

Deductive Program Verification with Frama-C

Maria João Frade

HASLab - INESC TEC
Departamento de Informática, Universidade do Minho

2022/2023

Roadmap

- **Introduction**
 - ▶ Frama-C; WP plugin; ACSL; memory models; annotations; properties; local properties status; runtime errors;
- **Program Specification**
 - ▶ function contracts; safety; behaviours; function calls; logical predicates; state labels;
- **Program Verification**
 - ▶ loops and proof; loop invariants; termination policy; loop variants; proof failures;
- **Other Features**
 - ▶ using axiomatics; algebraic modeling; ghost code;

Bibliography

- **Introduction to C program proof with Frama-C and its WP plugin.**
Allan Blanchard, July 1, 2020.
- **Frama-C / WP Plug-in Manual - Frama-C 26.1 (Iron).**
Patrick Baudin, François Bobot, Loïc Correnson, Zaynah Dargaye, Allan Blanchard.
- **ACSL-by-Example - Towards a Formally Verified Standard Library.**
Jens Gerlach, November 2020.

Frama-C (FRAmework for Modular Analysis of C programs)

- Frama-C is an open-source platform for [static analysis of C code](#).
- Developed at CEA LIST and INRIA Saclay.
- Frama-C 1st release: 2008. Previous: Why+Caduceus (early 2000's), CAVEAT (90's).
- **Plugin architecture.** Various plugins: value analysis, deductive verification, slicing, dependency analysis, impact analysis, metrics calculation, ...
- Includes [ACSL specification language](#).
- **Extensible and collaborative platform.**
 - ▶ One can add new plugins.
 - ▶ Allows collaboration of analyses over the same code.
 - ▶ Inter-plugin communication through ACSL formulas.
- <http://frama-c.com/>

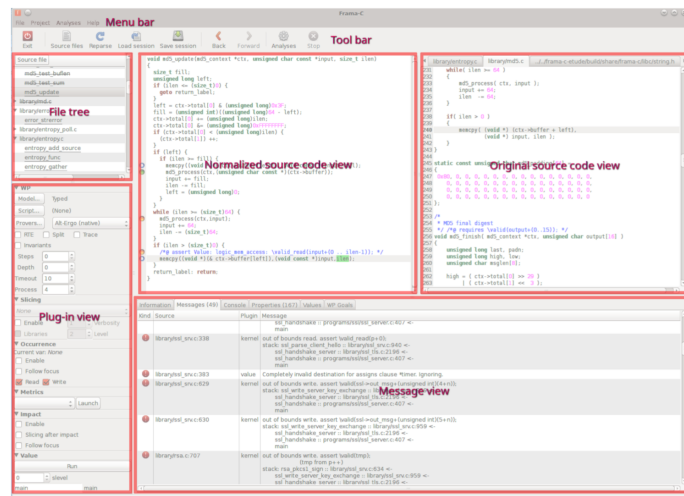
Deductive verification with Frama-C

- Frama-C has two plugins for deductive verification.
 - ▶ **Jessie** (developed at INRIA, 2009)
 - ▶ **WP** (developed at CEA LIST, 2012)
- Both plugins are based on **Hoare logic** and **weakest precondition calculus**.
 - ▶ **Jessie** relies on a separation memory model and operates by compiling to Why. Not actively maintained.
 - ▶ **WP** focuses on parametrization w.r.t. the memory model (different models are available).
- Both plugins allow to prove that C functions satisfy their specification (expressed in ACSL).
 - ▶ **Proofs are modular**: the specifications of the called functions are used to establish the proof without looking at their code.

Preparing the sources

- For the creation of an analysis project, Frama-C performs several steps for preparing the sources to be analyzed.
 - ▶ **Pre-processing phase**. Frama-C performs some pre-processing of the source code.
 - ▶ **Merging phase**. Frama-C parses, type-checks and links the code. It also performs these operations for the ACSL annotations.
 - ▶ **Normalization phase**. Frama-C performs a number of local code transformations aiming at making further work easier for the analyzers.
- Analyses usually take place on the **normalized version of the source code**.
- Normalization gives a program which is **semantically equivalent** to the original one.

frama-c-gui



WP plugin

- Proof of **safety and functional properties** of C annotated programs.
- Implements an **Weakest Precondition calculus parameterized by a memory model** (to represent pointers and heap values).
- **WP operates as follows**:
 - ▶ receives as **input** a normalized C program with ACSL annotations;
 - ▶ generates the **verification conditions (VCs)** via WP VCgen;
 - ▶ discharges the VCs using **external theorem provers** via Why3.
- The WP plugin is **cooperative**, i.e., it allows to combine WP calculus with other techniques available via other plugins.

