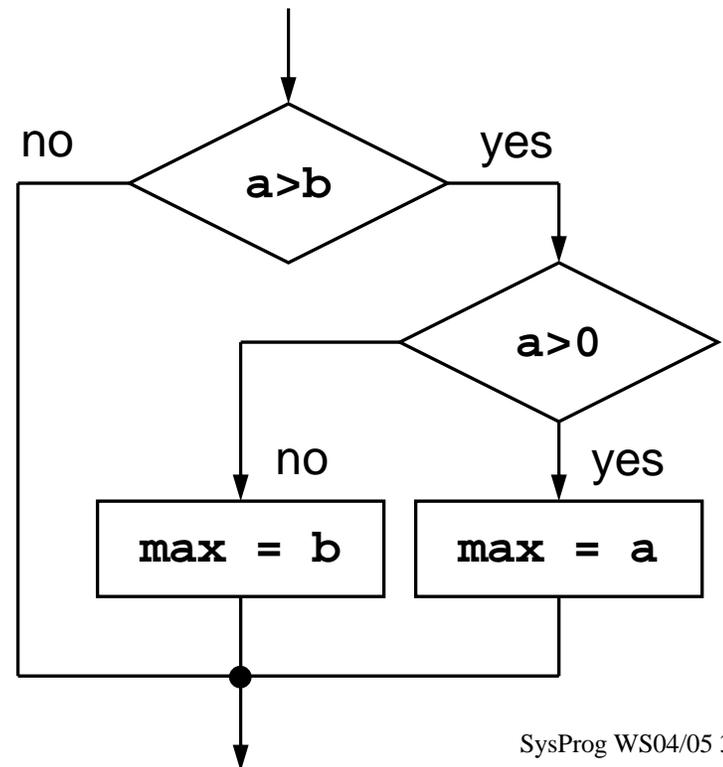
```
if (a>b)
    if (a > 0)  max = a;
else
    max = b;
```

Where does the else belong to?
Note that indentation is irrelevant for the compiler.

Interpretation 1:



Interpretation 2:



The „Dangling-Else“-Problem (2)

This ambiguity (also present in Java) is solved by the compiler by using interpretation 2. The compiler always associates the closest previous else-less if.

For interpretation 1 blocks need to be explicitly used.

```
if (a>b) {  
    if (a > 0) max = 0;  
} else {  
    max = b;  
}
```

If you do not want to run into these troubles always use braces { } .

Conditional Operator

The conditional operator

```
(boolean_expression) ? expression1 : expression2;
```

is a short-hand form of the following if-else statement:

e.g.:

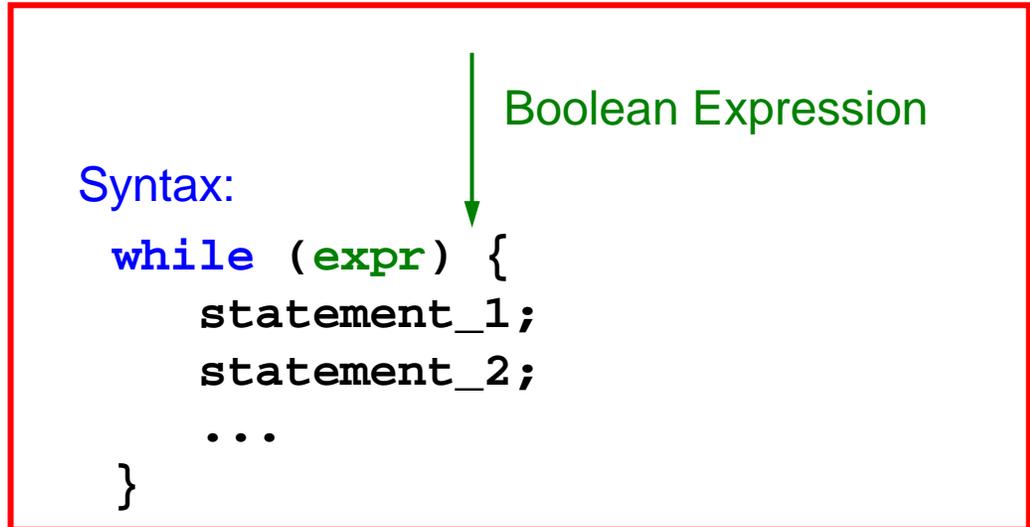
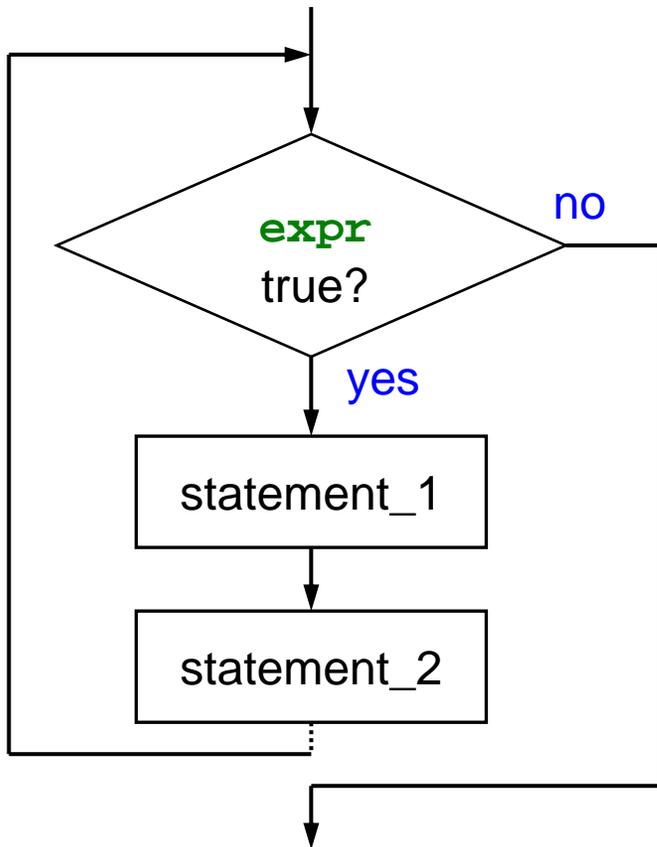
```
if (a>b)
    max = a;
else
    max = b;
```



```
max = (a>b)?a:b;
```

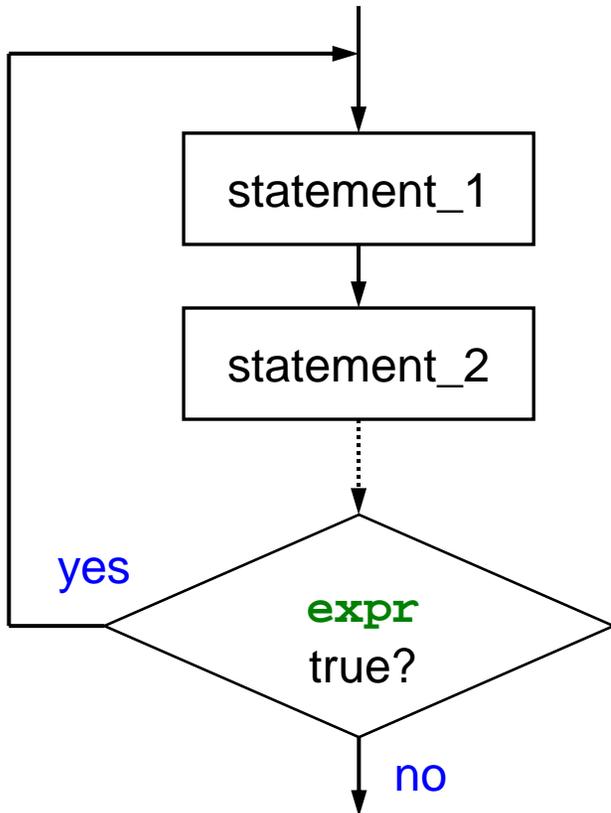
while Loops

- while-loops are executed as long as the expression is true
- the expression is tested before entering the loop



do-while Loops

- do-while loops are executed as long as the expression is true
- the expression is tested at the of the loop body.
- the body do-while loops is executed at least once.



