

Introduction to C program proof with Frama-C and its WP plugin

January 17, 2020

Contents

1. Introduction	4
2. Program proof and our tool for this tutorial: Frama-C	6
2.1. Program proof	6
2.1.1. Ensure program reliability	6
2.1.2. A bit of context	7
2.1.3. Hoare triples	9
2.1.4. Weakest precondition calculus	10
2.2. Frama-C	11
2.2.1. Frama-C? WP?	11
2.2.2. Installation	12
2.2.3. Verify installation	14
2.2.4. (Bonus) Some more provers	16
3. Function contract	20
3.1. Contract definition	20
3.1.1. Postcondition	21
3.1.2. Precondition	28
3.1.3. Some elements about the use of WP and Frama-C	32
3.1.4. Exercises	32
3.2. Well specified function	35
3.2.1. Correctly write what we expect	35
3.2.2. Pointers	37
3.2.3. Writing the right contract	44
3.2.4. Exercises	44
3.3. Behaviors	48
3.3.1. Exercises	49
3.4. WP Modularity	52
3.4.1. Exercises	54
4. Basic instructions and control structures	58
4.0.1. Inference rules	58
4.0.2. Hoare triples	59
4.1. Basic concepts	60
4.1.1. Assignment	60
4.1.2. Composition of statements	64
4.1.3. Conditional rule	65
4.1.4. Bonus Stage - Consequence rule	67
4.1.5. Bonus Stage - Constancy rule	68
4.1.6. Exercices	69

4.2.	Loops	71
4.2.1.	Induction and invariant	71
4.2.2.	The assigns clause ... for loops	75
4.2.3.	Partial correctness and total correctness - Loop variant	76
4.2.4.	Create a link between postcondition and invariant	79
4.2.5.	Early termination of loop	80
4.2.6.	Exercises	82
4.3.	More examples on loops	84
4.3.1.	Example with read-only arrays	84
4.3.2.	Examples with mutable arrays	86
4.3.3.	Exercises	88
4.4.	Function calls	91
4.4.1.	Calling a function	91
4.4.2.	Recursive functions	95
5.	ACSL - Properties	98
5.1.	Some logical types	98
5.2.	Predicates	98
5.2.1.	Syntax	99
5.2.2.	Abstraction	101
5.2.3.	Exercises	102
5.3.	Logic functions	104
5.3.1.	Syntax	104
5.3.2.	Recursive functions and limits of logic functions	105
5.3.3.	Exercises	106
5.4.	Lemmas	109
5.4.1.	Syntax	109
5.4.2.	Example: properties of linear functions	110
5.4.3.	Example: arrays and labels	111
5.4.4.	Exercises	111
6.	ACSL - Logic definitions and code	116
6.1.	Inductive definitions	116
6.1.1.	Syntax	116
6.1.2.	Recursive predicate definitions	120
6.1.3.	Example: sort	121
6.1.4.	Exercises	124
6.2.	Axiomatic definitions	128
6.2.1.	Syntax	128
6.2.2.	Recursive function or predicate definitions	129
6.2.3.	Consistency	130
6.2.4.	Example: counting occurrences of a value	131
6.2.5.	Example: The strlen function	133
6.2.6.	Exercises	136
6.3.	Ghost code	138
6.3.1.	Syntax	138
6.3.2.	Ghost code validity	139
6.3.3.	Make a logical state explicit	140

6.3.4.	Exercises	142
6.4.	Hidden content	143
6.4.1.	Coq Proof of the <code>no_changes</code> lemma	143
6.4.2.	Specified sort functions	144
6.4.3.	An important axiom	144
6.4.4.	Sum axioms	145
7.	Proof methodologies	147
7.1.	Absence of runtime errors: Minimal contracts	147
7.1.1.	Principle	147
7.1.2.	Example: the search function	148
7.1.3.	Advantages and limitations	151
7.1.4.	Exercises	151
7.2.	Guiding assertions and triggering of lemmas	153
7.2.1.	Proof context	153
7.2.2.	Triggering lemmas	159
7.2.3.	A more complex example: sort, again	161
7.2.4.	How to correctly use assertions?	168
7.2.5.	Exercises	168
7.3.	More on ghost code: lemma functions and lemma macros	172
7.3.1.	Proof by induction	172
7.3.2.	Lemma function	175
7.3.3.	Lemma macro	179
7.3.4.	Limitations	183
7.3.5.	Back to the insertion sort	184
7.3.6.	Exercises	192
8.	Conclusion	197
8.1.	Going further	198
8.1.1.	With Frama-C	198
8.1.2.	With deductive proof	199

1. Introduction



This document is a beta version of the tutorial. The source code of the tutorial is available on GitHub, as well as the solutions to the different exercises (including the Coq proofs of some properties).

If you find some errors, please do not hesitate to post an issue or a pull request at: https://github.com/AllanBlanchard/tutoriel_wp



In this tutorial, some examples and some elements of organization are similar to the ones used in the [TAP 2013 tutorial](#) by Nikolai Kosmatov, Virgile Prevosto and Julien Signoles of the CEA LIST, since it is quite didactic. It also contains examples taken from [ACSL By Example](#) by Jochen Burghardt, Jens Gerlach, Kerstin Hartig, Hans Pohl and Juan Soto from the Fraunhofer. For formal aspects, I verified my statements and explanations using the course on Why3 given by Andrei Paskevich [at EJCP 2018](#). The remaining ideas come from my personal experience with Frama-C and WP.

The versions of the tools considered in this tutorial are:

- Frama-C 20 Calcium
- Why3 1.2.0
- Alt-Ergo 2.0.0
- Coq 8.9.0 (for provided scripts, not used in the tutorial)
- Z3 4.9.4 (in one example, it is not absolutely necessary to follow the course)

Depending on the versions used by the reader some differences could appear in what is proved and what is not. Some presented features are only available in recent versions of Frama-C.

The only requirement to this tutorial is to have a basic knowledge of the C language, and at least to be familiar with the notion of pointer.

1. Introduction

Despite its old age, C is still a widely used programming language. Indeed, no other language can pretend to be available on so many different (hardware and software) platforms, its low-level orientation and the amount of time invested in the optimization of its compilers allows to generate very light and efficient machine code (if the code allows it of course), and that there are a lot of experts in C language, which is an important knowledge base.

Furthermore, a lot of systems rely on a huge amount of code historically written in C, that needs to be maintained and sometimes fixed, as it would be far too costly to rewrite these systems.

But anyone who has already developed with C also knows that it is very hard to perfectly master this language. There are numerous reasons, but the complexity of the ISO C, and the fact that it is extremely permissive, especially about memory management, make the development of robust C program very hard, even for an experienced programmer.

However, the C language is often chosen for critical systems (avionics, railway, armament, ...) where it is appreciated for its good performances, its technological maturity and the predictability of its compilation.

In such cases, the needs in terms of test coverage of the source code become extremely high. Thus, the question “is our software tested enough?” becomes a question to which it is very hard to answer. Program proof can help us. Rather than testing all possible and (un)imaginable inputs of the program, we will *statically* and *mathematically* prove that there cannot be any problem at runtime.

The goal of this tutorial is to use Frama-C, a tool developed at CEA LIST, and WP, its deductive proof plugin, to learn the basics about C program proof. More than the use of the tool itself, the goal of this tutorial is to convince that it is more and more possible to write programs without any programming error, but also to sensitize to simple notions that allows to better understand and write programs.



Many thanks to the different beta-testers for their constructive feedback:

- Taurre [↗](#)
- barockobamo [↗](#)
- Vayel [↗](#)
- Aabu [↗](#)

I thank ZesteDeSavoir validators who helped me improve again the quality of this tutorial:

- Taurre [↗](#) (again)
- Saroupille [↗](#)

Finally, many thanks to Jens Gerlach for his help during the translation of this tutorial from French to English, and to Rafael Bachmann for his review and remarks.

2. Program proof and our tool for this tutorial: Frama-C

The goal of this first part is, in the first section, to introduce the idea of program proof without giving too much details, and then, in the second section, to give the necessary instructions to install Frama-C and some automatic provers that we will use in this tutorial.

2.1. Program proof

2.1.1. Ensure program reliability

It is often difficult to ensure that our programs have only correct behaviors. Moreover, it is already complex to establish good criteria that make us confident enough to say that a program works correctly:

- beginners simply “try” to use their programs and consider that these programs work if they do not crash (which is not a really good indicator in C language),
- more advanced developers establish some test cases for which they know the expected result and compare the output they obtain,
- most companies establish complete test bases, that cover as much code as they can ; which are systematically executed on their code. Some of them apply test driven development,
- in critical domains, such as aerospace, railway or armament, source code needs to be certified using standardized processes with very strict criteria about coding rules and code covering by the test.

In all these ways to ensure that a program produces only what we expect it to produce, a word appears to be common: *test*. We *try* different inputs in order to isolate cases that are problematic. We provide inputs that we *estimate to be representative* of the actual use of the program (note that unexpected use cases are often not considered whereas there are generally the most dangerous ones) and we verify that the results we get are correct. But we cannot test *everything*. We cannot try *every* combination of *every* possible input of a program. It is then quite hard to choose good tests.

The goal of program proof is to ensure that, for any input provided to a given program, if it respects the specification, then the program will only well-behave. However, since we cannot test everything, we will formally, mathematically, establish a proof that our software can only exhibit specified behaviors, and that runtime-errors are not part of these behaviors.

A very well-known quote from Dijkstra precisely express the difference between test and proof:

2. Program proof and our tool for this tutorial: Frama-C

Program testing can be used to show the presence of bugs, but never to show their absence!

Dijkstra

2.1.1.1. The developer's Holy Grail: the bug-free software

Every time we read news about attacks on computer systems, or viruses, or bugs leading to crashes in well known apps, there is always the one same comment “the bug-free/perfectly secure program does not exist”. And, if this sentence is quite true, it is a bit misunderstood.

First, there is a difference between safety and security. Loosely speaking, in security a malicious entity that can attack the system, while in safety, we want to verify when used in a conform way, the system well behaves. Thus, safety is a first before we can get security¹. Second, and more important, we do not really define what we mean by “bug-free”. Creating software always relies at least on two steps: we establish a specification of what we expect from the program and then we produce the source code of the program that must respect this specification. Both of these steps can lead to the introduction of errors.

In this tutorial, we will show how we can prove that an implementation verifies a given specification. But what are the arguments of program proof, compared to program testing? First, the proof is complete, it cannot forget some corner case if the behavior is specified (program test cannot be complete, being exhaustive would be far too costly). Second, the obligation to formally specify with a logic formulation requires to exactly understand what we have to prove about our program.