Memory models

- The essence of a weakest precondition calculus is [to translate code annotations into mathematical properties](#).
- To apply the WP calculus to programs [dealing with pointers](#) one has to have a **memory model**, which defines a [mapping from values inside the C memory heap to mathematical terms](#).
- The WP has been designed to support different memory models:
 - ▶ **Hoare Model**. A very efficient model that generates concise proof obligations. It simply maps each C variable to one pure logical variable. However, the heap cannot be represented in this model.
 - ▶ **Typed Model**. Heap values are stored in several separated global arrays, one for each atomic type and an additional one for memory allocation. Pointer values are translated into an index into these arrays.
 - ▶ **Mixed models**. It is possible to mix the above models whenever the address of a variable is never taken in a program. Such an optimization is activated by default in the WP plugin, since it is very effective.

Arithmetic models

- The WP plugin is able to take into account the precise semantics of integral and floating-point operations of C programs.
- The WP has been designed to support different arithmetic models:
 - ▶ **Machine Integer Model**. The kernel options are used to determine if an operation is allowed to overflow or not. Using `-rte` or `-wp-rte` will generate all the necessary assertions to be verified.
 - ▶ **Natural Model**. Integer operations are performed on mathematical integers.
 - ▶ **Float Model**. floating-point values are represent in a special theory with dedicated operations over float and double values and conversion from and to their real representation via rounding.
 - ▶ **Real Model**. floating-point operations are transformed on reals, with no rounding.

Program annotations

- Frama-C supports writing [code annotations with the ACSL language](#).
- The purpose of annotations is [to formally specify the properties of C code](#).
- Annotations can originate from a number of different sources:
 - ▶ **the user** who writes his own annotations;
 - ▶ **some plugins** may generate code annotations (cf. RTE plugin);
 - ▶ **the kernel** of Frama-C, that attempts to generate as precise an annotation as it can, when none is present.

Properties

- A **property** is a [logical statement bound to a precise code location](#).
- A property might originate from an ACSL code annotation or by a plugin-dependent meta-information.

Property validity

Consider a program point i , and call T the set of traces that run through i .

- A logical property P is **valid at i** if it is valid on all $t \in T$.
- Conversely, any trace u that does not validate P , stops at i : properties are **blocking**.

Local property status

- An important part of the interactions between Frama-C components rely on their capacity to emit a **judgment on the validity of a property P at program point i** .
- In Frama-C nomenclature, this judgment is called a **local property status**, which has two parts: **the status and a list of dependencies**.
 - ▶ The first part of a local status ranges over the following values:
 - ★ **True** when the property is true for all traces;
 - ★ **False** when there exists a trace that falsifies the property;
 - ★ **Maybe** when the emitter e cannot decide the status of P .
 - ▶ An emitter can add a list of **dependencies**, which is the set of properties whose validity may be necessary to establish the judgment.
 - ★ The dependencies are meant *as a guide* to safety engineers. They are neither correct, nor complete.

An example: **abs**

A specification for the **abs** function could be (**abs_init.c**):

```
/*@ ensures (x >= 0 ==> \result == x) &&
           (x < 0 ==> \result == -x);
    assigns \nothing;
*/
int abs (int x) {
    if (x >= 0) return x ;
    return -x ;
}
```

We can run Frama-C to determine if the implementation is correct against the specification using

- the **command line interface** of Frama-C, or
- the **graphical user interface** of Frama-C.

Invoking WP

```
$ frama-c -wp abs_init.c
[kernel] Parsing abs_init.c (with preprocessing)
[wp] Warning: Missing RTE guards
[wp] 3 goals scheduled
[wp] Proved goals:    3 / 3
Qed:                3 (4ms)
```

It results in 3 VCs. The all discharged internally by the Qed simplifier of WP.

- Notice the **warning "Missing RTE guards"**, emitted by the WP plugin.
 - ▶ The WP calculus implemented **relies on the hypothesis that programs are runtime-error free**.
 - ▶ By default, the WP plugin does not generate any proof obligation for verifying the absence of runtime errors in the code. They can be proved by generating all the necessary annotations with the **RTE plugin**.

Runtime errors

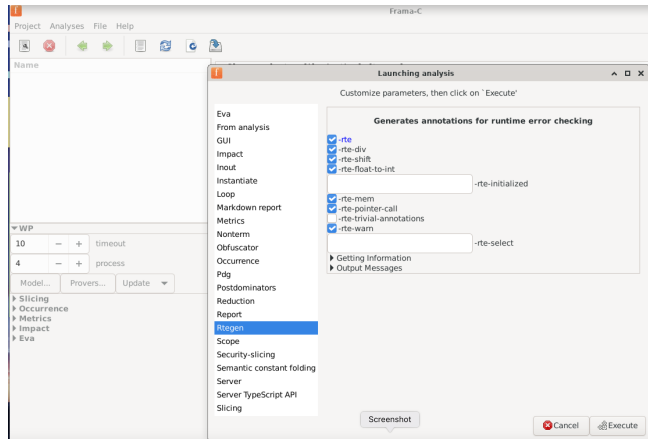
```
$ frama-c -wp -wp-rte abs_init.c
[kernel] Parsing abs_init.c (with preprocessing)
[rte] annotating function abs
[wp] 4 goals scheduled
[wp] [Alt-Ergo] Goal typed_abs_assert_rte_signed_overflow : Unknown (52ms)
[wp] Proved goals:    3 / 4
Qed:                3 (4ms)
Alt-Ergo:            0 (unknown: 1)
```