Syntax:

```
do {  
    statement_1;  
    statement_2;  
    ...  
} while (expr);
```

Boolean
Expression

for Loops

For loops consist of the following for components:

- **Initialisation statement:** executed before loop is entered the first time (and before continuation statement is evaluated)
- **Continuation condition:** Evaluated before entering the loop. Loop-body is only executed if the condition evaluates to true.
- **loop statement:** executed after the last statement of the loop body
- **loop body:** statements that are part of the loop

Syntax:

executed before entering the first time	loop while true	execute at end of loop body
↓	↓	↓
<code>for (init_stmt; cont_cond; loop_stmt) {</code>		
<code> statement_1;</code>		
<code> ...</code>		
<code>}</code>		

for Loops (2)

Example:

```
#include <stdio.h>
main()
{
    int i;          // i is loop variable
    for (i=1; i<=10; i++) {
        printf("%d: %d\n", i, i*i);
    }
}
```

- ① initialise i with 1
- ② exit loop if $i > 10$
- ③ increment i after each iteration

Switch-Case Statement

C-Provides a elegant construct for the sequence of if statements (see „251-0053 System Programming“ lecture class06b):

```
int day; // day of week 0=Monday, 1=Tuesday, 2=...
if (day == 0) {
    printf("Monday");
} else if (day == 1) {
    printf("Tuesday");
} else if (day == 2) {
    printf("Wednesday");
} else if (day == 3) {
    printf("Thursday");
} else if (day == 4) {
    printf("Friday");
} else if (day == 5) {
    printf("Saturday");
} else if (day == 6) {
    printf("sunday");
} else {
    printf("invalid day of week ");
}
```

Switch-Case Statement (2)

The Boolean expression of each if statements contains the same variable (expression). `day` has valid values in the range 0..6. In the switch-block there is one case for each possible value. Switch-case might lead to more efficient machine code (if jump tables are used).

```
int day; // day of week 0=Monday, 1=Tuesday, 2=...
switch(day) {
case 0:
    printf("Monday");
    break;
case 1:
    printf("Tuesday");
    break;
...
case 6:
    printf("Sunday");
    break;
default:
    printf("invalid day of week");
}
```

Switch-Case Statement (3)

General syntax:

```
switch (expr) {  
  case value1:  
    statement_sequence1;  
    break;  
  case value2:  
    statement_sequence2;  
    break;  
  case value3:  
    statement_sequence3;  
    return;  
  .  
  .  
  .  
  default:  
    default_statements;  
}
```

The value of **expr** is compared with every literal **value_x** starting from the beginning of the switch-block.

All statements after the first matching case will be executed until **break** or **return** is reached.

The optional **default** case is executed if there was no matching case found.

Note 1: If breaks or returns are omitted we can fall through multiple case statements until either a break/return is found or the end of the switch case statement is reached

Note 2: If the **default** case is omitted the switch statement does nothing if none of the cases were matched.

Example For Switch-Case Statement

Calculator: `c = compute('+', 46, 53);`

```
int compute(char op, int a, int b) {
    switch(op) {
        case '-':
            b *= -1;
        case '+':
            return a+b;
        case '*':
            return a*b;
        case ':':
        case '/':
            return a/b;
        default:
            printf("invalid op %c", c);
            return ERR;
    }
}
```

If `op == '-'` the statement for case `+'` is also executed

return `a/b` for both cases `:'` and `/'`

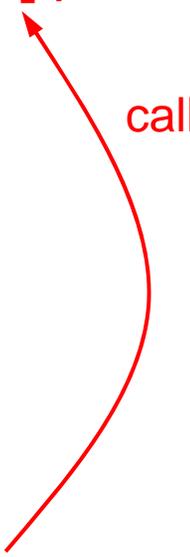
Subroutines in C

Example:

```
#include <stdio.h>

int add(int x, int y)
{
    return x+y;
}

main()
{
    int a, b, c;
    a = 7;
    b = 4;
    c = add(a, b);
    printf("%d+%d=%d\n", a, b, c);
}
```



file: add.c

Subroutine (function) `add` has two `int`-parameters `x`, `y`. It adds these two and returns the sum (`int` type) to the caller.

The function is called on line `c = add(a, b);`. Here the formal parameter `x` and `y` of the function are replaced by the actual parameters (expressions `a` and `b`). The result of the function is stored in variable `c`.

Functions in C

General syntax of a function declaration in C:

```
resultType functionName(parameterType param1, ... )
{
    // function body
    ...
    return exprOfResultType;
}
```

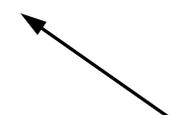
- If `resultType` is not specified the compiler assumes type `int`.
- Return value is set in expression of `return` statement
- Statement `return` immediately leaves function
- Warning: gcc does not complain if there is a flow of execution that does not lead to a `return` statement! Return value is not determined.
- Functions can have zero or more formal arguments
 - Comma-separated list of type-name pairs
 - C also allows open-parameter lists (e.g.: in `printf`)

Procedures and Functions

- Everything is a function
- Procedures are functions that return “nothing”
 - Use keyword **void** instead of return type in function declaration
 - A void function does not need to call **return**, however it may so in case of early abort (expression in return statement is then omitted).

```
void write_int(int nr)
{
    FILE *f;
    if ((f=fopen("myfile.txt","w")) == 0) {
        // error while opening file
        return;
    }
    ...
}
```

“early abort” on error



Functions and Return Types

- If the return type of a function is omitted the compiler automatically assumes `int` return type.

Example:

```
main()  
{  
    printf("Hello World\n");  
}
```

Is identical to:

```
int main()  
{  
    printf("Hello World\n");  
}
```

- Note: In both cases main returns an int. However there is **no return statement!** Value returned at runtime is undefined.

Prototypes in C

Compile the following program:

```
#include <stdio.h>

main()
{
    float a, b, c;
    a = 3.0f;  b = 4.0f;
    c = mult(a,b);
    printf("%f*f=%f\n", a, b, c);
}

float mult(float a, float b)
{
    return a*b;
}
```

file: ctest.c

Compiler error:

```
gcc -o ctest ctest.c

ctest.c:11: warning type
mismatch with previous
implicit declaration

ctest.c:7: warning previous
implicit declaration of `mult`

ctest.c:11: warning: `mult`
was previously implicitly
declared to return `int`
```

Reason: The compiler does not know the result type of function `mult` in `c=mult(a,b)`. It thus **assumes `int`** (implicit declaration). Later on it finds the declaration of `mult` with `float`.

Prototypes in C (2)

Solution 1: Swap functions `main` and `mult` in source code such that the `mult` function is declared before being used in `main`.

```
#include <stdio.h>

float mult(float a, float b)
{
    return a*b;
}

main()
{
    float a, b, c;
    a = 3.0f; b = 4.0f;
    c = mult(a,b);
    printf("%f*f=%f\n", a, b, c);
}
```

Program compiles
without errors.

Prototypes in C (3)

Solution 2: Leave order of implementations unchanged but specify only signature of function `mult` before `main`. Thus the compiler also knows the return type of the function when it is used.

Declaration of function signatures are called prototypes.

```
#include <stdio.h>

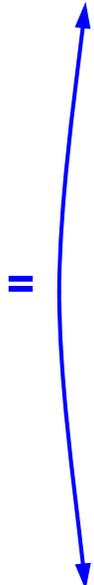
// Prototype of function mult
float mult(float a, float b);

main()
{
    float a, b, c;
    a = 3.0f;  b = 4.0f;
    c = mult(a,b);
    printf("%f*%f=%f\n", a, b, c);
}

float mult(float a, float b)
{
    return a*b;
}
```

Program compiles
without errors.