One could cynically say that program proof shows that “the implementation does not contain bugs that do not exist in the specification”. But, well, it is already a big step compared to “the implementation does not contain too many bugs that do not exist in the specification”, since for example it is generally not the fault of the specification if we try to access a nul pointer. Moreover, there also exist approaches that allow to analyze specifications to find errors or under-specified behaviors. For example, with model checking techniques, we can create an abstract model from the specification and produce the set of states that can be reached according to this model. By characterizing what is an error state, we can determine if reachable states are error states.

2.1.2. A bit of context

Formal methods, as we name them, allow in computer science to rigorously, mathematically, reason about programs. There exist a lot of formal methods that can take place at different levels from program design to implementation, analysis and validation, and for all systems that allow to manipulate information.

Here, we will focus on a method that allows to formally verify that our programs have only correct behaviors. We will use tools that are able to analyze a source code and to determine

¹Depending on the field you are working in, the term “safety” may have a very different sense. Namely, a *safe* system must be a system that can never cause injuries to human beings. And thus in this case, the opposite applies: we can never have safety if security is not guaranteed before.

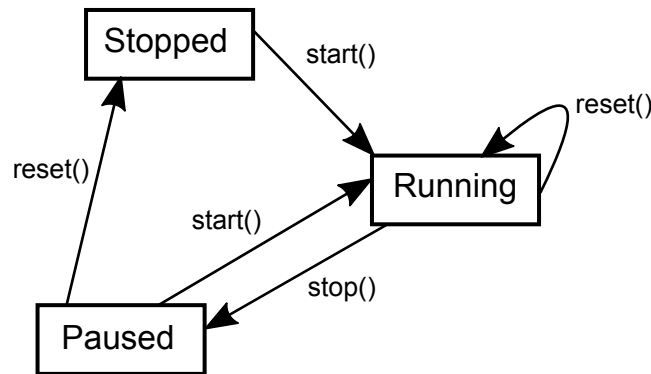
2. Program proof and our tool for this tutorial: Frama-C

whether a program correctly implements what we want to express. The analyzer we will use provides a static analysis, that we can oppose to dynamic analysis.

In static analysis, the analyzed program is not executed. We reason on a mathematical model of the states it can reach during its execution. On the opposite, dynamic analyses such as program testing, require to execute the analyzed source code. Note that there exist formal dynamic analysis methods, for example automatic test generation, or code monitoring techniques that allow to instrument a source code to verify some properties about it during execution (correct memory use, for example).

Talking about static analyses, the model we use can be more or less abstract depending on the techniques, this is always an approximation of possible states of the program. The more the approximation is precise, the more the model is concrete, the more the approximation is vague, the more it is abstract.

To illustrate the difference between concrete and abstract model, we can have a look at the model of a simple chronometer. A very abstract model of a chronometer could be the one presented here:



We have a model of the behavior of our chronometer with the different states it can reach according to the different actions we can perform. However, we have not modeled how these states are implemented in the program (is this a C enumeration? a particular program point in the source code?), nor how the computation of the elapsed time is done (a single variable? multiple ones?). It would then be difficult to specify properties about our program. We could add some information:

- State stopped at 0 : time = 0s
- State running : time > 0s
- State stopped : time > 0s

Which gives us a more concrete model but that is still not precise enough to ask interesting questions like: “is it possible for the program to continue updating the time variable while the chronometer is on the Stopped state?”, as we do not model how the time measurement is updated by the chronometer.

On the opposite, with the source code of the program, we have a concrete model of the chronometer. The source code expresses the behavior of the chronometer since it will allow us to produce the executable. But this is still not the more concrete model! For example, the

2. Program proof and our tool for this tutorial: Frama-C

executable in machine code format, that we obtain after compilation, is far more concrete than our program.

The more a model is concrete, the more it precisely describes the behavior of our program. The source code more precisely describes the behavior than our diagram, but it is less precise than the machine code. However, the more the model is precise, the more it is difficult to have a global view of the defined behavior. Our diagram is understandable in the blink of an eye, the source code requires more time, and for the executable ... Every single person that has already opened an executable with a text editor by error knows that it is not really pleasant to read¹.

When we create an abstraction of a system, we approximate it, in order to limit the knowledge we have about it and make our reasoning easier. A constraint we must respect, if we want our analysis to be correct, is to never under-approximate behaviors: we would risk to remove a behavior that contains an error (and thus miss it during the analysis). However, when we over-approximate the behaviors of the program, we can add behaviors that cannot happen, and if we add too many of them, we could not be able to prove that our program is correct, since some of them could be faulty behaviors.

In our case, the model is quite concrete. Every type of instruction, of control structure, is associated to a precise semantics, a model of its behavior in a pure logic, mathematical, world. The logic we use here is a variant of the Hoare logic, adapted to the C language and all its complex subtleties (which makes this model concrete).

2.1.3. Hoare triples

Hoare logic is a program formalization method proposed by Tony Hoare in 1969 in a paper entitled *An Axiomatic Basis for Computer Programming*. This method defines:

- axioms, that are properties we admit, such as “the skip action does not change the program state”,
- rules to reason about the different allowed combinations of actions, for example “the skip action followed by the action A” is equivalent to “the action A”.

The behavior of the program is defined by what we call “Hoare triples”:

$$\{P\} C \{Q\}$$

Where P and Q are predicates, logic formulas that express properties about the memory at particular program points. C is a list of instructions that defines the program. This syntax expresses the following idea: “if we are in a state where P is verified, after executing C and if C terminates, then Q is verified for the new state of the execution”. Put in another way, P is a sufficient precondition to ensure that C will bring us to the postcondition Q . For example, the Hoare triples that corresponds to the skip action is the following one:

$$\{P\} \text{ skip } \{P\}$$

¹ There also exists formal methods which are interested in understanding how executable machine code work, for example in order to understand what malwares do or to detect security breaches introduced during compilation.

2. Program proof and our tool for this tutorial: Frama-C

When we do nothing, the postcondition is the precondition.

Along this tutorial, we will present the semantics of different program constructs (conditional blocks, loops, etc) using Hoare logic. So, let us skip these details now since we will work on it later. It is not necessary to memorize these notions nor to understand all the theoretical background, but it is still useful to have some ideas about the way our tool works.

All of this gives us the basics that allow us to say “here is what this action does” but it does not give us anything to mechanize a proof. The tool we will use rely on a technique called weakest precondition calculus.

2.1.4. Weakest precondition calculus

The weakest precondition calculus is a form of predicate transform semantics proposed by Dijkstra in 1975 in *Guarded commands, non-determinacy and formal derivation of programs*.

This title can appear complex but actually the content of the article is in fact quite simple. We have seen before that Hoare logic gives us rules that explain the behavior of the different actions of a program, but it does not say how to apply these rules to establish a complete proof of the program.

Dijkstra reformulate the Hoare logic by explaining, in the triple $\{P\} C \{Q\}$, how the instruction, or the block of instructions, C transforms the predicate P in Q . This kind of reasoning is called *forward-reasoning*. We calculate from the precondition and from one or multiple instructions, the strongest postcondition we can reach. Informally, considering what we have in input, we calculate what we will get in output. If the postcondition we want is as strong or weaker, then we prove that there are no unexpected behaviors.

For example:

```
1  int a = 2;  
2  a = 4;  
3  //calculated postcondition: a == 4  
4  //expected postcondition : 0 <= a <= 30
```

Ok, 4 is an allowed value for **a**.

The form of predicate transformer semantics which we are interested in works the opposite way, we speak about *backward-reasoning*. From the wanted postcondition and the instructions we are reasoning about, we find the weakest precondition that ensures this behavior. If our actual precondition is at least as strong, that is to say, if it implies the computed precondition, then our program is correct.

For example, if we have the instruction:

$$\{P\} x := a \{x = 42\}$$

What is the weakest precondition to validate the postcondition $\{x = 42\}$? The rule will define that P is $\{a = 42\}$.

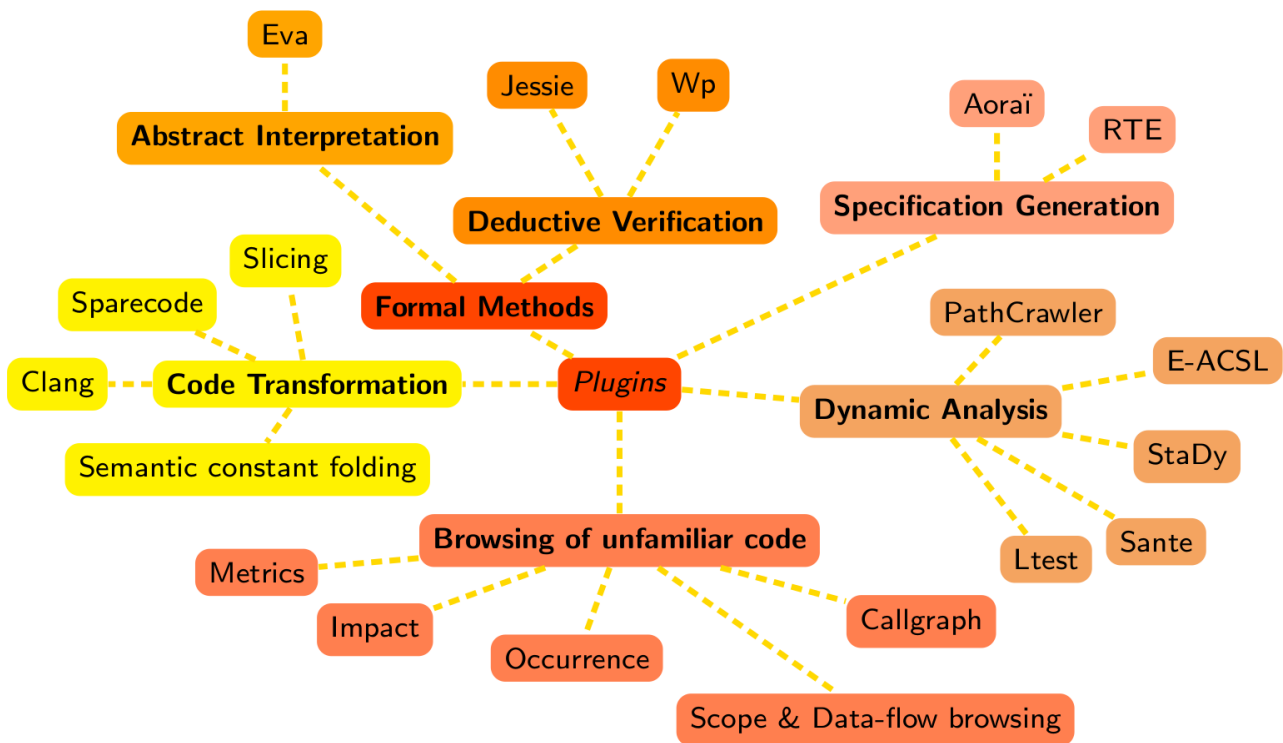
For now, let us forget about it, we will come back to these notions as we use them in this tutorial to understand how our tools work. So now, we can have a look at these tools.