It results in 4 VCs. Three VCs discharged internally by the Qed simplifier and one sent to Alt-Ergo (with an inconclusive response).

Why does the proof fail?

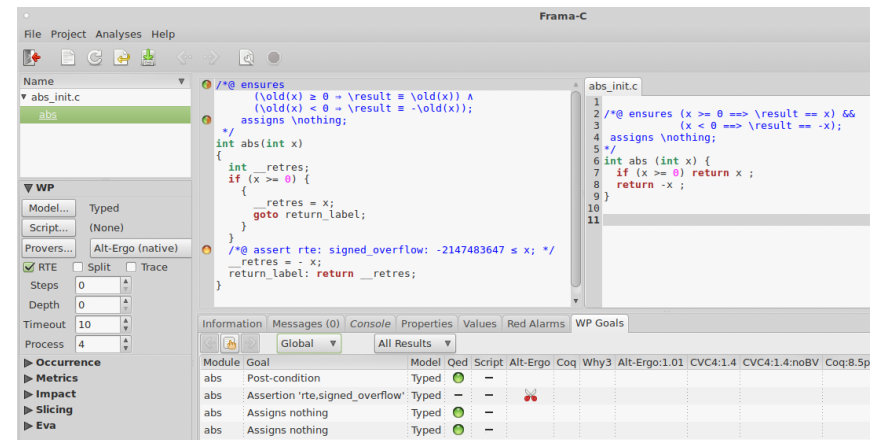
Frama-C GUI and the RTE plugin

- Call the Frama-C GUI by typing `frama-c-gui`
- Configure the system to launch the analysis RTE (Run Time Errors) by activating in the **Rtegen** the flag `-rte`.



Frama-C GUI

- Load the source file `abs_init.c`



- Alternatively `frama-c-gui abs_init.c -wp -wp-rte`

Frama-C GUI

- The options from the WP side panel correspond to some options of the plugin command-line.
- The *status of each code annotation* is reported in the left margin. The meaning of icons is the same for all plugins in Frama-C.

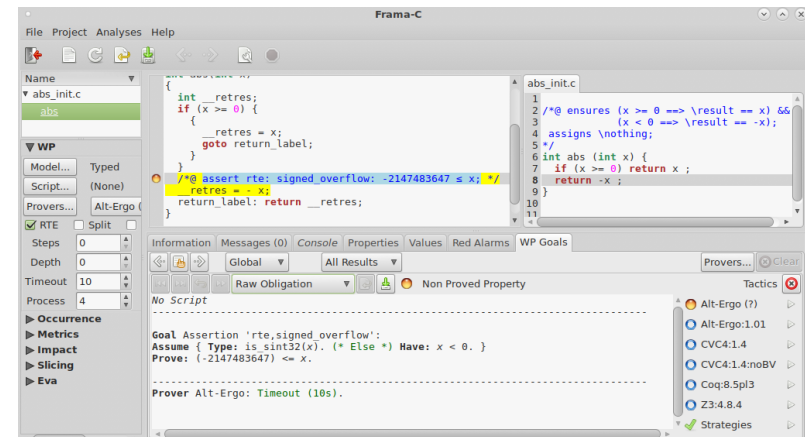
Icons for properties:

- No proof attempted.
- The property has not been validated.
- The property is *valid* but has dependencies.
- The property and *all* its dependencies are *valid*.

- **Left-clicking** on an object in the normalized code view displays information about it in the "Information" tab of the Messages View and displays the corresponding object of the original source view.
- **Right-clicking** on an object opens a contextual menu. Items of this menu depend on the kind of the selected object and on plugin availability.
- **Ctrl-clicking** on the original source code opens an external editor.

Proof editor

This panel focus on one goal generated by WP, and allow the user to **visualize the logical sequent to be proved**, and to interactively decompose a complex proof into smaller pieces by applying *tactics*.