=



Prototypes in C (4)

The names of the function parameters can be omitted in the prototype declaration. The compiler only needs to know the order and the types.

```
#include <stdio.h>

// Prototype of function mult
float mult(float, float);

main()
{
    float a, b, c;
    a = 3.0f;  b = 4.0f;
    c = mult(a,b);
    printf("%f*%f=%f\n", a, b, c);
}

float mult(float a, float b)
{
    return a*b;
}
```

Program compiles
without errors.

Local and Global Variables

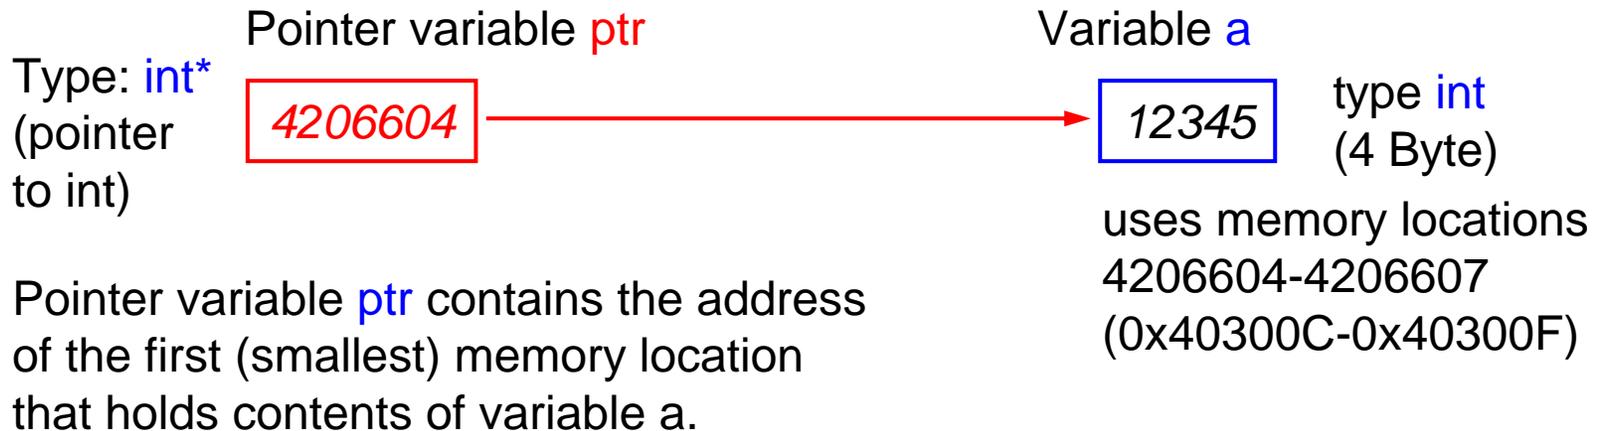
- Variables (defined in a function) local variables can only be accessed in that function.
- In general local variables cannot maintain content between different invocations of the function
 - This is not true for local variables declared `static`

```
double z=5.0;           ← global variable z

void foo()
{
    int x;
    static int y=0;     ← variables x, y local to function foo
                        ← x maintains content between invocations
    ...
    y++;               ← only initialized the first time
}                    ← counts number of invocations
```

Pointer – Definition

Def: A **Pointer** is a variable or expression that refers to some memory location.



The pointer variable itself request memory space too.

For 32-bit wide address spaces (x86) the pointer size is 4 bytes. On a 64-bit system (x86_64, IA64, PPC64) 8 bytes are needed.

Note: casts from a pointer type to `int` are dangerous. The following code works on x86 but not on x86_64 (`int` is only 4 byte on both platforms):

```
int aVar;  
int address = (int)&aVar;
```

`&aVar` retrieves address of `aVar`

Variables, Addresses and Pointers

A simple C program with three local variables:

```
main()
{
    int a = 5;
    int b = 6;
    int* c;

    // c points to a
    c = &a;
}
```

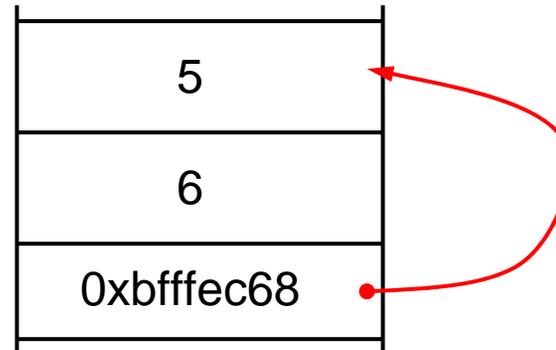
defines `c` as a pointer to a variable of type `int`

`&` address operator retrieves memory location of variable `a`

Variable

↓
a (0xbffec68)
b (0xbffec64)
c (0xbffec60)

Memory



Defining a Pointer in ANSI-C

There are two **identical** styles for declaring a pointer:

<code>int* x;</code>		<code>int *x;</code>
<code>float* y;</code>	or alternatively	<code>float *y;</code>
<code>char* z;</code>	(Asterisk to the variable)	<code>char *z;</code>
USW.		USW.

Note:

- Both types are really the same!
- When a pointer variable is declared space is only allocated for the storage of the pointers content (a memory address) but not for the variable the pointer points to.
- A pointer is bound to the data type that was specified in the definition.
- Use `void*` as a **general pointer type** that can point to any variable/address.
- Every pointer can be cast in a `void*` pointer and vice versa.

Reference and Dereference

Reference & Retrieve the memory address of a variable

```
int a = 6;  
int* c = &a; // &a is the memory location of variable a
```

Dereference * Accessing the variable (content) the pointer points to
(Indirection)

```
int a = 6;  
int* c = &a;
```

```
*c = 7; /* Changes content of variable a by using  
its address stored in pointer c */
```

equivalent to

```
a = 7;
```

Using Pointers for Out-Parameters in Functions

Example: Lets write a function `swap(a, b)` that exchanges the content of the two variables `a` and `b`.

However the following code has a flaw:

```
void swap(int x, int y)
{
    int tmp;
    tmp = x;
    x = y;
    y = tmp;
}
```

When calling `swap(a,b)` the formal parameters `x, y` of the function are replaced by the values of `a` and `b`

The swap-code in the function does indeed exchange the content of `x` and `y`.

But how is the caller affected by this change?

In order to really exchange the content of the variables their addresses must be passed to the function not their values. Call **by-reference** instead of call **by-value** is required.

⇒ Use `int*` instead of `int`.

Using Pointers for Out-Parameter in Functions (2)

Change the programs as follows:

```
void swap(int* x, int* y)
{
    int tmp;
    tmp = *x;
    *x = *y;
    *y = tmp;
}
```

When calling swap use **address operator &** to get the addresses of the two variables.

```
swap(&a, &b);
```

Function Pointers

- A function pointer is a variable that points to a function.
- Definition of a function pointer:

```
returnType (*functPtr)(paramType1, paramType2)
```

- Example: `fptr` is a function pointer that can point to any function that takes an `int` and `float` argument and returns a `float`.

```
float (*fptr)(int, float)
```

- Assigning a function pointer

```
float func(int i, float f) {  
    return i+f;  
}  
...  
fptr = func;           // assignment ptr to function
```

- Calling a function the function pointer points to

```
fptr(42, 3.14f);
```

Function Pointer – Example

- Calculator

```
int add(int x, int y) { return x+y; }
int sub(int x, int y) { return x-y; }
int mul(int x, int y) { return x*y; }
int div(int x, int y) { return x/y; }
```

each operation is
enumerated
0=add, 1=sub, 2=mul,
3=div

```
int evaluate(unsigned int op, int x, int y) {
```

```
    int (*eval[])(int, int) = { add, sub, mul, div };
```

```
    if (op>3) {
        printf("error invalid function");
        return 0;
    }
```

array function pointers
to function with signature
int f(int, int).

```
    return eval[op](x, y);
```

one function for each
enumerated operation.

```
}
```

```
main()
```

```
{
```

```
    printf("%d\n", evaluate(3, 42, 3));
```

```
}
```

Const and Pointers

- In C a pointer can be declared “constant” in two ways
- Constant pointer
 - Syntax: `TYPE * const ptrName = &aTYPEVar;`
 - Pointer variable: **constant** (cannot be overwritten)
 - Data the pointer points to: **not constant**
 - Example