2.2. Frama-C



2.2.1. Frama-C? WP?

Frama-C (FRAmework for Modular Analysis of C code) is a platform dedicated to the analysis of C programs created by the CEA LIST and Inria. It is based on a modular architecture allowing to use different (collaborating or not) plugins. The default plugins comprises different static analyses (that do not execute source code), dynamic analyses (that requires code execution), or combining both.



Frama-C provides a specification language called ACSL (“Axel”) for ANSI C Specification Language and that allows us to express the properties we want to verify about our programs. These properties are written using code annotations in comment sections. If one has already used Doxygen, it is quite similar, except that we write logic formulas and not text. During this tutorial, we will extensively write ASCL code, so let us just skip this for now.

The analysis we will use in this tutorial is provided by the WP plugin (for Weakest Precondition), a deductive verification plugin. It implements the technique we mentioned earlier: from ACSL annotations and the source code, the plugin generates what we call verification conditions, that are logic formulas that must be verified to be satisfiable or not. This verification can be performed manually or automatically, here we use automatic tools.

2. Program proof and our tool for this tutorial: Frama-C

We will use a SMT solver ([satisfiability modulo theory](#) [↗] , we do not detail how it works). This solver is [Alt-Ergo](#) [↗] , that was initially developed by the Laboratoire de Recherche en Informatique d’Orsay, and is today maintained by OCamlPro.

2.2.2. Installation

Frama-C is a software developed on Linux and OSX. Its support is thus better on those operating system. Nevertheless, it is possible to install Frama-C on Windows and in theory, its use will be identical to its use on Linux. However:



- the tutorial presents the use of Frama-C on Linux (or OSX) and the author did not experiment the differences that could exist with Windows,
- in recent versions of Windows 10, a possibility is to use Windows Subsystem for Linux, in combination with a Xserver installed on Windows to get GUI,
- the “Bonus” section of this part could not be accessible on Windows.

2.2.2.1. Linux

2.2.2.1.1. Using package managers On Debian, Ubuntu and Fedora, there exist packages for Frama-C. In such a case, it is enough to type a command like:

```
1 apt-get/yum install frama-c
```

However, these repositories are not necessarily up to date with the latest version of Frama-C. Most of the tutorial should be usable for the version available on your Linux distribution if it is recent enough. However, some specific features could be missing.

Go to the section “Verify installation” to perform some tests about the installation.

2.2.2.1.2. Via opam A second option is to use Opam, a package manager for Ocaml libraries and applications.

First of all, Opam must be installed (see its documentation). Then, some packages from the Linux distribution must be installed before installing Frama-C. On most systems, it is possible to ask Opam to install dependencies of the packages we want to install. For this, first install the `depext` tool of Opam:

```
1 opam install depext
```

And then use the `depext` tool to install the dependencies of Frama-C:

2. Program proof and our tool for this tutorial: Frama-C

```
1 opam depext frama-c
```

If the depext tool does not support your distribution, the following libraries and software should be installed:

- GTK2 (development library)
- GTKSourceview 2 (development library)
- GnomeCanvas 2 (development library)
- autoconf

On recent versions of some distributions, GTK2 might not be available, in this case, or if you want GTK3 and not GTK2, the packages GTK2, GTKSourceview2 and GnomeCanvas2 must be replaced with GTK3 and GTKSourceview 3.

Once these packages are installed, we can install Frama-C and Alt-Ergo.

```
1 opam install frama-c
2 opam install alt-ergo-free
```

Note that for versions older than Potassium (Argon and before), if you want to use the Why3 platform as a backend (introduced later), the version 0.88.3 should be used for Why3.

Go to the section “Verify installation” to perform some tests about the installation.

2.2.2.1.3. Via “manual” compilation The packages we have listed in the Opam section are required (of course, Opam itself is not). It requires a recent version of Ocaml and its compiler (including a compiler to native code). It is also necessary to install Why3 (at least version 1.2.0), which is available either on Opam or on their website ([Why3](http://why3.org/)).

After having extracted the folder available here : <http://frama-c.com/download.html> (Source distribution). Navigate to the folder and then execute the command line:

```
1 autoconf && ./configure && make && sudo make install
```

Go to the section “Verify installation” to perform some tests about the installation.

2.2.2.2. OSX

On OSX, the use of Homebrew and Opam is recommended to install Frama-C. The author does not use OSX, so here is a shameful copy and paste of the installation guide of Frama-C for OSX.

General Mac OS tools for OCaml:

2. Program proof and our tool for this tutorial: Frama-C

```
1 > xcode-select --install
2 > open http://brew.sh
3 > brew install autoconf opam
```

Graphical User Interface:

```
1 > brew install gtk+ --with-jasper
2 > brew install gtksourceview libgnomecanvas graphviz
3 > opam install lablgtk ocamlgraph
```

Recommended for Frama-C:

```
1 > brew install gmp
2 > opam install zarith
```

Necessary for Frama-C/WP:

```
1 > opam install alt-ergo-free
2 > opam install frama-c
```

2.2.2.3. Windows

Currently, the best way to install Frama-C on Windows is to use the Windows Subsystem for Linux. Basically, once the Linux system is installed, one should install Opam and follow the instructions provided in the Linux section. To get a usable graphical user interface, a X server should be installed on Windows.

2.2.3. Verify installation

In order to verify that the installation has been correctly performed, we will use the following code:

```
1 /*@
2   requires \valid(a) && \valid(b);
3   assigns *a, *b;
4   ensures *a == \old(*b);
5   ensures *b == \old(*a);
6 */
7 void swap(int* a, int* b){
8     int tmp = *a;
9     *a = *b;
10    *b = tmp;
11 }
12
13 int main(){
14     int a = 42;
```

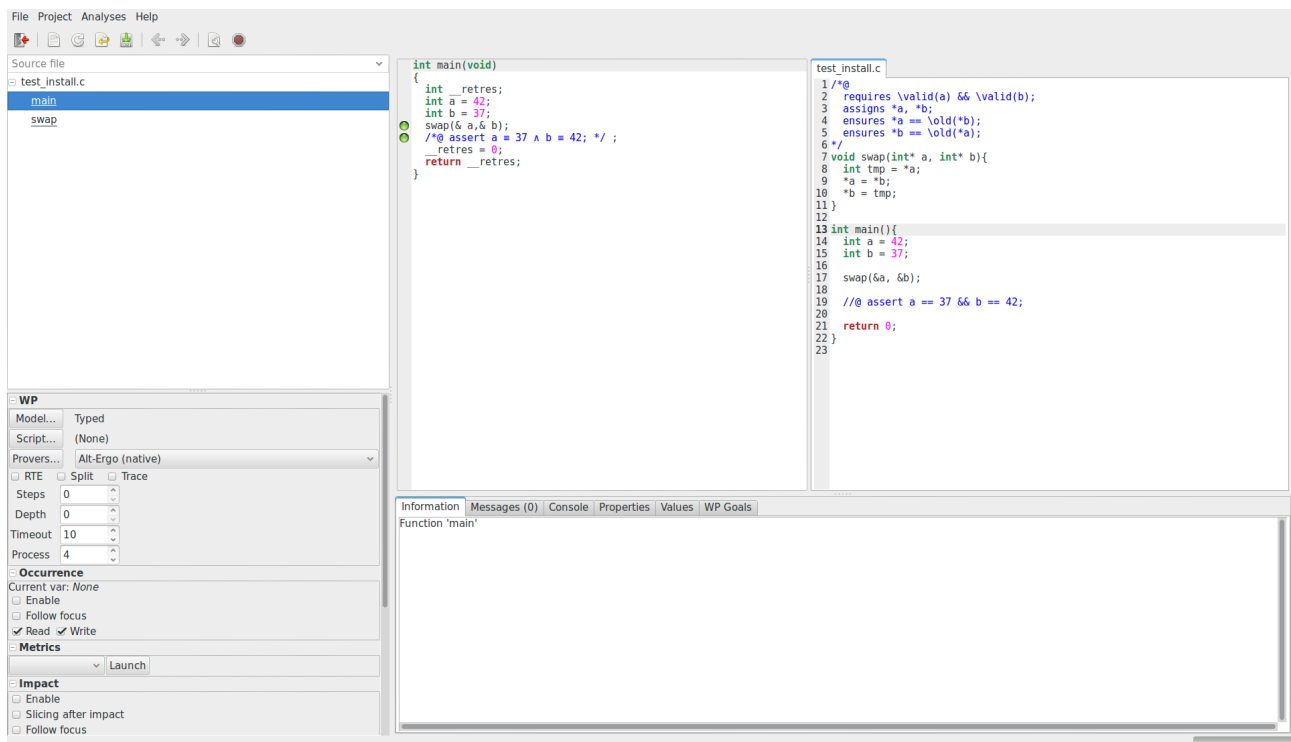
2. Program proof and our tool for this tutorial: Frama-C

```
15  int b = 37;
16
17  swap(&a, &b);
18
19  //@ assert a == 37 && b == 42;
20
21  return 0;
22 }
```

Copy and paste this code in a file named `main.c`. Then, from a terminal, in the folder where the file has been created, we start Frama-C with the following command line:

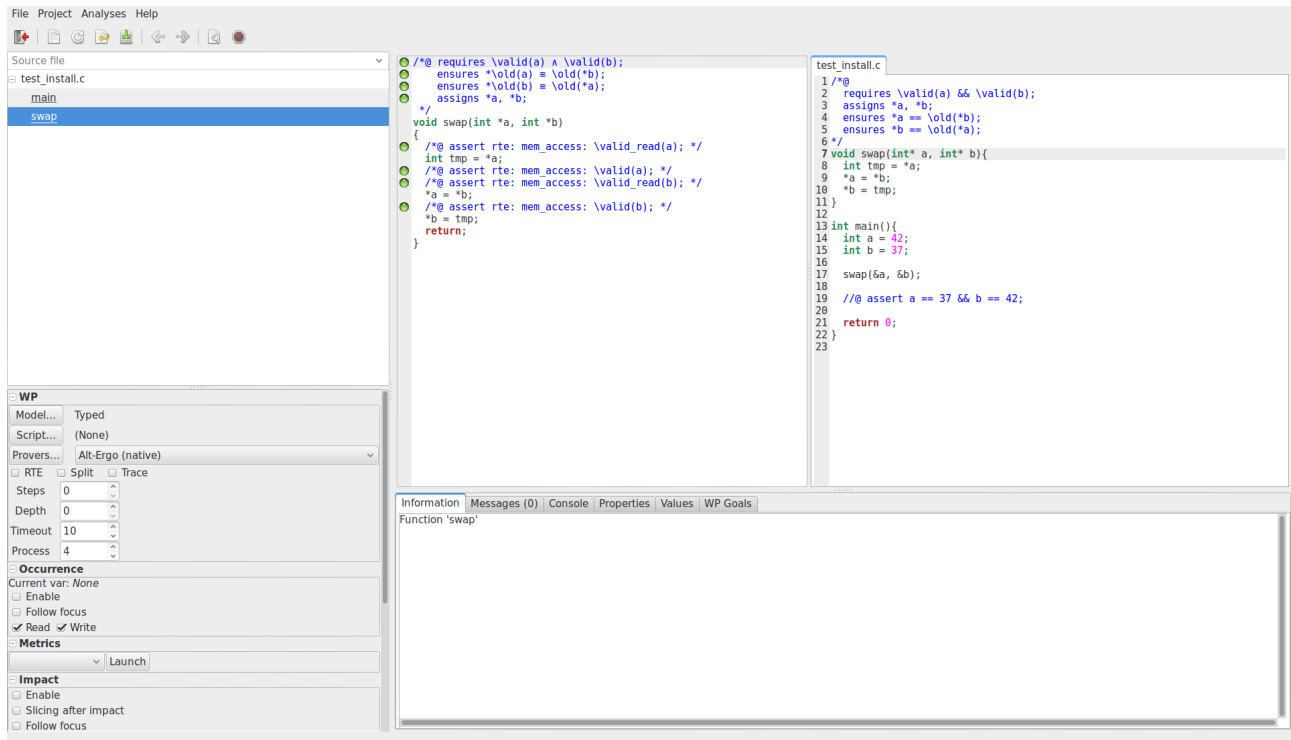
```
1 frama-c-gui -wp -rte main.c
```

The following window should appear:



Clicking `main.c` in the left side panel to select it, we can see its content (slightly) modified, and some green bullets on different lines as illustrated here:

2. Program proof and our tool for this tutorial: Frama-C



The graphical user interface of Frama-C does not allow source code edition.



For color blinds, it is possible to start Frama-C with another theme where color bullets are replaced:

```
1 frama-c-gui -gui-theme colorblind
```

2.2.4. (Bonus) Some more provers

This part is optional, nothing in this section should be particularly useful *in the tutorial*. However, when we start to be interested in proving more complex programs, it is often possible to reach the limits of Alt-Ergo, which is available by default, and we would thus need some other provers. For basic properties, almost all solvers have the same capabilities, for more complex ones, each solvers has its predilection domains.

2.2.4.1. Why3

[Why3](#) is a deductive proof platform developed by the LRI in Orsay. It provides a programming language and a specification language, as well as a module that allows to interact with a wide variety of automatic and interactive provers. This point is the one that interests us here. WP uses Why3 as a backend to talk with external provers.

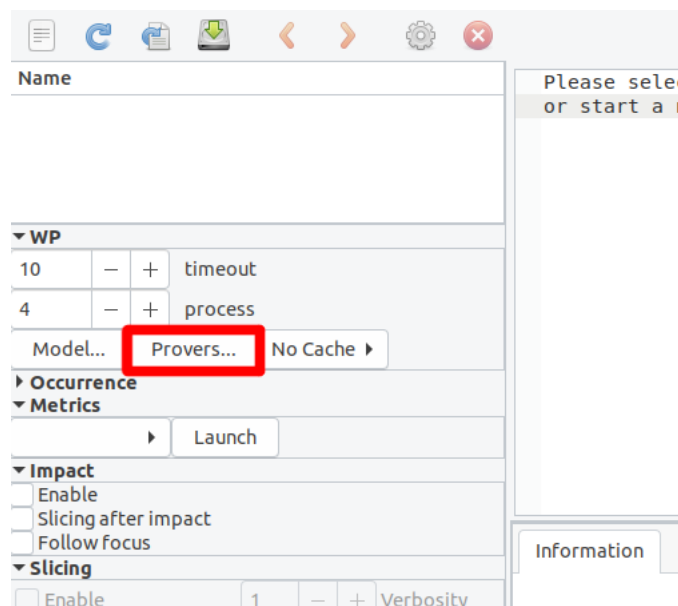
2. Program proof and our tool for this tutorial: Frama-C

On their website, we can find the list of [supported provers](#) . We recommend to install [Z3](#) which is developed by Microsoft Research, and [CVC4](#) which is developed by many research teams (New York University, University of Iowa, Google, CEA List). Those two provers are very efficient and somewhat complementary.

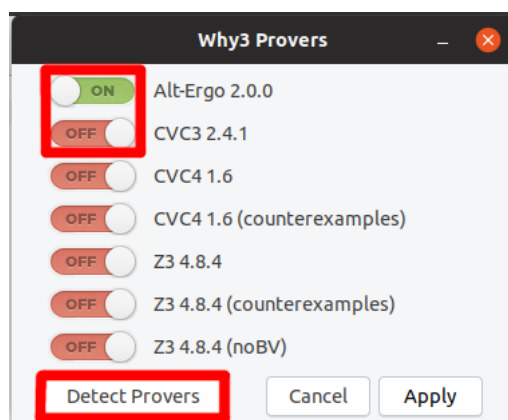
New provers can be installed any time after the installation of Frama-C. However, Why3 provers list must be updated:

```
1 why3 config --detect
```

And then activated in Frama-C. In the left side panel, in the WP part, clic "Provers..." :



And then "Detect" in the window that appears. Once it is done, provers can be activated thanks to the button next to their name.



2. Program proof and our tool for this tutorial: Frama-C

2.2.4.2. Coq

Coq, which is developed by Inria, is a proof assistant. Basically, we write the proofs ourselves in a dedicated language and the assistant verifies (using typing) that the proof is actually a valid proof.

Why would we need such a tool? Sometimes, the properties we want to prove can be too complex to be solved automatically by SMT solvers, typically when they requires careful inductive reasoning with precise choices at each step. In this situation, WP allows us to generate verification conditions translated in Coq language, and to write the proof ourselves.

To learn Coq, we would recommend [this tutorial](#) .

i

If Frama-C has been installed using the package manager of a Linux distribution, Coq could be automatically installed.

If one needs more information about Coq and its installation, this page can help: [The Coq Proof Assistant](#) .

The current support of Coq is deprecated in Frama-C but is still available. As we provide Coq scripts that can be used, we provide these instruction so that this way of proving programs can be tested by users. To run Frama-C with Coq support, use:

```
1  frama-c -wp-provers=native:coq
```

2. Program proof and our tool for this tutorial: Frama-C

Our tools are now installed and ready to be used.

The goal of this part, apart of the installation of our tools was to put in relief two main ideas:

- program proof is a way to ensure without executing them, that our programs only have correct behaviors, that are described by our specification,
- it is still our work to ensure that this specification is correct.

3. Function contract

It is time to enter the heart of the matter. Rather than starting with the basic notions of the C language, as we would do for a tutorial about C, we will start with functions. First because it is necessary to be able to write functions before starting this tutorial (to be able to prove that a code is correct, being able to write it correct is required), and then because it will allow us to directly prove some programs.

After this part about functions, we will on the opposite focus on simple notions like assignments or conditional structures, to understand how our tool really works.

In order to be able to prove that a code is valid, we first need to specify what we expect of it. Building the proof of our program consists in ensuring that the code we wrote corresponds to the specification that describes its job. As we previously said, Frama-C provides the ACSL language to let the developer write contracts about each function (but that is not its only purpose, as we will see later).

3.1. Contract definition

The goal of a function contract is to state the properties of the input that are expected by the function, and in exchange the properties that will be assured for the output. The expectation of the function is called the **precondition**. The properties of the output are called the **postcondition**.

These properties are expressed with ACSL. The syntax is relatively simple if one has already developed in C language since it shares most of the syntax of boolean expressions in C. However, it also provides:

- some logic constructs and connectors that do not exist in C, to ease the writing of specifications,
- built-in predicates to express properties that are useful about C programs (for example: a valid pointer),
- as well as some primitive types for the logic that are more general than primitive C types (for example: mathematical integer).

We will introduce along this tutorial a large part of the notations available in ACSL.

ACSL specifications are introduced in our source code using annotations. Syntactically, a function contract is integrated in the source code with this syntax:

3. Function contract

```
1  /*@
2    //contract
3  */
4  void foo(int bar){
5
6  }
```

Notice the `@` at the beginning of the comment block, this indicates to Frama-C that what follows are annotations and not a comment block that it should simply ignore.

Now, let us have a look at the way we express contracts, starting with postconditions, since it is what we want our function to do (we will later see how to express precondition).

3.1.1. Postcondition

The postcondition of a function is introduced with the clause `ensures`. Let us illustrate its use with the following function that returns the absolute value of an input value. One of its postconditions is that the result (which is denoted with the keyword `\result`) is greater or equals to 0.

```
1  /*@
2    ensures \result >= 0;
3  */
4  int abs(int val){
5    if(val < 0) return -val;
6    return val;
7  }
```

(Notice the `;` at the end of the line).

But this it is not the only property to verify. We also need to specify the general behavior of a function returning the absolute value. That is: if the value is positive or 0, the function returns the same value, else it returns the opposite of the value.

We can specify multiple postconditions, first by combining them with a `&&` as we do in C, or by introducing a new `ensures` clause, as we illustrate here:

```
1  /*@
2    ensures \result >= 0;
3    ensures (val >= 0 ==> \result == val) &&
4              (val < 0 ==> \result == -val);
5  */
6  int abs(int val){
7    if(val < 0) return -val;
8    return val;
9  }
```

This specification is the opportunity to present a very useful logic connector provided by ACSL and that does not exist in C: the implication $A \Rightarrow B$, that is written `A ==> B` in ACSL. The truth table of the implication is the following:

3. Function contract

A	B	$A \Rightarrow B$
F	F	T
F	T	T
T	F	F
T	T	T

That means that an implication $A \Rightarrow B$ is true in two cases: either A is false (and in this case, we do not check the value of B), or A is true and then B must also be true. Note that it means that $A \Rightarrow B$ is equivalent to $\neg A \vee B$. The idea finally being “I want to know if when A is true, B also is. If A is false, I don’t care, I consider that the complete formula is true”. For example, “if it rains, I want to check that I have an umbrella, if it does not, I do not care, everything is fine”.