Proof editor

- The logical sequent consists of a formula to **Prove** under the hypotheses listed in the **Assume** section. Each hypothesis can consists of :

Type: formula expressing a typing constraint;
Init: formula characterizing global variable initialisation;
Have: formula from an assertion or an instruction in the code;
When: condition from a simplification performed by Qed;
If: structured hypothesis from a conditional statement;
Either: structured disjunction from a switch statement.
Stmt: labels and C-like instructions representing the memory updates

- There are several modes to display the current goal:

Autofocus: filter out clauses not mentioning *focused* terms (see below);
Full Context: disable autofocus mode — all clauses are visible;
Unmangled Memory: autofocus mode with low-level details of memory model;
Raw Obligation: no autofocus and low-level details of memory model.

ACSL (ANSI C Specification Language)

- Aims at specifying behavioral properties of C source code (inspired in JML).
- Based on the notion of **contract**.
 - ▶ Each function contract (safety included) is verified independently.
 - ▶ The correctness of all the remaining functions in the program is assumed.
- Specifications are given as **annotations in comments** written directly in C source files (`/*@ */` and `//@`).
- **Basic features**
 - ▶ First-order logic + Pure C expressions
 - ▶ C types + \mathbb{Z} + \mathbb{R}
 - ▶ Built-in predicates and logic functions, particularly for pointers:
`\valid(p)`, `\valid(p+(0..n))`, `\separated(p+(0..5),q+(0..3))`,
`\block_length(p)`, etc.
 - ▶ Users defined predicates and logic functions.

ACSL contracts

- The contract of a function is composed of two main parts:
 - ▶ **Precondition clauses** introduced by the **requires** keyword.
 - ▶ **Postcondition clauses** introduced by the **ensures** keyword.
 - ★ The **assigns** keyword allows to specify which (non-local) memory locations can be modified by the function. If the function does not assign any memory location **assigns \nothing** clause can be added.
 - ★ The **\result** keyword is used to refer to the value returned by the function.
 - ▶ Several clauses on the same type are equivalent to a single clause expressing the conjunction of the corresponding properties.
- Deductive verification with the plugin WP can be used to prove that:
 - ▶ each function respects its contract
 - ▶ each call to a function satisfies the precondition before the call.
- The ACSL keyword **assert** introduces an assertion, i.e. a property that must be checked at a particular program point.

Proof modularity

- When trying to prove a property, the contracts of the called functions are used to establish the proof without looking at their code.
- If the contract does not describe precisely the function behavior, it may be impossible to prove some expected properties.

Consider the following code with an imprecise specification for the max function:

```
/*@ ensures \result >= a && \result >= b ;  
*/  
int max(int a, int b) {  
    return (a >= b) ? a : b ;  
}  
  
extern int x ;  
  
int main() {  
    x = 3;  
    int r = max(4, 2);  
    //@ assert r == 4 ;  
    //@ assert x == 3 ;  
}
```

Use `frama-c-gui` to analyze the file `max_init.c` with WP.

Example 1

```
/*@
  ensures \result >= a && \result >= b ;
*/
int max(int a, int b) {
  return (a >= b) ? a : b ;
}

extern int x ;

int main() {
  x = 3;
  int r = max(4, 2);
  //@ assert r == 4 ;
  //@ assert x == 3 ;
}
```

- The max function can be proved to correctly implement the specification that has been stated.
- However the assertion `r == 4` will not be proved.
- We can reinforce the postcondition with

```
ensures \result == a || \result == b ;
```

Example 1

```
/*@
  ensures \result >= a && \result >= b ;
  ensures \result == a || \result == b ;
*/
int max(int a, int b) {
  return (a >= b) ? a : b ;
}

extern int x ;

int main() {
  x = 3;
  int r = max(4, 2);
  //@ assert r == 4 ;
  //@ assert x == 3 ;
}
```

- However the assertion `x == 3` will not be proved. (*why?*)
- The assigns clause is not provided, so the function is just assumed to assign any memory location.
- We can reinforce the specification with `assigns \nothing ;`

Example 1

```
/*@
  ensures \result >= a && \result >= b ;
  ensures \result == a || \result == b ;
  assigns \nothing ;
*/
int max(int a, int b) {
  return (a >= b) ? a : b ;
}

extern int x ;

int main() {
  x = 3;
  int r = max(4, 2);
  //@ assert r == 4 ;
  //@ assert x == 3 ;
}
```

With this specification both assertions will be proved by WP.

Example 2

Run WP with `abs_init.c`

```
/*@ ensures (x >= 0 ==> \result == x) &&
         (x < 0 ==> \result == -x);
    assigns \nothing;
*/
int abs (int x) {
  if (x >= 0) return x ;
  return -x ;
}
```

Why does the proof fail?

- For `x==INT_MIN`, `-x` overflows
- Example: for 32-bit, `x==INT_MIN=-231` while `x==INT_MAX=231 - 1`

Example 2 (overflow safety)

```
#include <limits.h>

/*@ requires x > INT_MIN;
    ensures (x >= 0 ==> \result == x) &&
           (x < 0 ==> \result == -x);
    assigns \nothing;
*/
int abs (int x) {
    if (x >= 0) return x ;
    return -x ;
}
```

It is sufficient to add as a precondition that x must be strictly greater than `INT_MIN` to guarantee that the underflow will never happen.