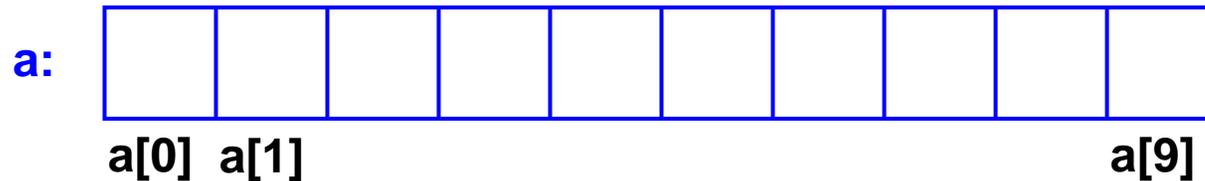
```
int a = 42, b = 42;
int* const ptr = &a;
*ptr = 1; // ok
ptr = &b; // wrong → gcc error: assignment of read-only variable 'ptr'
```
- Pointer to constant data
 - Syntax: `const TYPE * ptrName = &aTYPEVar;`
 - Pointer variable: **not constant**
 - Data the pointer points to: **constant** (cannot be overwritten)
 - Example

```
int a = 42, b = 42;
const int* ptr = &a;
*ptr = 1; // wrong → gcc error: assignment of read-only location
ptr = &b; // ok
```

Declaring Arrays in C

Following declaration defines an array with 10 elements of type `MyType`. The fields can be accessed by an index, `a[0]`, `a[1]`, ..., `a[9]`.

```
MyType a[10];
```



`a[i]` is the *i*th element des Arrays, *i* is called index.

Initialisation of Arrays

An array can be initialised when it is declared.

Example:

```
int primes[] = {2, 3, 5, 7, 11, 13};
```

Creates an array with 6 elements and initialises it with the specified values.

Equivalent code:

```
int primes[6];  
primes[0] = 2;  
primes[1] = 3;  
primes[2] = 5;  
primes[3] = 7;  
primes[4] = 11;  
primes[5] = 13;
```

Variant with Braces is probably better.

Accessing Arrays



If you run the following program you get a strange output

```
#include <stdio.h>

int a[12], b;

main()
{
    int i;
    b = 5;
    printf("b=%d\n", b);

    for (i=0; i<=12; i++) {
        a[i] = i;
    }
    printf("b=%d\n", b);
}
```

Output:

```
b=5
b=12
```

What's wrong here?

fix: use < instead of <= !

The last element written is a[12]. But array was initialised as with length 12, thus last element is a[11]. We are **writing accross the array borders** and therefore into variable b!

Accessing Arrays (2)

In C there is no boundary check when accessing arrays, i.e., the runtime system does not check if index i of $a[i]$ points to a valid array element. It just adds the offset (index) to the array's base address.

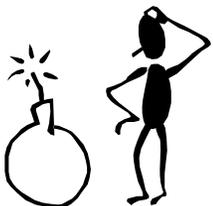
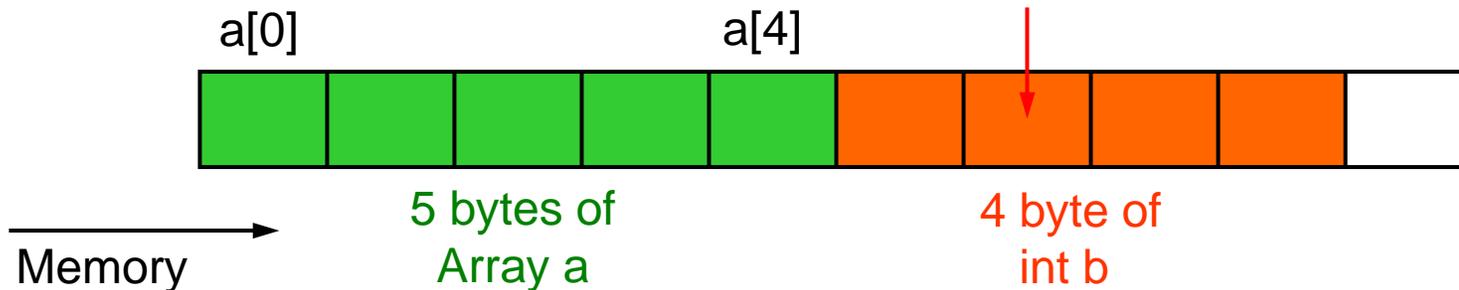
Example:

```
char a[5];  
int b = 0;
```

```
a[6] = 2;
```

Overwrites the second byte
of `int b`!

`a[6]` \Rightarrow `b = 512` (little endian)



Warning: These bugs are really hard to find!

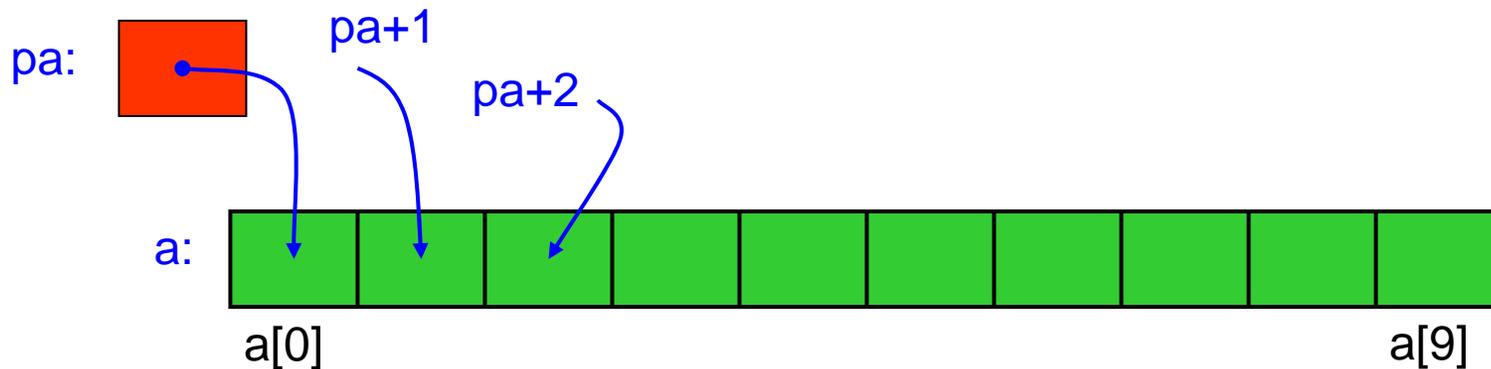
Pointer and Arrays

Arrays and Pointer are related. Instead of accessing an array element by its index also a pointer can be used:

```
int a[10];  
int* pa;  
pa = &a[0];
```

Accessing the first array element as

```
x = *pa; // equivalent to x=a[0]
```



For accessing the *i*th element, i.e., `a[i]`:

```
y = *(pa+i);
```

Note: That `pa+1` is not one added to the value of `pa`. Instead the address is `pa + sizeof(*pa)`.

Pointer and Arrays (2)

An array can be considered as a pointer to the first element of the array. Therefore the following expressions are equivalent:

<code>pa = &a[0];</code>	<code>↔</code>	<code>pa = a;</code>
<code>&a[i]</code>	<code>↔</code>	<code>a+i</code>
<code>pa[i]</code>	<code>↔</code>	<code>*(pa+i)</code>
<code>char s[]</code>	<code>↔</code>	<code>char* s</code>

An array can be passed (by-reference) to a function.