Another available connector is the equivalence $A \Leftrightarrow B$ (written `A <==> B` in ACSL), and it is stronger. This is the conjunction of the implication in both ways $(A \Rightarrow B) \wedge (B \Rightarrow A)$. This formula is true in only two cases: A and B are both true, or both false (it can be seen as the negation of the exclusive or). Continuing with our example, “I do not only want to know that I have an umbrella when it rains, I also want to be sure that I have one only when it rains”.

i

Let us give a quick reminder about all truth tables of usual logic connectors in first order logic ($\neg =$ `!`, $\wedge =$ `&&`, $\vee =$ `||`) :

A	B	$\neg A$	$A \wedge B$	$A \vee B$	$A \Rightarrow B$	$A \Leftrightarrow B$
F	F	T	F	F	T	T
F	T	T	F	T	T	F
T	F	F	F	T	F	F
T	T	F	T	T	T	T

We can come back to our specification. As our files become longer and contains a lot of specifications, it can be useful to name the properties we want to verify. So, in ACSL, we can specify a name (without spaces) followed by a `:`, before stating the property. It is possible to put multiple levels of names to categorize our properties. For example, we could write this:

```
1 /*@
2   ensures positive_value: function_result: \result >= 0;
3   ensures (val >= 0 ==> \result == val) &&
4     (val < 0 ==> \result == -val);
5 */
6 int abs(int val){
7   if(val < 0) return -val;
8   return val;
9 }
```

3. Function contract

In most of this tutorial, we will not name the properties we want to prove, since they will be generally quite simple and we will not have too many of them, names would not give us much information.

We can copy and paste the function `abs` and its specification in a file `abs.c` and use Frama-C to determine if the implementation is correct with respect to the specification. We can start the GUI of Frama-C (it is also possible to use the command line interface of Frama-C but we will not use it during this tutorial) by using this command line:

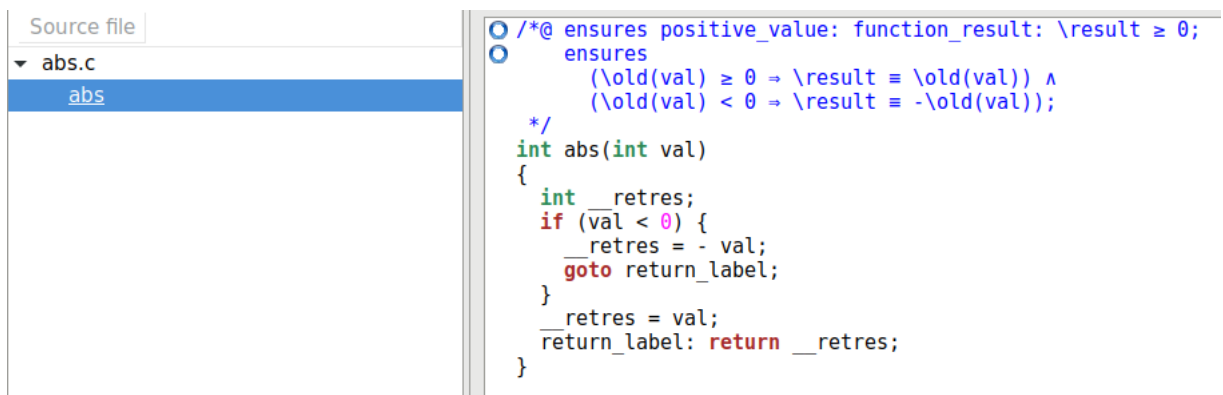
```
1 $ frama-c-gui
```

Or by opening it from the graphical environment.

It is then possible to click on the button “Create a new session from existing C files”, files to analyze can be selected by double-clicking it, the OK button ending the selection. Then, adding other files is done by clicking Files > Source Files.

Notice that it is also possible to directly open file(s) from the terminal command line passing them to Frama-C as parameter(s):


```
1 $ frama-c-gui abs.c
```



The window of Frama-C opens and in the panel dedicated to files and functions, we can select the function `abs`. At each `ensures` line, we can see a blue circle, it indicates that no verification has been attempted for these properties.

We ask the verification of the code by right-clicking the name of the function and “Prove function annotations by WP”:

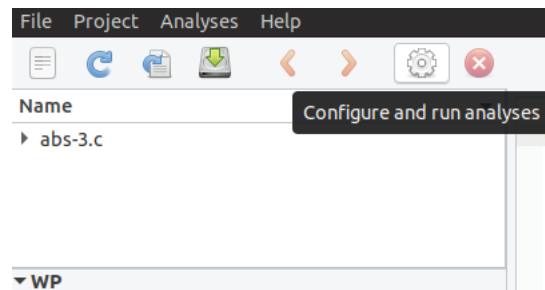
3. Function contract



We can see that blue circles become green bullets, indicating that the specification is indeed ensured by the program. We can also prove properties one by one by right-clicking on them and not on the name of the function.


But is our code really bug free? WP gives us a way to ensure that a code respects a specification, but it does not directly check the absence of runtime errors (RTE) if we do not ask it. Another plugin of Frama-C, called RTE, can be used to generate some ACSL annotations that can be verified by the other plugins. Its purpose is to add, in the program, some controls to ensure that the program cannot create runtime errors (integer overflow, invalid pointer dereferencing, 0 division, etc).

To activate these controls, we must activate this option in the configuration of WP. This is done by first clicking on the plugin configuration button:



And then adding the option `-wp-rte` in the options related to WP:

3. Function contract



We can also ask WP to add the RTE controls in a function by right-clicking on its name and then click “Insert wp-rte guards”.

Finally, we execute the verification again (we can also click on the “Reparse” button of the toolbar, it will delete existing proofs).


We can then see that WP fails to prove the absence of arithmetic overflow for the computation of `-val`. And, indeed, on our architectures, $-\text{INT_MIN} (-2^{31}) > \text{INT_MAX} (2^{31} - 1)$.

```

/*@ ensures positive_value: function_result: \result ≥ 0;
    ensures
        (\old(val) ≥ 0 ⇒ \result ≡ \old(val)) ∧
        (\old(val) < 0 ⇒ \result ≡ -\old(val));
*/
int abs(int val)
{
    int __retres;
    if (val < 0) {
        /*@ assert rte: signed_overflow: -2147483647 ≤ val; */
        __retres = - val;
        goto return_label;
    }
    __retres = val;
    return_label: return __retres;
}

```



We can notice that the overflow risk is real for us, since our computers (for which the configuration is detected by Frama-C) use the [Two's complement](#)  implementation of integers, which does not define the behavior of overflows.

Here, we can see another type of ACSL annotation. By the line `//@ assert property ;`, we can ask the verification of a property at a particular program point. Here, the RTE plugin inserted it for us, since we have to verify that `-val` does not produce an overflow, but we can also add such an assertion manually in the source code.

In this screenshot, we can see two new colors for our bullets: green+brown and orange.

If the proof has not been entirely redone after adding the runtime error checks, these bullets must still be green. Indeed, the corresponding proofs have been realized without the knowledge of the property in the assertion, so they cannot rely on this unproved property.

When WP transmits a verification condition to an automatic prover, it basically transmits two kinds of properties : G , the goal, the property that we want to prove, and $A_1 \dots A_n$, the different assumptions we can have about the state of the memory at the program point where we want to verify G . However, it does not receive (in return) the properties that have been used by the prover to validate G . So, if A_3 is an assumption, and if WP did not succeed in getting a proof of A_3 , it indicates that G is true, but only assuming that A_3 is true, for which no proof has been established so far.

The orange color indicates that no prover could determine if the property is verified. There are two possible reasons:

- the prover did not have the information needed to terminate the proof,
- the prover did not have enough time to compute the proof and met a timeout (which can be configured in the WP panel).

In the bottom panel, we can select the “WP Goals” tab, it shows the list of verification conditions, and for each prover the result is symbolized by a logo that indicates if the proof has been tried and if it succeeded, failed or met a timeout. Note that it may require to select “All Results” in the squared field to see all verification conditions.

3. Function contract

```

/*@ ensures positive_value: function_result: \result ≥ 0;
ensures
  (\old(val) ≥ 0 → \result ≡ \old(val)) ∧
  (\old(val) < 0 → \result ≡ -\old(val));
*/
int abs(int val)
{
  int __retres;
  if (val < 0) {
    /*@ assert rte: signed_overflow: -2147483647 ≤ val; */
    retres = - val;
    goto return_label;
  }
  __retres = val;
  return_label: return __retres;
}

```

Module	Goal	Model	Qed	Script	Alt-Ergo 2.0.0
abs	Post-condition 'positive_value,function_result'	Typed	●	▶	
abs	Post-condition	Typed	●	▶	
abs	Assertion 'rte,signed_overflow'	Typed	—	—	✂

In the first column, we have the name of the function the verification condition belongs to. The second column indicates the name of the verification condition. For example here, our postcondition is named `postcondition 'positive_value,function_result'`, we can notice that if we select a property in this list, it is also highlighted in the source code. Unnamed properties are automatically named by WP with the kind of wanted property. In the third column, we see the memory model that is used for the proof, we will not talk about it in this tutorial. Finally, the last columns represent the different provers available through WP.

In the list of provers, the first element is Qed. This is not really a prover. It is a tool used by WP to simplify properties before sending them to external provers. Then we find Script, this is a way to finish some proofs by hand when automatic solvers do not succeed in doing them. Finally, we find Alt-Ergo which is one of the solvers we use. Notice that for the RTE property, some scissors are indicated: it means that the solver was stopped due to timeout.

If we double-click on the property “absence of overflow” (highlighted in blue in the last screenshot), we can see the corresponding verification condition (if it is not the case, make sure that the value “Raw obligation” is selected in the blue squared field):

Goal Assertion 'rte,signed_overflow':
 Assume { Type: is_sint32(val 0). (* Then *) Have: val_0 < 0. }
 Prove: (-2147483647) <= val_0.

Prover z3: Timeout (Qed:4ms) (10s).
 Prover Alt-Ergo: Unknown (Qed:4ms) (52ms).

3. Function contract

This is the verification condition generated by WP about our property and our program, we do need to understand everything here, but we can get the general idea. It contains (in the “Assume” part) the assumptions that we have specified and those that have been deduced by WP from the instructions of the program. It also contains (in the “Prove” part) the property that we want to verify.