Notice that it is also necessary to include the header where `INT_MIN` is defined.

Example 2 (overflow safety)

Module	Goal	Model	Qed	Script	Alt-Ergo	Coq	Why3	Alt-Ergo:1.01	CVC4:1.4	CVC4:1.4:nbBV	Coq:8.5p
abs	Post-condition	Typed	✓	—	—	—	—	—	—	—	—
abs	Assertion 'rte.signed_overflow'	Typed	✓	—	—	—	—	—	—	—	—
abs	Assigns nothing	Typed	✓	—	—	—	—	—	—	—	—
abs	Assigns nothing	Typed	✓	—	—	—	—	—	—	—	—

ACSL annotations

- Global annotations
 - ▶ function contracts
 - ▶ global invariants
 - ▶ type invariants
 - ▶ logic specifications
- Statement annotations
 - ▶ loop annotations
 - ▶ assert clauses
 - ▶ statement contracts
 - ▶ ghost code

- We will learn the language through examples.

An example: maxarray

```
/*@ requires 0 < size && \valid(u+(0..size-1));
    ensures 0 <= \result < size;
    ensures
        \forall integer a; 0 <= a < size ==> u[a] <= u[\result];
    assigns \nothing;
*/
int maxarray(int u[], int size) {
    int i = 1;
    int max = 0;

    /*@ loop invariant \forall integer a;
        0 <= a < i ==> u[a] <= u[max];
        loop invariant 0 <= max < i <= size;
        loop assigns max, i;
        loop variant size-i; */
    while (i < size) {
        if (u[i] > u[max]) max = i;
        i++;
    }
    return max;
}
```


An example: `maxarray`

- Run Frama-C with the file `maxarray.c`
- Explore the WP plugin with this example.
- Observe VCs generated.
 - ▶ There are VCs related to the verification of a **function's default behavior**, which includes verification of its postcondition, frame condition, loop invariants and intermediate assertions.
 - ▶ There are VCs guarding against **safety violations** such as null-pointer dereferencing, buffer overflow, arithmetic overflow, division by zero, termination, etc.

Program Specification

Example 3 (memory safety)

Why does the proof of this program fail?

```
/*@ ensures \result >= *p && \result >= *q;  
    ensures \result == *p || \result == *q;  
*/  
int max_ptr (int *p, int *q) {  
    if (*p >= *q) return *p ;  
    return *q ;  
}
```

Run WP with `max_ptr_init.c`

Nothing ensures that `p` and `q` are valid pointers!

WP automatically generates VCs to check memory access validity.

Example 3 (frame conditions)

Is this specification complete?

```
/*@ requires \valid(p) && \valid(q);  
    ensures \result >= *p && \result >= *q;  
    ensures \result == *p || \result == *q;  
*/  
int max_ptr (int *p, int *q);
```

Give a valid implementation that does not work properly...

We should say that the function cannot modify `*p` and `*q`.

Example 3 (frame conditions)

This is the completely specified program:

```
/*@ requires \valid(p) && \valid(q);
    ensures \result >= *p && \result >= *q;
    ensures \result == *p || \result == *q;
    assigns \nothing;
*/
int max_ptr (int *p, int *q) {
    if (*p >= *q) return *p;
    return *q;
}
```

- Lists of assigned variables explicitly included in contracts are called **frame conditions**.
- Avoids to state that for any unchanged global variable v , we have **ensures $\text{\old}(v) == v$**

Example 3 (memory safety)

- If we have a closer look to the assertions that WP adds in the `max_ptr` function comprising RTE verification, we can notice that there exists another version of the `\valid` predicate, denoted `\valid_read`.
- The predicate `\valid_read` indicates that a pointer can be dereferenced, but only to read the pointed memory.
- Try the following annotation:

```
/*@ requires \valid_read(p) && \valid_read(q);
    ensures \result >= *p && \result >= *q;
    ensures \result == *p || \result == *q;
*/
int max_ptr (int *p, int *q);
```

Behaviours

The specification can be done by cases.

- Global preconditions and postconditions applies to all cases.
- Behaviours define contracts in particular cases.
- For each case (**behavior**):
 - ▶ **assumes** clause defines the subdomain.
 - ▶ the behaviour's precondition is defined by **requires** clauses.
 - ▶ the behaviour's postcondition is defined by **ensures**, **assigns** clauses.
- **complete behaviors** states that given behaviors cover all cases.
- **disjoint behaviors** states that given behaviors do not overlap.