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

```
f(a);    // or equivalent  
f(&a[0]);
```

Function f:

```
void f(int* array)  
{  
    ...  
}  
  
// or equivalent  
void f(int array[])  
{  
    ...  
}
```

Arithmetic with Pointers

Pointers can be added and subtracted:

```
int a[10];  
int* pa = a;  
  
x = *pa;    // x = a[0]  
  
pa++;  
x = *pa;    // x = a[1]  
  
pa = pa+5;  // x = a[6]  
x = *pa;  
  
pa = pa-4;  // x = a[2]  
x = *pa;  
  
pa += 3;    // x = a[5]  
x = *pa;
```

Dynamic Allocation of Arrays

- Sometimes the size of an array is unknown at compile time
⇒ dynamic allocation at runtime
- Dynamic allocation with runtime support
Function `malloc` from Standard-C-Library fetches memory from operating system
- Dynamically allocated memory on heap
- In order to use `malloc` the prototype must be present, it can be included from `stdlib.h` header file.

`#include <stdlib.h>` ← At the beginning of the C file

Prototype of `malloc`:

```
void* malloc(size_t size);
```



↑ number of bytes to allocate
return value: pointer to first byte of array-block, or 0 (NULL) if allocation failed.

`malloc` returns a `void*` pointer (general untyped pointer) to the first element of the array. For later assignment to the array pointer a cast is necessary from `void*` to the effective pointer type of the array.

Dynamic Allocation of Arrays (2)

Example:

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

```
main()
```

```
{
```

```
    int* a;
```

```
    // allocates 4 MB main memory 1 Mio. ints)
```

```
    a = (int*)malloc(4*1024*1024);
```



Malloc returns a `void*` Pointer. Since we work with `int*` we need to cast the universal type to the more specific type `int*`.

```
    // deallocates (frees) array
```

```
    free(a);
```

```
}
```

Deallocation of previously allocated Memory

- A previously allocated block of memory eventually must be returned to the operating system.
 - implicitly when the program exits and the process terminates
 - explicitly by calling a function `free`
- Function `free`, is also part of the Standard-C-Library
- Only memory allocated earlier with `malloc` can be freed
- Always the entire memory block has to be freed (otherwise use `realloc`)

```
void free(void* ptr);
```



Pointer to first byte of memory block to be deallocated.

Note: This address has to be the address return by malloc when the block was allocated.

C-Strings

A String is an array consisting out of `char` elements. String literals are written in double-quotes: `"Hello World"`

The end of a C-String must be marked with a `Null-Character '\0'` (with a `String-Terminator`).

```
char text[] = "Hello World";
```

Is implicitly added

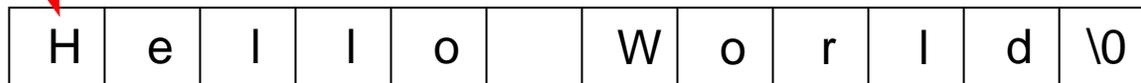
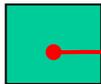
text:



Also possible is the following declaration:

```
char* text = "Hello World"
```

text:



The C-Language does not provide any other features for working with C-Strings, e.g., copying strings or retrieve the length of strings.

⇒ Instead use the function from the Standard-C-Library.

C-Strings (2)

Standard-C-Library provides a series of functions that work on C-Strings. In order to use these functions you need the prototypes from the `strings.h` header file.

```
#include <strings.h> ← At the beginning of the C file
```

Some useful functions:

```
// copies string source to destination
char* strcpy(char* destination, char* source);

// get number of characters in the string
size_t strlen(const char* str);

// compare string str1 with string str2
int strcmp(const char* str1, const char* str2);
```

Multi-dimensional Arrays

- Declaration

`TYPE array2d[num_dim1][num_dim2];` 2-dimensional

`TYPE array3d[num_dim1][num_dim2][num_dim3];` 3-dimensional

- `num_dim1, ..., num_dim3` is number of indices in each dimension
- Array elements are accessed as

`array2d[i][j]` and `array3d[i][j][k]`

with $0 \leq i < \text{num_dim1}$, $0 \leq j < \text{num_dim2}$ and $0 \leq k < \text{num_dim3}$

- Array elements are **stored in row-wise order**
- Initialisation of multi-dimensional arrays

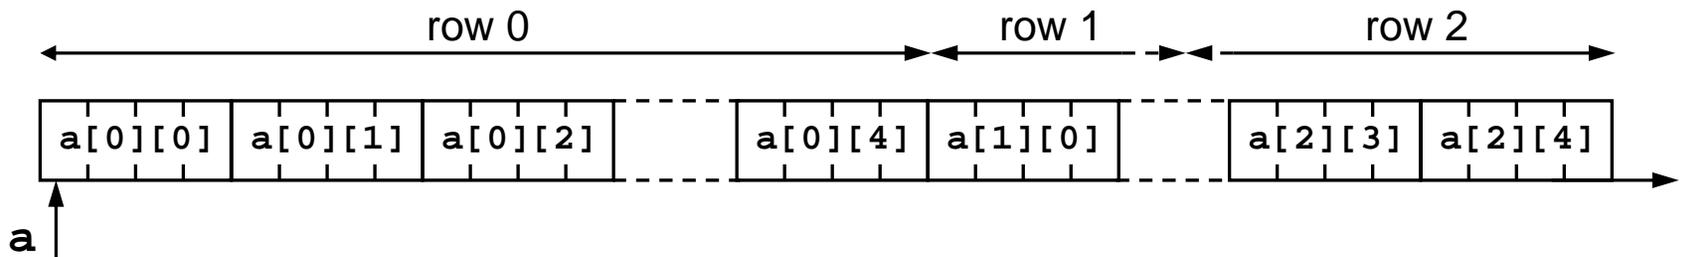
```
double matrix[4][4] = { {1.0, 0.0, 0.0, 0.0},  
                        {0.0, 1.0, 0.0, 0.0},  
                        {0.0, 0.0, 1.0, 0.0},  
                        {0.0, 0.0, 0.0, 1.0} };
```

can be omitted,
i.e. []

specifies storage size
(row-stride) required by the
compiler to compute address

2-dimensional Arrays

- Example: `int a[3][5];` (int: 4 byte)
- Row-wise storage:
 - stride (distance) between elements `a[i][j]` and `a[i][j+1]`: 4 bytes
 - stride between elements `a[i][j]` and `a[i+1][j]`: 4*5 bytes



- Index computation:
 - Array declaration: `TYPE array[n][m]`
 - Element access: `array[i][j]`
 - Pointer to element `array[i][j]`
`TYPE *element = (TYPE*)array+i*m+j`
 - Byte address to element `array[i][j]`
`char *byteaddr = ((char*)array) + (i*m + j)*sizeof(TYPE)`
- number of bytes
required for a variable
of type **TYPE**

Pointers and Multi-dimensional Arrays

- In order to compute the address of an element the compiler needs to know the number of columns (2-dim)
- Then how to pass array to a function (**TYPE*** does not work)?
- Example: `int array[3][15]` is to be passed to function `f` which is then invoked as `f(array)`
- Possible declarations of function `f`
 - `f(int x[3][15]) { ... }`
 - `f(int x[][15]) { ... }`
 - `f(int (*x)[15]) { ... }`

identical
since the number of
rows is irrelevant
- Note that for the last declaration the parentheses around `*x` **cannot** be omitted then
 - `f(int *x[15]) { ... }`
 - specifies an array of 15 pointer to integer instead of a pointer to an array with 15 integer!

Multi-dimensional Arrays – Limitations

- ANSI C[89] standard does not allow multi-dimensional arrays whose dimensions are determined at runtime[1], i.e., row-stride must be known at compile time.