What does WP do using these properties? In fact, it transforms them into a logic formula and then asks to different provers if it is possible to satisfy this formula (to find for each variable, a value that can make the formula true), and it determines if the property can be proved. But before sending the formula to provers, WP uses a module called Qed, which is able to perform different simplifications about it. Sometimes, as this is the case for the other properties about `abs`, these simplifications are enough to determine that the property is true, in such a case, WP do not need the help of the automatic solvers.

When automatic solvers cannot ensure that our properties are verified, it is sometimes hard to understand why. Indeed, provers are generally not able to answer something other than “yes”, “no” or “unknown”, they are not able to extract the reason of a “no” or an “unknown”. There exist tools that can explore a proof tree to extract this kind of information, currently Frama-C does not provide such a tool. Reading verification conditions can sometimes be helpful, but it requires a bit of practice to be efficient. Finally, one of the best way to understand the reason why a proof fails is to try to do it interactively with Coq. However, it requires to be quite comfortable with this language to be able to understand the verification conditions generated by WP, since these conditions need to encode some elements of the C semantics that can make them quite hard to read.

If we go back to our view of the verification conditions (see the red squared button in the previous screenshot), we can see that our hypotheses are not sufficient to determine that the property “absence of overflow” is true (which is actually impossible), so we need to add some hypotheses to guarantee that our function will well-behave: a precondition.

3.1.2. Precondition

Preconditions are introduced using `requires` clauses. As we could do with `ensures` clauses, we can compose logic expressions and specify multiple preconditions:

```
1 /*@
2   requires 0 <= a < 100;
3   requires b < a;
4 */
5 void foo(int a, int b){
6
7 }
```

Preconditions are properties about the input (or about global variables) that we assume to be true when we analyze the function. We will verify that they are indeed true only at program points where the function is called.

In this small example, we can also notice a difference with C in the writing of boolean expressions. If we want to specify that `a` is between 0 and 100, we do not have to write

3. Function contract

`0 <= a && a < 100`, we can directly write `0 <= a < 100` and Frama-C will perform necessary translations.

If we come back to our example about the absolute value, to avoid the arithmetic overflow, it is sufficient to state that `val` must be strictly greater than `INT_MIN` to guarantee that the overflow will never happen. Thus, we add it as a precondition of the function (notice that it is also necessary to include the header where `INT_MIN` is defined):

```
1 #include <limits.h>
2
3 /*@
4  requires INT_MIN < val;
5
6  ensures \result >= 0;
7  ensures (val >= 0 ==> \result == val) &&
8          (val < 0 ==> \result == -val);
9 */
10 int abs(int val){
11     if(val < 0) return -val;
12     return val;
13 }
```



Reminder: The Frama-C GUI does not allow source code modification.

Once we have modified the source code with our precondition, we click on “Reparses” and we can ask again to prove our program. This time, everything is validated by WP, our implementation is proved:

```
● /*@ requires val > -2147483647 - 1;
● ensures positive_value: function_result: \result ≥ 0;
● ensures
    (\old(val) ≥ 0 ⇒ \result == \old(val)) ∧
    (\old(val) < 0 ⇒ \result == -\old(val));
*/
int abs(int val)
{
    int __retres;
    if (val < 0) {
        ● /*@ assert rte: signed_overflow: -2147483647 ≤ val; */
        __retres = - val;
        goto return_label;
    }
    __retres = val;
    return_label: return __retres;
}
```

We can also verify that a function that would call `abs` correctly respects the required precondition:

```
1 void foo(int a){
2     int b = abs(42);
3     int c = abs(-42);
4     int d = abs(a); // False : "a" can be INT_MIN
5     int e = abs(INT_MIN); // False : the parameter must be strictly greater than INT_MIN
6 }
```

3. Function contract

```
void foo(int a)
{
  int b = abs(42);
  int c = abs(-42);
  int d = abs(a);
  int e = abs(-2147483647 - 1);
  return;
}
```

Note that we can click on the bullet next to the function call to see the list of preconditions and check which ones are not validated. Here, there is only one precondition, but when there are multiple ones it is useful to check what is exactly the problem.

```
void foo(int a)
{
  int b = abs(42);
  int c = abs(-42);
  /* preconditions of abs:
     requires -2147483647 - 1 < a; */
  int d = abs(a);
  int e = abs(-2147483647 - 1);
  return;
}
```

We can modify this example by reverting the last two instructions. If we do this, we can see that the call `abs(a)` is validated by WP if it is placed after the call `abs(INT_MIN)` ! Why?

We must keep in mind that the idea of the deductive proof is to ensure that if preconditions are verified, and if our computation terminates, then the postcondition is verified.

If we give to a function an input that surely breaks the precondition, we can deduce that everything can happen, including obtain false in postcondition. More precisely, here, after the call, we just suppose that the precondition is still true (as the function does not modify anything in memory), thus we suppose that `INT_MIN < INT_MIN` which is obviously false. Knowing this, we can prove absolutely everything because this “false” becomes an assumption of every call that follows. Knowing false, we can prove everything, because if we have a proof of false, then false is true, as well as true is true. So everything is true.

Taking our modified program, we can convince ourselves of this fact by looking at the verification conditions generated by WP for the bad call and the subsequent call that becomes verified:

3. Function contract

The top screenshot shows the QED GUI with the following code in the editor:

```
void foo(int a)
{
  int b = abs(42);
  int c = abs(-42);
  int e = abs(-2147483647 - 1);
  /* preconditions of abs:
     requires -2147483647 - 1 < a; */
  int d = abs(a);
  return;
}
```

The bottom panel shows the 'WP Goals' tab with the following information:

- Goal Instance of 'Pre-condition' (call 'abs'):
- Prove: false.
- Prover Alt-Ergo: Unknown (53ms).

The bottom screenshot shows the same code, but the goal is now 'Proved Goal'.

The bottom panel shows the 'WP Goals' tab with the following information:

- Goal Instance of 'Pre-condition' (call 'abs'):
- Prove: true.
- Prover Qed: Valid.

We can notice that for function calls, the GUI highlights the execution path that leads to the call for which we want to verify the precondition. Then, if we have a closer look at the call `abs(INT_MIN)`, we can notice that, simplifying, Qed deduced that we try to prove “False”. Consequently, the next call `abs(a)` receives in its assumptions the property “False”. This is why Qed can immediately deduce “True”.

The second part of the question is then: why our first version of the calling function (`abs(a)` and then `abs(INT_MIN)`) did not have the same behavior, indicating a proof failure on the second call? The answer is simply that for the call `abs(a)` can we add the assumption `a < INT_MIN`, and while we do not have a proof that it is true, we do not have a proof that is false neither. So, while `abs(INT_MIN)` necessarily gives us the knowledge of “false”, the call `abs(a)` does not, since it can succeed.

3. Function contract

Produce a correct specification is then crucial. Typically, by stating false precondition, we can have the possibility to create a proof of false:

```
1  /*@
2    requires a < 0 && a > 0;
3    ensures  \false;
4  */
5  void foo(int a){
6
7  }
```

If we ask WP to prove this function, it will accept it without any problem since the assumption we give in precondition is necessarily false. However, we will not be able to give an input that respects the precondition so we will be able to detect this problem by carefully reading what we have specified.

Some notions we will see in this tutorial can expose us to the possibility to introduce subtle incoherence. So, we must always be careful specifying a program.

3.1.3. Some elements about the use of WP and Frama-C

In the two preceding sections, we have seen a lot of notions about the use of the GUI to start proofs. In fact, we can ask WP to immediately prove everything at Frama-C's startup with the option `-wp`:

```
1  $ frama-c-gui file.c -wp
```

Which will collect all properties to be proved in `file.c`, generate all verification conditions and try to discharge them.

About runtime-errors, it is generally advised to first verify the program without generating RTE assertions, and then to generate them to terminate the verification with WP. It allows WP to “focus” on the functional properties in a first step without having in its knowledge base purely technical properties, that are generally not useful for the proof of functional properties. Again, it is possible to directly produce this behavior using the command line:

```
1  $ frama-c-gui file.c -wp -then -wp -wp-rte
```

“Start Frama-C with WP, then create assertions to verify the absence of RTE and start WP again”.

3.1.4. Exercises

While these exercises are not absolutely necessary to read the next chapters of the tutorial we strongly suggest to practice them. Note that we also suggest to, at least, read the fourth exercise that introduces a notation that can be sometimes useful.

3. Function contract

3.1.4.1. Addition

Write the postcondition of the following addition function:

```
1 int add(int x, int y){  
2     return x+y ;  
3 }
```

And run the command:

```
1 frama-c-gui your-file.c -wp
```

Once the function is successfully proved to respect the contract, run:

```
1 frama-c-gui your-file.c -wp -wp-rte
```

It should fail, adapt the contract by adding the right precondition.

3.1.4.2. Distance

Write the postcondition of the following distance function, by expressing the value of `b` in terms of `a` and `\result`:

```
1 int distance(int a, int b){  
2     if(a < b) return b - a ;  
3     else return a - b ;  
4 }
```

And run the command:

```
1 frama-c-gui your-file.c -wp
```

Once the function is successfully proved to respect the contract, run:

```
1 frama-c-gui your-file.c -wp -wp-rte
```

It should fail, adapt the contract by adding the right precondition.

3. Function contract

3.1.4.3. Alphabet Letter

Write the postcondition of the following function that return true if the character received in input is an alphabet letter. Use the equivalence operator `<==>`.

```
1  int alphabet_letter(char c){
2      if( ('a' <= c && c <= 'z') || ('A' <= c && c <= 'Z') ) return 1 ;
3      else return 0 ;
4  }
5
6  int main(){
7      int r ;
8
9      r = alphabet_letter('x') ;
10     //@ assert r ;
11     r = alphabet_letter('H') ;
12     //@ assert r ;
13     r = alphabet_letter(' ') ;
14     //@ assert !r ;
15 }
```

And run the command:

```
1  frama-c-gui your-file.c -wp -wp-rte
```

All verification conditions should be proved, including the assertions in the main function.

3.1.4.4. Days of the month

Write the postcondition of the following function that returns the number of days in function of the received month (notice that we consider that they are numbered from 1 to 12), for February, we only consider the case when it has 28 days, we will see later how to solve this problem:

```
1  int day_of(int month){
2      int days[] = { 31, 28, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31 } ;
3      return days[month-1] ;
4  }
```

And run the command:

```
1  frama-c-gui your-file.c -wp
```

Once the postcondition is proved, run the command:

```
1  frama-c-gui your-file.c -wp -wp-rte
```

If it fails, complete the precondition in order to solve the problem.