Example 4

Try **abs_behavior_init.c** and fix the problems.

```
#include <limits.h>
/*@ requires x > INT_MIN;
    assigns \nothing;

    behavior pos :
        assumes x > 0;
        ensures \result == x;

    behavior neg :
        assumes x < 0;
        ensures \result == -x;

    complete behaviors;
    disjoint behaviors;
*/
int abs (int x) {
    if (x >= 0) return x;
    return -x;
}
```

Example 4

This is the completely specified program:

```
#include <limits.h>
/*@ requires x > INT_MIN;
    assigns \nothing;

    behavior pos :
        assumes x >= 0;
        ensures \result == x;

    behavior neg :
        assumes x < 0;
        ensures \result == -x;

    complete behaviors;
    disjoint behaviors;
*/
int abs (int x) {
    if (x >= 0) return x;
    return -x;
}
```

Functions calls

Suppose function `g` contains a call to function `f`.

```
void g(...) {
    ...
    f(...);
    ...
}
```

Suppose we try to prove the caller `g`.

The function call is handled as follows:

- Before the call to `f` in `g`, the precondition of `f` must be ensured by `g`.
 - ▶ VCs are generated to prove that the precondition of `f` is respected.
- After the call to `f` in `g`, the postcondition of `f` is supposed to be true.
 - ▶ the postcondition of `f` is assumed in the proof below.
- Only a contract and a declaration of `f` are required.

Example 5

Run WP with `abs_calls.c` and see the problem.

```
#include <limits.h>
/*@ requires x > INT_MIN;
    ensures (x >= 0 ==> \result == x) &&
           (x < 0 ==> \result == -x);
    assigns \nothing;
*/
int abs (int x);

/*@ requires -1000 <= x <= 1000;
    requires -1000 <= y <= 1000;
    ensures \result >= 0;
    assigns \nothing;
*/
int sumabs(int x, int y){
    int a = abs(x);
    int b = abs(y);
    return a+b;
}

/*@ assigns \nothing;
void foo(int a){
    int b = abs(42);
    int c = abs(-50);
    int d = abs(a);          // False : "a" can be INT_MIN
}
}
```

Example 6

Try `max_abs_init.c` and fix the problems.

```
#include <limits.h>
/*@ requires x > INT_MIN;
    ensures (x >= 0 ==> \result == x) && (x < 0 ==> \result == -x);
    assigns \nothing; */
int abs (int x);

/*@ ensures \result >= x && \result >= y;
    ensures \result == x || \result == y;
    assigns \nothing; */
int max (int x, int y);

/*@ ensures \result >= x && \result >= -x && \result >= y && \result >= -y;
    ensures \result == x || \result == -x || \result == y || \result == -y;
    assigns \nothing; */
int max_abs(int x, int y){
    x = abs(x);
    y = abs(y);
    return max(x,y);
}
```

Example 6

This is a solution:

```
#include <limits.h>

/*@ requires x > INT_MIN;
    ensures (x >= 0 ==> \result == x) && (x < 0 ==> \result == -x);
    assigns \nothing;
*/
int abs (int x);

/*@ ensures \result >= x && \result >= y;
    ensures \result == x || \result == y;
    assigns \nothing ;
*/
int max (int x, int y);

/*@ requires x > INT_MIN;
    requires y > INT_MIN;
    ensures \result >= x && \result >= -x && \result >= y && \result >= -y;
    ensures \result == x || \result == -x || \result == y || \result == -y;
    assigns \nothing ;
*/
int max_abs(int x, int y){
    x = abs(x);
    y = abs(y);
    return max(x,y);
}
```

Example 7

- Run WP with **unref_init.c**

```
/*@ requires \valid(p);
*/
int unref(int* p){
    return *p;
}

int const value = 42;

int main(){
    int i = unref(&value);
}
```

- Dereferencing p is valid, however the precondition of unref will not be verified by WP since dereferencing value is only legal for a read-access.
- Fix the problem.

Example 8

Run WP with **proc_mem_init.c**

```
/*@ requires \valid(a) && \valid(b);
    ensures *a==10 && *b==20;
    assigns *a, *b;
*/
void proc(int *a, int *b) {
    *a = 10;
    *b = 20;
}
```

Why it fails?

Recall that WP memory model does not make any assumptions about memory regions, and they can overlap.

With this in mind, fix the problem.

Example 8 (solution)

```
/*@ requires \valid(a) && \valid(b);
    requires a != b;
    ensures *a==10 && *b==20;
    assigns *a, *b;
*/
void proc(int *a, int *b) {
    *a = 10;
    *b = 20;
}
```

Example 9

Write the ACSL specification corresponding to the following informal specification of function `find_array`.