Example:

```
- void foo() {  
    int n,m,i,j;  
    printf("n m: ");  
    scanf("%d %d", &n, &m);  
    int a[n][m]; ← size unknown at  
                  compile time  
    for (i=0; i<n; i++)  
        for (j=0; j<m; j++)  
            a[i][j] = ...  
}
```

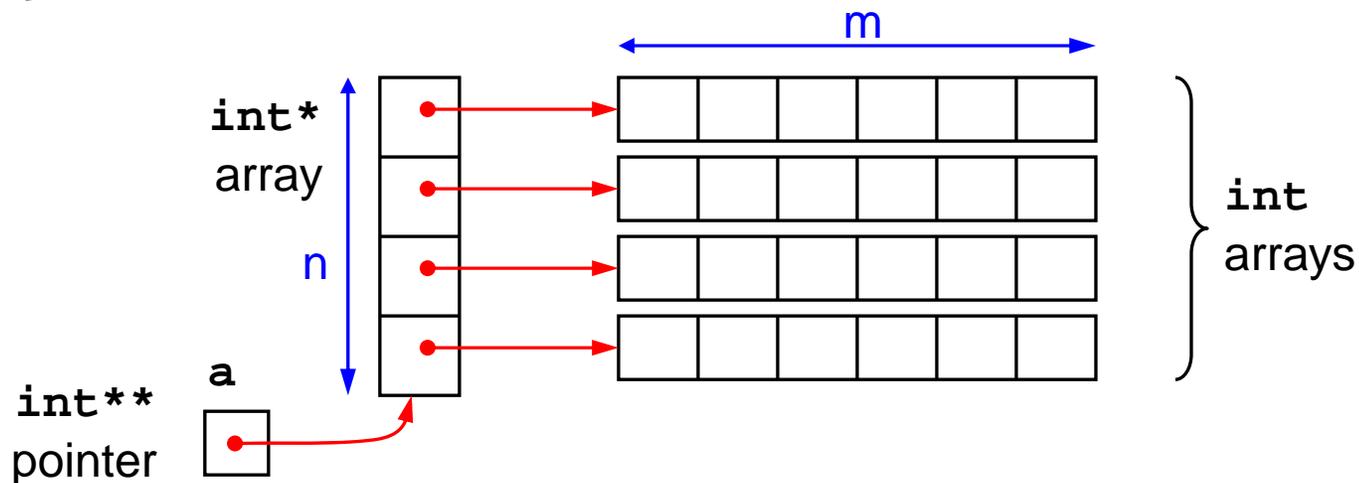
- Permitted in C99 and does compile on GCC (even if -std=c89 is specified)
- Works for local arrays that can be allocated on stack. But what about global arrays (stored in data segment (.bss or .data) but not on the stack)?

[1] Ritchie D.M.: Variable-Size Arrays. *Journal of C Language Translation*. Vol. 2 No. 2, 1989

Alternative 1: Arrays of Pointers

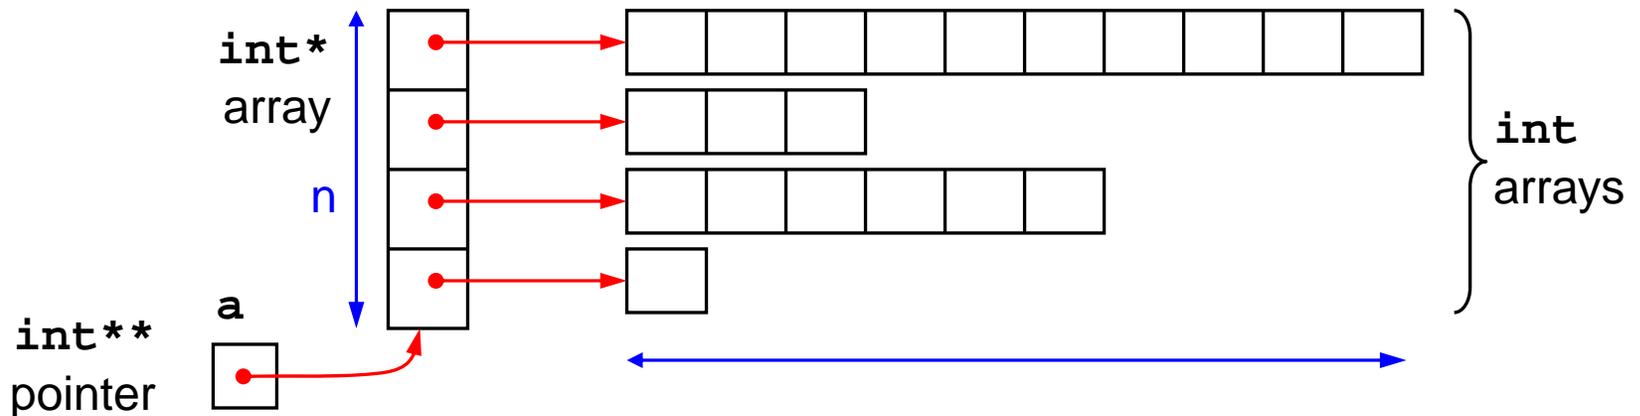
- Define multi-dimensional array as arrays of pointers
- Example 2-dimensional array

```
int i, n, m, **array;
printf("n m: ");
scanf("%d %d", &n, &m);
array = (int**)malloc(n*sizeof(int*));
for (i=0; i<n; i++) {
    array[i] = (int*)malloc(m*sizeof(int));
    for (j=0; j<m; j++) array[i][j] = f(i,j);
}
```



Alternative 1: Arrays of Pointers (2)

- Advantage
 - Permits dynamic array allocation for global multi-dimensional arrays
 - Allows definition “non-rectangular” arrays



- Disadvantage
 - Complicated allocation and deallocation
 - Memory leak if deallocation (with `free`) is forgotten
 - Element access requires two memory operations
 - Row-arrays not necessarily stored on consecutive memory location → non-uniform strides → negative effect on caching

Alternative 2: “Manual” Multidimensional Arrays

- Idea for array `TYPE x[n][m][r]`
 - map n-dimensional array to a 1-dimensional array and manually compute index
 - 1. Create 1-dimensional array of size $n*m*r$ elements
`TYPE *x = (TYPE*)malloc(n*m*r*sizeof(TYPE));`
 - 2. Compute index `x[i][j][k]` manually
for row-wise storage: `x[i*m*r+j*r+k]`
for column-wise storage: `x[i+j*n+k*m*n]`
- Suggestion: define a macro
 - `#define array(i,j,k) x[(i)*m*r+(j)*r+(k)]`
 - Usage: `y = array(1,2,3)`
- Advantage
 - Direct control about storage locations of elements, possibility to exploit characteristics of caching → better performance possible

Complicated Pointer Declarations

- `char **argv`
argv: pointer to pointer to char
- `int (*x)[13]`
x: pointer to array[13] of int
- `int *x[13]`
x: array[13] of pointer to int
- `void *comp()`
comp: function return pointer to void
- `void (*comp)()`
comp: pointer to function returning void
- `char ((*x())())[5]`
x: function returning pointer to array[] of pointer to function returning char
- `char ((*x[3])())[5]`
x: array[3] of pointer to function returning pointer to array[5] of char
 - Isn't C beautiful?

Enumeration Types

An enumeration is a list of constant integer values. A variable of this enumeration type can be a value of the enumeration list.

tag is the name of the enumeration

```
enum tag { list of names };
```

Example:

Tag 'direction' is the name of the enumeration

(north, south, east, west)

```
enum direction { north, south, east, west };
```

```
enum direction d; ← Variable d is an enumeration of  
'direction'
```

```
direction = east; ← Any value of the enumeration list  
can be used in the assignment. The  
symbolic name is substituted by the integer  
constant.
```

```
if (direction == south)  
    printf("wrong way!");
```

Enumeration Types (2)

Enumeration are represented as integer values. In symbolic names are enumerated automatically in the definition statement.

```
enum direction { north, south, east, west };
```

```
printf("%d\n", north);    ⇒ 0  
printf("%d\n", south);   ⇒ 1  
printf("%d\n", east);    ⇒ 2  
printf("%d\n", west);    ⇒ 3
```

Optionally the corresponding integer value can be specified for each symbolic name in the enumeration list.