3. Function contract

You might have notice that writing the postcondition can be a bit painful. Let us try to simplify that. ACSL provide the notion of set, and the operator `\in` that can be used to check whether a value is in a set or not.

For example:

```
1 //@ assert 13 \in { 1, 2, 3, 4, 5 } ; // FALSE
2 //@ assert 3  \in { 1, 2, 3, 4, 5 } ; // TRUE
```

Modify the postcondition by using this notation.

3.1.4.5. Last angle of a triangle

This function receives two values of angle in input and returns the value of the last angle considering that the sum of the three angles must be 180. Write the postcondition that expresses that the sum of the three angles is 180.

```
1 int last_angle(int first, int second){
2     return 180 - first - second ;
3 }
```

And run the command:

```
1 frama-c-gui your-file.c -wp
```

Once the function is successfully proved to respect the contract, run:

```
1 frama-c-gui your-file.c -wp -wp-rte
```

If it fails, add the right precondition. Note that the values of the different angles should not be more than 180 degrees.

3.2. Well specified function

3.2.1. Correctly write what we expect

This is certainly the hardest part of our work. Programming is already an effort that consists in writing algorithms that correctly answer to our need. Specifying requests the same kind of work, except that we do not try to express *how* we answer to our need but *what* is exactly our need. To prove that our code implements what we need, we must be able to describe exactly what we need.

From now, let us use an other example, the `max` function:

3. Function contract

```
1 int max(int a, int b){
2     return (a > b) ? a : b;
3 }
```

The reader could write and prove their own specification. Let us use this one:

```
1 /*@
2     ensures \result >= a && \result >= b;
3 */
4 int max(int a, int b){
5     return (a > b) ? a : b;
6 }
```

If we ask WP to prove this code, it succeeds without any problem. However, is our specification really what we need? We can try to prove this calling code:

```
8 void foo(){
9     int a = 42;
10    int b = 37;
11    int c = max(a,b);
12
13    //@assert c == 42;
14 }
```

There, it fails. In fact, we can go further by modifying the body of the `max` function and notice that the following code is also correct with respect to the specification:

```
1 #include <limits.h>
2
3 /*@
4     ensures \result >= a && \result >= b;
5 */
6 int max(int a, int b){
7     return INT_MAX;
8 }
```

While being a correct specification of `max`, our specification is however too permissive. We have to be more precise. We do not only expect the result to be greater or equal to both parameters, but also that the result is one of them:

```
1 /*@
2     ensures \result >= a && \result >= b;
3     ensures \result == a || \result == b;
4 */
5 int max(int a, int b){
6     return (a > b) ? a : b;
7 }
```

3. Function contract

This specification can also be proved correct by WP, but now we can also prove the assertion in our function `foo`, and we cannot prove anymore an implementation that would just return the value `INT_MAX`.

3.2.2. Pointers

If there is one notion that we permanently have to confront with in C language, this is definitely the notion of pointer. Pointers are quite hard to manipulate correctly, and they still are the main source of critical bugs in programs, so they benefit of a preferential treatment in ACSL.

We can illustrate with a swap function for C integers:

```
1  /*@
2   ensures *a == \old(*b) && *b == \old(*a);
3  */
4  void swap(int* a, int* b){
5      int tmp = *a;
6      *a = *b;
7      *b = tmp;
8  }
```

3.2.2.1. History of values in memory

Here, we introduce a first built-in logic function of ACSL: `\old`, that allows us to get the old (that is to say, before the call) value of a given element. So, our specification defines that the function must ensure that after the call, the value of `*a` is the old value of `*b` and conversely.

The `\old` function can only be used in the postcondition of a function. If we need this kind of information somewhere else, we use `\at` that allows us to express that we want the value of a variable at a particular program point. This function receives two parameters. The first one is the variable (or memory location) for which we want to get its value and the second one is the program point (as a C label) that we want to consider.

For example, we could write:

```
2  int a = 42;
3  Label_a:
4  a = 45;
5
6  //@assert a == 45 && \at(a, Label_a) == 42;
```

Of course, we can use any C label in our code, but we also have 6 built-in labels defined by ACSL that can be used:

- `Pre` / `Old` : value before function call,
- `Post` : value after function call,
- `LoopEntry` : value at loop entry

3. Function contract

- `LoopCurrent` : value at the beginning of the current step of the loop,
- `Here` : value at the current program point.

i

The behavior of `Here` is, in fact, the default behavior when we consider a variable. Its use with `\at` generally allows us to ensure that what we write is not ambiguous, and is more readable, when we express properties about values at different program points in the same expression.

Whereas `\old` can only be used in function postconditions, `\at` can be used anywhere. However, we cannot use any program point with respect to the type annotation we are writing. `Old` and `Post` are only available in function postconditions, `Pre` and `Here` are available everywhere. `LoopEntry` and `LoopCurrent` are only available in the context of loops (which we will detail later in this tutorial).

Note that one must take care to use `\old` and `\at` for values that make sense. This is why for example in a contract, all values received in input are put into `\old` when used in postcondition, the “new” value of the input variables not make any sense for the caller of the function as they are not accessible: they are local to the called function. For example, if we have look at the contract of the swap function transformed by Frama-C, we can see that in the postcondition, the pointers are enclosed into `\old` :

```
/*@ requires \valid(a) ^ \valid(b);
   ensures *\old(a) == \old(*b) ^ *\old(b) == \old(*a);
   */
void swap(int *a, int *b)
```

For the built-in `\at`, we have to take care of this more explicitly. In particular, the specified label must make sense with respect to the scope of the value. For example, in the following program, Frama-C detects that we ask the value of the variable `x` at a program point where it does not exist:

```
1 void example_1(void){
2   L: ;
3   int x = 1 ;
4   //@ assert \at(x, L) == 1 ;
5 }
```

Console

```
[kernel] Parsing at-2.c (with preprocessing)
[kernel:annot-error] at-2.c:6: Warning:
  unbound logic variable x. Ignoring code annotation
[kernel] User Error: warning annot-error treated as fatal error.
[kernel] User Error: stopping on file "at-2.c" that has errors. Add '-kernel-msg-key pp'
  for preprocessing command.
```

Cancel

3. Function contract

However, in some other cases, we only reach a proof failure since determining that the value does not exist at some particular label cannot be done by a syntactic analysis. For example, if the variable is declared but undefined or if we want the value of a pointed value:

```
7 void example_2(void){
8   int x ;
9   L:
10  x = 1 ;
11  //@ assert \at(x, L) == 1 ;
12 }
13
14 void example_3(void){
15   L: ;
16   int x = 1 ;
17   int *ptr = &x ;
18   //@ assert \at(*\at(ptr, Here), L) == 1 ;
19 }
```

Here, it is easy to see the problem. However, the considered label is propagated to subexpressions. Thus, sometimes, terms that seem to be innocent can have a surprising behavior if we do not keep this fact in mind. For example, in the following example:

```
21 /*@ requires x + 2 != p ; */
22 void example_4(int* x, int* p){
23   *p = 2 ;
24   //@ assert x[2] == \at(x[2], Pre) ;
25   //@ assert x[*p] == \at(x[*p], Pre) ;
26 }
```

The first assertion is proved, and while the second assertion seems to express the same property, it cannot be proved. Because, it in fact does not express the same property. The expression `\at(x[*p], Pre)` must be seen as `\at(x[\at(*p)], Pre)` as the label is propagated to the subexpression `*p`, for which we do not know the value at label `Pre` (since it is not specified).

For the moment, we do not need `\at` but it can often be useful, if not essential, when we want to make our specification precise.

3.2.2.2. Pointers validity

If we try to prove that the swap function is correct (comprising the verification of absence of runtime errors), our postcondition is indeed verified but WP fails to prove that some runtime-error cannot happen, since we perform access to some pointers that we did not indicate to be valid pointers in the precondition of the function.

We can express that the dereferencing of a pointer is valid using the `\valid` predicate of ACSL which receives the pointer in input:

3. Function contract

```
3  /*@
4   requires \valid(a) && \valid(b);
5   ensures  *a == \old(*b) && *b == \old(*a);
6  */
7  void swap(int* a, int* b){
8      int tmp = *a;
9      *a = *b;
10     *b = tmp;
11 }
```

Once we have specified that the pointers we receive in input must be valid, dereferencing is assured to not produce undefined behaviors.

As we will see later in this tutorial, `\valid` can take more than one pointer in parameter. For example, we can give it an expression such as: `\valid(p + (s .. e))` which means “for all `i` between included `s` and `e`, `p+i` is a valid pointer”. This kind of expression is extremely useful when we have to specify properties about arrays in specifications.

If we have a closer look at the assertions that RTE adds in the swap function when we ask the verification of absence of runtime errors, we can notice that there exists another version of the `\valid` predicate, denoted `\valid_read`. As opposed to `\valid`, the predicate `\valid_read` indicates that a pointer can be dereferenced, but only to read the pointed memory. This subtlety is due to the C language, where the downcast of a const pointer is easy to write but is not necessarily legal.

Typically, in this code:

```
1  /*@ requires \valid(p); */
2  int unref(int* p){
3      return *p;
4  }
5
6  int const value = 42;
7
8  int main(){
9      int i = unref(&value);
10 }
```

Dereferencing `p` is valid, however the precondition of `unref` is not verified by WP since dereferencing `value` is only legal for a read-access. A write access would result in an undefined behavior. In such a case, we can specify that the pointer `p` must be `\valid_read` and not `\valid`.

3.2.2.3. Side Effects

Our `swap` function is provable with regard to the specification and potential runtime errors, however is our specification precise enough? We can slightly modify our code to check this (we use `assert` to verify some properties at some particular points):

3. Function contract

```
1  int h = 42;
2
3  /*@
4   requires \valid(a) && \valid(b);
5   ensures  *a == \old(*b) && *b == \old(*a);
6  */
7  void swap(int* a, int* b){
8      int tmp = *a;
9      *a = *b;
10     *b = tmp;
11 }
12
13 int main(){
14     int a = 37;
15     int b = 91;
16
17     //@ assert h == 42;
18     swap(&a, &b);
19     //@ assert h == 42;
20 }
```

The result is not exactly what we expect:

```
int main(void)
{
    int __retres;
    int a = 37;
    int b = 91;
    /*@ assert h == 42; */ ;
    swap(&a, &b);
    /*@ assert h == 42; */ ;
    __retres = 0;
    return __retres;
}
```

Indeed, we did not specify the allowed side effects for our function. In order to specify side effects, we use an `assigns` clause which is part of the postcondition of a function. It allows us to specify which **non-local** elements (we verify side effects) can be modified during the execution of the function.