`find_array(arr, len, x)` returns any index `idx` of the sorted array `arr` of length `len` such that `arr[idx] == x`. If such an index does not exist, it returns `-1`.

```
int find_array(int* arr, int len, int x);
```

Example 9

Here is a correct answer:

```
/*@ requires len >= 0;
    requires \valid(arr+(0..(len-1)));
    requires \forall integer i, j;
        0<=i<=j<len ==> arr[i]<=arr[j];
    ensures (\exists integer i; 0<=i<len && arr[i]==x)
        ==> 0<=\result<len && arr[\result]==x;
    ensures (\forall integer i; 0<=i<len ==> arr[i]!=x)
        ==> \result==-1;
    assigns \nothing;
*/
int find_array(int* arr, int len, int x);
```

Example 9

Here is a correct answer which defines two behaviours:

```
/*@ requires \forall integer i, j;
    0 <= i <= j < len ==> arr[i] <= arr[j];
    requires len >= 0;
    requires \valid(arr+(0..(len-1)));
    assigns \nothing;

    behavior belongs:
        assumes \exists integer i;
            0 <= i < len && arr[i] == x;
        ensures 0 <= \result < len;
        ensures arr[\result] == x;

    behavior not_belongs:
        assumes \forall integer i;
            0 <= i < len ==> arr[i] != x;
        ensures \result == -1;
*/
int find_array(int* arr, int len, int x);
```

Example 9 (logical predicates)

We can define two logical predicates:

- `sorted` which states that a given array is sorted
- `elem` which states that an element belongs to a given array

```
/*@ predicate sorted(int* arr, integer length) =
    \forall integer i, j; 0<=i<=j<length
        ==> arr[i]<=arr[j];

    predicate elem(int v, int* arr, integer length) =
        \exists integer i; 0<=i<length && arr[i]==v;
*/
```

Modify the previous specification to use these predicates.

Example 9 (using the logical predicates)

A correct answer, with behaviours, using the logical predicates defined:

```
/*@ requires sorted(arr, len);
    requires len >= 0;
    requires \valid(arr+(0..(len-1)));

    assigns \nothing;

    behavior belongs:
        assumes elem(x, arr, len);
        ensures 0 <= \result < len;
        ensures arr[\result] == x;

    behavior not_belongs:
        assumes ! elem(x, arr, len);
        ensures \result == -1;
*/
int find_array(int* arr, int len, int x);
```

Assert annotations

- **assert** p means that p must hold in the current state (the sequence point where the assertion occurs).
- **for** id_1, \dots, id_n : **assert** p associates the assertion to the named behaviours id_i . It means that this assertion must hold only for the considered behaviours.

Exemplo 10

Try **foo.assert_init.c** and fix the problems.

```
/*@ requires \valid_read(p) && \valid_read(q);
    ensures \result >= *p && \result >= *q;
    ensures \result == *p || \result == *q;
*/
int max_ptr (int *p, int *q);

void foo(){
    int a = 42;
    int b = 37;

    b += 10;
    int c = max_ptr(&a, &b);
    //@ assert c == 47;
}
```

State label mechanism

- Specification may require values at different program points.
- ACSL uses a **state label mechanism** that allows to refer to the value of a variable in any point of the program.
- Use $\backslash\text{at}(e, L)$ to refer to the value of expression e at label L .
- **Predefined logic labels:**
 - ▶ **Here** refers to the point where the annotation appears;
 - ▶ **Old** or **Pre** refers to the point before function call;
 - ▶ **Post** refers to the point after function call;
 - ▶ **LoopEntry** refers to the point at loop entry;
 - ▶ **LoopCurrent** refers to the point at the beginning of the current step of the loop.
- $\backslash\text{old}(e)$ is equivalent to $\backslash\text{at}(e, \text{Old})$, and $\backslash\text{at}(e, \text{Here})$ to e .

Exemplo 10

Check **foo_assert_at.c**.

```
/*@ requires \valid_read(p) && \valid_read(q);
    ensures \result >= *p && \result >= *q;
    ensures \result == *p || \result == *q;
    assigns \nothing;
*/
int max_ptr (int *p, int *q);

void foo(){
    int a = 42;
    int b = 37;

    Label_b:
    b += 10;
    int c = max_ptr(&a,&b);
    //@ assert c == 47;
    b = c+b;
    //@ assert b == 94 && \at(b,Label_b) == 37;
}
```

Program Verification

Example 11 (binary search)

A possible implementation of the specification given for **find_array** is

```
int find_array(int* arr, int len, int x) {
    int low = 0;
    int high = len - 1;
    while (low <= high) {
        int mean = (low + high) / 2;
        if (arr[mean] == x) return mean;
        if (arr[mean] < x) low = mean + 1;
        else high = mean - 1;
    }
    return -1;
}
```

- Check **binary_search_init.c**.
- Prove the correction of this implementation w.r.t. its specification.
 - ▶ 7 unknown VCs remain (all located inside loop or after loop)
- This pinpoints a classic difficulty: **reasoning about loops**