Example: Definition for missing Boolean type in C as enumeration:

```
enum bool { true=1, false=0 };
```

```
enum bool test = true;  
printf("%d\n", test);
```

Enumeration Types (3)

Anonymous enums do not have a type name. Thus no other variables can exist except the ones in the declaration statement of the enumeration list.

```
enum { SEEKING, FOLLOWING, STOP } state, anotherState;
```

In general a tag name and a number of variables of that enum type can be specified in the declaration statement:

```
enum state_type { SEEKING, FOLLOWING, STOP }  
state, anotherState;
```

This is identical to the following three lines:

```
enum state_type { SEEKING, FOLLOWING, STOP };  
enum state_type state;  
enum state_type anotherState;
```

Enumeration Types (4)

- Note since an enum is in fact an integer variable of type enum it basically can have any integer assigned to it.
- Generally the compilers do not complain if a value is assigned that is not in the enumeration list (even though they might in order to assure type consistency)

```
enum bool { true=1, false=0 };  
enum bool value;
```

`value = 5;` Valid even though 5 is no in the enumeration list

- The GCC does not check that the value is part of the enumeration list.
- An enumeration type enhances legibility of the code, e.g. the debugger can replace integer values by the symbolic names.

Type Names and Aliases (typedef)

`typedef` creates a new name for an existing type.

existing type *new alias, i.e., name of new type*

```
typedef type new-name;
```

Example:

```
typedef unsigned char BYTE;      ← BYTE is now an unsigned  
typedef int seconds;              8 Bit type
```

```
BYTE biteMe;                      ← Variable biteMe is of type BYTE  
seconds waitTime;                and thus also of type  
                                    “unsigned char”.
```

Therefore the variable declaration is equivalent to:

```
unsigned char myBite;  
int waitTime;
```

Typedef and Enumeration Types

Variables of enumeration type have to be declared with keyword `enum` followed by the tag name.

Using `typedef` an “real” type alias can be defined. The new type can be used as any other type (e.g. `int`, `float`). The keyword `enum` no longer has to be used in the variable declaration:

```
enum bool { true=1, false=0 };
typedef enum bool boolean;

boolean aBoolean;           // uses typedef alias
enum bool anotherBoolean;  // uses original enum definition
```

Since the enumeration type tag ‘bool’ is no longer used, one can declare the enumeration as an anonymous enum:

```
typedef enum { true=1, false=0 } boolean;
boolean aBoolean;
```

Structs – Example

For a geometry package a point object is used that is represented by two coordinates. The components are combined to a single entity, the **struct**. (related to OO, in fact in C++ **struct** is the same keyword of **class** except that all members are public).

```
struct point {  
    int x;  
    int y;  
};
```

As for the enumeration types, for a struct a tag name has to be specified. The struct with tag name “point” then consists of two integer members x and y.

Variable myPoint of type point is declared as (similar to enums):

```
struct point mypoint;
```

The members of a struct can be accessed by their name preceded by a dot ‘.’:

```
mypoint.x = 4;  
mypoint.y = 3;
```

Structs

As for enumeration types, in the declaration statement also variables can be defined.

```
struct {  
    int x;  
    int y;  
} A, B;
```

Variables A and B are structs with components x and y. Important: The struct itself is nameless, i.e. no further variables of this struct can be defined.

```
struct point {  
    int x;  
    int y;  
} A, B;
```

Variables A and B are both types of struct point with two integer components x and y.

Struct variables can be initialised in the declaration statement. The initial values are specified in braces. The order of the values is the same as the member list in the definition of the struct.

```
struct point A = { 3, 4 };
```

← identical →

```
{ struct point A;  
  A.x = 3;  
  A.y = 4;
```

Structs – Summary

The general syntax of struct definition is:

```
struct tag-name {  
    type1 component1;  
    type2 component2;  
    ...  
    typen componentn  
} variable1, variable2;
```

<i>tag-name</i>	name of the struct
<i>type1</i>	type of component1
<i>componente1</i>	name of member
<i>variable1</i>	variable1 is of this struct type

Later further variables of this struct can be defined using keyword **struct** and the tag name:

```
struct tag-name variable-name;
```

The member of the struct can be access the variable name followed by a dot ‘.’ and the name of the member.

```
variable.component
```

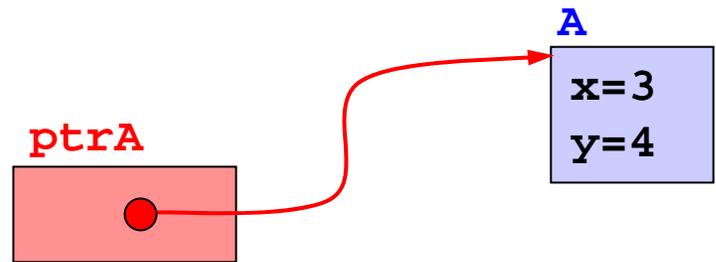
Struct variables can be initialised in the declaration statement. The initial values are specified in braces in the same order as the member in the struct definition.

```
struct tag-name variable-name = { initvall1, ... };
```

Structs und Pointers

Pointer to struct variables are also possible:

```
struct point A = { 3, 4 };  
struct point* ptrA = &A;
```



However there is an important difference. Members in struct variables are accessed using a dot '.' between variable name and the name of the struct member.

```
A.x = 5;  
A.y = 6;
```

But using a pointer to a structs the arrow '->' has to be used for the accessing the member.

```
ptrA->x = 2;  
ptrA->y = 5;
```

Since pointer `ptrA` points to `A` the members of variable are manipulated.

```
printf("( %d, %d)\n", A.x, A.y); ⇒ (2, 5)
```

Memory Layout of Structs

Example:

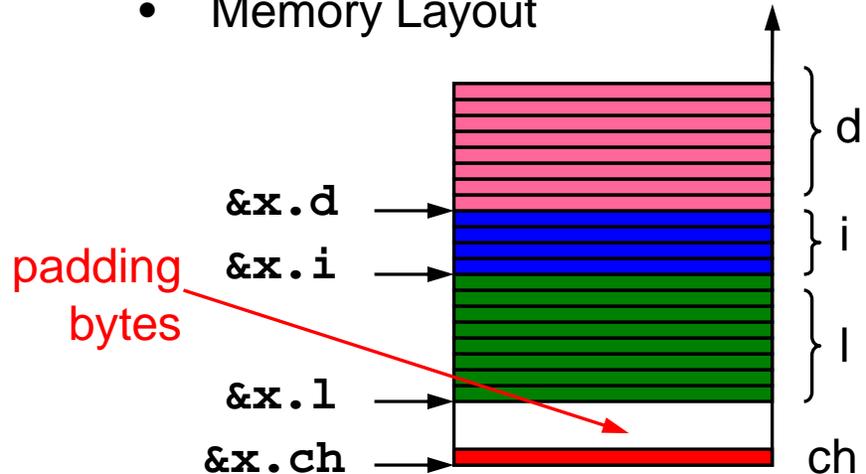
```
struct foo {
    char ch;
    long long int l;
    int i;
    double d;
};

main() {
    struct foo x;
    printf("ch:%p\n",&x.ch);
    printf("l: %p\n",&x.l);
    printf("i: %p\n",&x.i);
    printf("d: %p\n",&x.d);
    printf("%d\n",sizeof(x));
}
```

compiled with GCC
(for MacOS X, PPC)

- Output
ch: bffff988
l: bffff98c
i: bffff994
d: bffff998
24

- Memory Layout



- GCC does not use dense packing, it inserts padding bytes between the struct members
- By using of padding bytes the members can be word-aligned which speeds up memory access.

Memory Layout of Structs (2)

Enforcing packed structs on GCC:

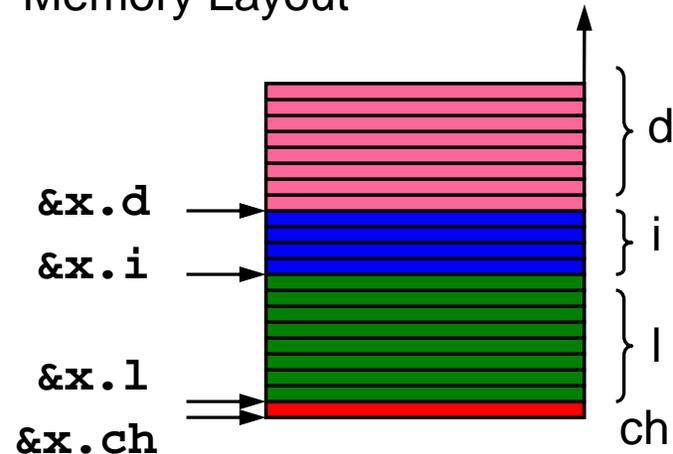
```
struct foo {  
    char ch;  
    long long int l;  
    int i;  
    double d;  
} __attribute__((__packed__));  
main() {  
    struct foo x;  
    printf("ch:%p\n",&x.ch);  
    printf("l: %p\n",&x.l);  
    printf("i: %p\n",&x.i);  
    printf("d: %p\n",&x.d);  
    printf("%d\n",sizeof(x));  
}  
__attribute__((__packed__))
```

instructs GCC not to insert padding bytes.

Note: this only works for GCC not necessarily for other compilers.

- Output
ch: bffff988
l: bffff989
i: bffff991
d: bffff995
21

- Memory Layout



Typedefs and Structs

As for enums, typedef can be used to create an alias to the struct type. Thus if the typedef type is used, the keyword `struct` can be omitted in the variable declaration.

```
typedef struct { // anonymous struct
    int x;
    int y;
} point; // typedef alias with name point
```

```
point A; // definition using struct and typedef name
point B = {3, 2};
```

Without using `typedef` the declaration is:

```
struct point {
    int x;
    int y;
};
```

```
struct point A; // point is tag name of the struct
struct point B;
```

Structs in Linked Lists

- There is an important difference between the tag name and the typedef name of a struct when the struct itself is used as a member of itself, e.g., in linked lists

- Example:

```
typedef struct {  
    node* previous;  
    node* next;  
    int value;  
} node;
```

- Cannot be compiled: “error: parse error before 'node'”. The identifier “node” is not known when it is used in the definition of the member **previous**.

- Instead the struct tag name has to be used:

```
typedef struct node_t {  
    struct node_t *previous;  
    struct node_t *next;  
    int value;  
} node;
```

Unions

unions provide a way to manipulate different kinds of data in a single area of storage. Unions consists of several members of different types that are referenced to by a name. But in contrast to structs all members of a union are mapped to the same memory location.

Example:

```
union int_float {  
    unsigned int int_value; // 32 bit signed integer  
    float float_value;      // 32 bit floating point number  
} x; ← x is a variable of this union type
```

tag name of this union type

As for structs, union members can be accessed with dot '.' followed by the union name (additionally for union pointers members have to be accessed with arrow '->').

```
x.float_value = 3.141592654f; // set as float  
printf("%f", x.float_value);  ⇒ 3.14593  
  
// read content as unsigned int (decimal and hex)  
printf("%u\n", x.int_value);  ⇒ 1078530011  
printf("0x%X\n", x.int_value); ⇒ 0x40490FDB
```

Unions (II)

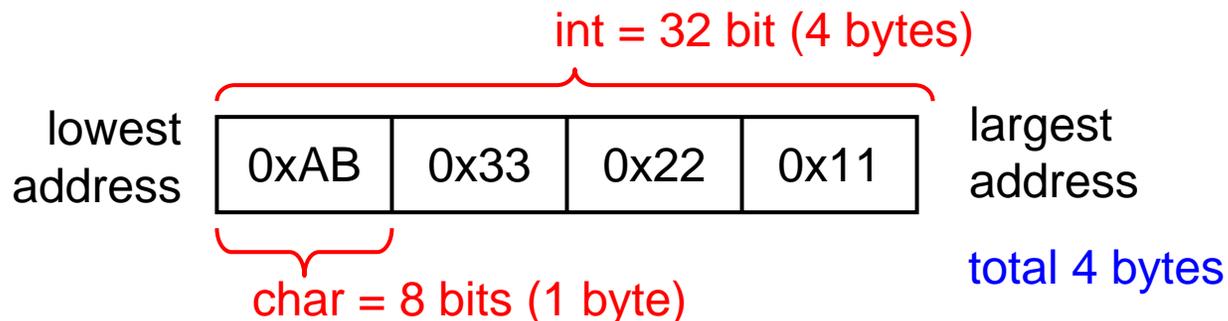
If the member types of a union differ in size the size of the union is equal to the size of the largest member type.

```
union int_char { ← the size of this union is 4 bytes (= 32 bits)
    unsigned char char_value; // 8 bit
    unsigned int int_value; // 32 bit
} x;
```

```
x.int_value = 0x11223344;
x.char_value = 0xAB;
```

```
printf("int: 0x%X\n", x.int_value); ⇒ 0x112233AB
printf("char: 0x%X\n", x.char_value); ⇒ 0xAB
printf("sizeof: %d\n", sizeof(x)); ⇒ 4 (bytes)
```

Representation of `x` in memory: Little Endian representation (Intel & Co.)



Unions – Summary

The general syntax of a union definition consists of a tag name, the member (with their types) and a list of variables of this union types.

```
union tag-name {  
    type1 member1;  
    type2 member2;  
    ...  
    typen membern  
} variable1, variable2;
```

tag-name

Type name of this union

type1

Type of member1

member1

member name

variable1

variable of this union type

Definition of variables of union with specified union tag name:

```
union tag-name variable-name;
```

Union members are accessed with a dot '.':

```
variable-name.member
```

and for pointers to unions (as with structs) with '->':

```
union tag-name* ptr;  
ptr->member;
```

Typedef und Union

As for enums and structs, a type alias can be defined. Instead of using the keyword union and the union name the typedef alias can be used to declare a variable.

```
typedef union {  
    unsigned char char_value;  
    unsigned int int_value;  
} int_char;
```

```
int_char x; // definition using typedef alias
```

Or using the union tag name:

```
union int_char {  
    unsigned char char_value;  
    unsigned int int_Value;  
};
```

```
union int_char x; // int_char is the union tag name
```

Application of Unions

- Representation of IP addresses in BSD sockets

```
struct sockaddr_in {
    short    sin_family;           // e.g. AF_INET for internet
    u_short  sin_port;           // port number
    struct   in_addr sin_addr;    // ip address
    char     sin_zero;
};
struct in_addr {
    union {
        struct { u_char s_b1, s_b2, s_b3, s_b4 };
        struct { u_short s_w1, s_w2 };
        u_long  S_addr;
    } S_un;
};
```

- Set IP “129.132.98.12”:

```
struct sockaddr_in ddr;
addr.sin_addr.S_un.s_b1 = 129;
addr.sin_addr.S_un.s_b2 = 132;
addr.sin_addr.S_un.s_b3 = 98;
addr.sin_addr.S_un.s_b4 = 12;
```

- Macros `htonl(...)` used to adapt endianness from host to network order, and `ntohl(...)` from network to host order.

Bit Fields

- When storage space is a premium, it may be necessary to pack several objects into a single machine word, e.g. a collection of bit flags.

- Example

```
union error_flags {
    unsigned char error_byte;
    struct {
        unsigned int  cntr:3; // 3 bits
        unsigned int  ovfl:1; // 1 bit
        unsigned int  regs:2; // 2 bits
        unsigned int  flgs:2; // 2 bits
    } error_bits;
} error;
error.error_byte = 0;
error.error_bits.ovfl = 1;
error.error_bits.regs = 3;
printf("0x%X", error.error_byte);
```

- Aligning of bit field depends also on endian-ess of the machine.

on PowerPC: \Rightarrow 0x1C

2^7 2^0							
0	0	0	1	1	1	0	0

on x86: \Rightarrow 0x38

2^7 2^0							
0	0	1	1	1	0	0	0

Summary Types

- Enum types are integer types whose values are specified as symbolic names in the enumeration list.
- Structs are a collection of variables of different types and size. Structs allow encapsulation of attributes to objects.
- Unions is a collection of variables of different type and size. But in contrast to structs all members of the union are mapped to the same storage area. The size of a union is equal to the largest size of its member types.
- Typedefs can be used to create alias for enum, struct and union types. When the typedef name is used the keywords enum, union and struct can be omitted. In fact with its typedef name a user defined type can be treated as any other data type (int, float, etc.)