Loops and proof

- Main difficulty: to find appropriate **loop invariants** for each loop of the program.
- The invariants are **the only thing** that is known about the state of the program after the loop.
 - ▶ They must thus be **strong enough** to allow us to prove postconditions.
 - ▶ But **not too strong**, or we will not be able to prove the invariants themselves.
- The proof of loop invariants is done **by induction**.
 - ▶ it must hold before the loop (0 iterations)
 - ▶ it must hold after $k+1$ iterations whenever it holds after k iterations
- The VCs for a loop invariant include two parts
 - ▶ **loop invariant initially holds**
 - ▶ **loop invariant is preserved** by any iteration

Loop invariants

Loop invariants may be tricky.

How to find a suitable loop invariant?

- identify **variables modified** in the loop
 - ▶ define their possible value intervals (relationships) after k iterations
 - ▶ **loop assigns** clause can be used to list variables that (might) have been assigned so far after k iterations
- identify **realized actions**, or properties already ensured by the loop
 - ▶ what part of the job already realized after k iterations?
 - ▶ why the next iteration can proceed as it does?
- a stronger property on each iteration may be required to prove the final result of the loop

Some experience may be necessary to find appropriate loop invariants.

Example 11 (loop invariants)

Write loop invariants to prove the **safety properties** for **find_array**

The following invariants show that **low** and **high** are within **arr**'s bounds:

```
/*@ loop invariant 0 <= low;  
    loop invariant high < len;  
*/
```

There is **still an arithmetic overflow VC unknown**. Why? **Fix the problem!**

```
int mean = low + (high-low) / 2;
```

Example 11

Write the loop invariants that allow to prove the **postconditions of behaviours** for **find_array**.

```
/*@ loop invariant \forall integer i;  
    0 <= i < low ==> arr[i] < x;  
    loop invariant \forall integer i;  
    high < i < len ==> arr[i] > x;  
*/
```

There are still 2 safety VCs not proved which concerns to loop termination.

Loop termination

- Program termination is undecidable.
- For proving loop termination one has to give a **loop variant**.
 - ▶ A loop variant is an **upper bound on the number of remaining loop iterations**.
 - ▶ A loop variant is an **integer expression with a non-negative value which decreases on each iteration of the loop**.
- To find a variant, **look at the loop condition**.

Example 11 (loop variant)

Provide a loop variant that ensures that the loop always terminates.

```
//@ loop variant high - low + 1;
```

Example 11 (loop assigns)

- Considering loops, WP only reasons about what is provided by the user to perform its reasoning.
- In this example, the invariant does not specify anything about the way the variables are assigned.
- ACSL allows to add [assigns annotations for loops](#). Any other variable is considered to keep its old value.

```
//@ loop assigns low, high, mean;
```

- ▶ When the `loop assigns` clause is omitted, the VCGen assumes it is equal to the frame condition of the routine.

Example 11 (solution)

```
int find_array(int* arr, int len, int x) {
  int mean;
  int low = 0;
  int high = len - 1;
  /*@ loop invariant 0 <= low;
     loop invariant high < len;
     loop invariant \forall integer i;
       0 <= i < low ==> arr[i] < x;
     loop invariant \forall integer i;
       high < i < len ==> arr[i] > x;
     loop assigns low, high, mean;
     loop variant high - low + 1; */
  while (low <= high) {
    mean = low + (high-low) / 2;
    if (arr[mean] == x) return mean;
    if (arr[mean] < x) low = mean + 1;
    else high = mean - 1;
  }
  return -1;
}
```

Proof failures

A proof of a VC can fail for **various reasons**

- [erroneous implementation](#)
- [incorrect specification](#)
- [missing or erroneous \(previous\) annotation](#)
- [complexity of the proof](#)
 - ▶ try different provers
 - ▶ split the VC in independent properties
 - ▶ try a longer timeout
 - ▶ additional statements (assert, lemma, ...) may help the provers
 - ▶ if nothing else helps try an interactive proof assistant...

Exercises

Write a contract and prove the correctness of the following code.

```
void change(int *a, int *b) {  
  int tmp = *a;  
  *a = *b;  
  *b = tmp;  
}
```

See [change_init.c](#).

Exercises

For each of the following functions:

- tests if an array has negative values

```
int negs(int A[], int N);
```

- returns the index where the minimum of an array is

```
int minarray(int A[], int N);
```

- tests if the segments [a..b] of two different arrays are equal

```
int equal_seg(int A[], int B[], int a, int b, int N);
```

- returns an index which value is x, if it exists; -1 otherwise

```
int where(int A[], int N, int x);
```

- 1 Write a ACSL contract.
- 2 Write the function definition and prove its safety and functional correctness.
- 3 Write a main function that invokes it and check it.