By default, WP considers that a function can modify everything in memory. So, we have to specify what can be modified by a function. For example, our `swap` function can be specified to modify the values pointed by the received pointers:

```
3  /*@
4   requires \valid(a) && \valid(b);
5
6   assigns *a, *b;
7
8   ensures  *a == \old(*b) && *b == \old(*a);
9  */
10 void swap(int* a, int* b){
11     int tmp = *a;
12     *a = *b;
13     *b = tmp;
14 }
```

If we ask WP to prove the function with this specification, it is validated (including with the variable added in the previous source code).

3. Function contract

Finally, we sometimes want to specify that a function is side effect free. We specify this by giving `\nothing` to `assigns`:

```
1  /*@
2   requires \valid_read(a);
3   requires *a <= INT_MAX - 5 ;
4
5   assigns \nothing ;
6
7   ensures \result == *a + 5 ;
8 */
9  int plus_5(int* a){
10     return *a + 5 ;
11 }
```

The careful reader will now be able to take back the examples we presented until now to integrate the right `assigns` clause.

3.2.2.4. Memory location separation

Pointers bring the risk of aliasing (multiple pointers can have access to the same memory location). For some functions, it does not cause any problem, for example when we give two identical pointers to the `swap` function, the specification is still verified. However, sometimes it is not that simple:

```
1  #include <limits.h>
2
3  /*@
4   requires \valid(a) && \valid_read(b);
5   assigns *a;
6   ensures *a == \old(*a) + *b;
7   ensures *b == \old(*b);
8 */
9  void incr_a_by_b(int* a, int const* b){
10     *a += *b;
11 }
```

If we ask WP to prove this function, we get the following result:

```
④ /*@ requires \valid(a) ^ \valid_read(b);
   ensures *\old(a) == \old(*a) + *\old(b);
   ensures *\old(b) == \old(*b);
   assigns *a;
   */
void incr_a_by_b(int *a, int const *b)
{
    *a += *b;
    return;
}
```

The reason is simply that we do not have any guarantee that the pointer `a` is different of the pointer `b`. Now, if these pointers are the same,

3. Function contract

- the property `*a == \old(*a) + *b` in fact means `*a == \old(*a) + *a` which can only be true if the old value pointed by `a` was 0, and we do not have such a requirement,
- the property `*b == \old(*b)` is not validated because we potentially modify this memory location.

?

Why is the `assigns` clause validated?

The reason is simply that `a` is indeed the only modified memory location. If `a != b`, we only modify the location pointed by `a`, and if `a == b`, this is still the case: `b` is not another location.

In order to ensure that pointers refer to separated memory locations, ACSL provides the predicate `\separated(p1, ..., pn)` that receives in parameter a set of pointers and is true if and only if these pointers are non-overlapping. Here, we specify:

```
1  #include <limits.h>
2
3  /*@
4   requires \valid(a) && \valid_read(b);
5   requires \separated(a, b);
6   assigns *a;
7   ensures *a == \old(*a) + *b;
8   ensures *b == \old(*b);
9  */
10 void incr_a_by_b(int* a, int const* b){
11     *a += *b;
12 }
```

And this time, the function is verified:

```
1  /*@ requires \valid(a) ^ \valid_read(b);
2   requires \separated(a, b);
3   ensures *\old(a) == \old(*a) + *\old(b);
4   ensures *\old(b) == \old(*b);
5   assigns *a;
6  */
7  void incr_a_by_b(int *a, int const *b)
8  {
9      *a += *b;
10     return;
11 }
```

We can notice that we do not consider the arithmetic overflow here, as we do not focus on this question in this section. However, if this function was part of a complete program, it would be necessary to define the context of use of this function and the precondition guaranteeing the absence of overflow.

3. Function contract

3.2.3. Writing the right contract

Writing a specification that is precise enough can sometimes be a bit tricky. Interestingly, a good way to check if a specification is precise enough is to write tests. And in fact, this is basically what we have done for our examples `max` and `swap`. We have written a first version of the specification and we have written some code with a call to the corresponding function to determine whether we could prove some properties that we expected to be easily provable from the contract of the function.

The most important idea is to determine the contract without taking in account the content of the function (at least, in a first step). Indeed, we are trying to prove the function, but maybe it contains a bug, so if we write the contract taking in account too directly its code, we have a risk to introduce the same bug, for example taking in account an erroneous conditional structure. In fact, it is generally a good practice to work with someone else. One specifies the function and the other implements it (even if they previously agreed on a common textual specification).

Once the contract have been stated, we work on the specifications that are due to the constraints of our language and our hardware. That mostly concerns the precondition of the function. For example, the absolute value does not really have a precondition, this is our hardware that adds the condition we have given in precondition due to the two's complement on which it relies. As we will see in the chapter 7, verifying the absence of runtime errors can also impact the postcondition. For now, let us leave this.

3.2.4. Exercises

3.2.4.1. Division and remaining

Specify the postcondition of the following function, that computes the results of the division of `a` by `b` and its remaining and stores it into two memory locations `p` and `q`:

```
1 void div_rem(unsigned x, unsigned y, unsigned* q, unsigned* r){
2     *q = x / y ;
3     *r = x % y ;
4 }
```

Run the command:

```
1 frama-c-gui your-file.c -wp
```

Once the function is successfully proved to respect the contract, run:

```
1 frama-c-gui your-file.c -wp -wp-rte
```

If it fails, complete the contract by adding the right precondition.

3. Function contract

3.2.4.2. Reset on condition

Provide a contract for the following function that reset its first parameter if and only if the second is true. Be sure to express that the second parameter remains unmodified:

```
1 void reset_1st_if_2nd_is_true(int* a, int const* b){
2     if(*b) *a = 0 ;
3 }
4
5 int main(){
6     int a = 5 ;
7     int x = 0 ;
8
9     reset_1st_if_2nd_is_true(&a, &x);
10    //@ assert a == 5 ;
11    //@ assert x == 0 ;
12
13    int const b = 1 ;
14
15    reset_1st_if_2nd_is_true(&a, &b);
16    //@ assert a == 0 ;
17    //@ assert b == 1 ;
18 }
```

Run the command:

```
1 frama-c-gui your-file.c -wp -wp-rte
```

3.2.4.3. Addition of pointed values

The following function receives two pointers as an input and returns the sum the pointed values. Write the contract of this function:

```
1 int add(int *p, int *q){
2     return *p + *q ;
3 }
4
5 int main(){
6     int a = 24 ;
7     int b = 42 ;
8
9     int x ;
10
11    x = add(&a, &b) ;
12    //@ assert x == a + b ;
13    //@ assert x == 66 ;
14
15    x = add(&a, &a) ;
16    //@ assert x == a + a ;
17    //@ assert x == 48 ;
18 }
```

Run the command:

3. Function contract

```
1 frama-c-gui your-file.c -wp -wp-rte
```

Once the function and calling code are successfully proved, modify the signature of the function `add` as follows:

```
1 void add(int* a, int* b, int* r);
```

Now the result of the sum should be stored at `r`. Accordingly modify the calls in the main function as well as the code and the contract of `add`.

3.2.4.4. Maximum of pointed values

The following functions computes the maximum of the values pointed by `a` and `b`. Write the contract of the function:

```
1 int max_ptr(int* a, int* b){
2     return (*a < *b) ? *b : *a ;
3 }
4
5 extern int h ;
6
7 int main(){
8     h = 42 ;
9
10    int a = 24 ;
11    int b = 42 ;
12
13    int x = max_ptr(&a, &b) ;
14
15    //@ assert x == 42 ;
16    //@ assert h == 42 ;
17 }
```

Run the command:

```
1 frama-c-gui your-file.c -wp -wp-rte
```

Once it is proved, modify the signature of the function as follows:

```
1 void max_ptr(int* a, int* b);
```

Now the function should ensure that after its execution `*a` contains the maximum of the input value, and `*b` contains the other value. Modify the code accordingly as well as the contract. Note that the variable `x` is not necessary anymore in the `main` function and that we can change the assertion on line 15 to reflect the new behavior of the function.

3. Function contract

3.2.4.5. Order 3 values

The following function should order the 3 input values in increasing order. Write the corresponding code and specification of the function:

```
1 void order_3(int* a, int* b, int* c){  
2   // CODE  
3 }
```

And run the command:

```
1 frama-c-gui your-file.c -wp -wp-rte
```

Remember that ordering values is not just ensuring that resulting values are sorted increasing order that each pointed value must be one the original ones. All original values should still be there after the sorting operation: new values are a permutation of the original ones. To express this idea, one can rely on the set datatype. For example, this property is true:

```
1 //@ assert { 1, 2, 3 } == { 2, 3, 1 };
```

We can use this to express that the set of original values and final values is the same. However, that is not the only thing to consider, as a set only contains one occurrence of each value. So, if `*a == *b == 1`, `{ *a, *b } == { 1 }`. Thus, we also have to consider three other particular cases:

- all original values equal
- two original values equal and the last is greater
- two original values equal and the last is lower

That should set one more constraint on the final values.

As an helper, one could use the following test program:

```
27 void test(){  
28   int a1 = 5, b1 = 3, c1 = 4 ;  
29   order_3(&a1, &b1, &c1) ;  
30   //@ assert a1 == 3 && b1 == 4 && c1 == 5 ;  
31  
32   int a2 = 2, b2 = 2, c2 = 2 ;  
33   order_3(&a2, &b2, &c2) ;  
34   //@ assert a2 == 2 && b2 == 2 && c2 == 2 ;  
35  
36   int a3 = 4, b3 = 3, c3 = 4 ;  
37   order_3(&a3, &b3, &c3) ;  
38   //@ assert a3 == 3 && b3 == 4 && c3 == 4 ;  
39  
40   int a4 = 4, b4 = 5, c4 = 4 ;  
41   order_3(&a4, &b4, &c4) ;  
42   //@ assert a4 == 4 && b4 == 4 && c4 == 5 ;
```

3. Function contract

```
43 }
```

If the specification is precise enough, each assertion should be proved. However, that does not mean that all cases have been considered by our tests, so do not hesitate to add other cases.

3.3. Behaviors

Sometimes, a function can have behaviors that can be quite different depending on the input. Typically, a function can receive a pointer to an optional resource: if the pointer is **NULL**, we have a certain behavior, which is different of the behavior expected when the pointer is not **NULL**.

We have already seen a function that have different behaviors: the **abs** function. Let us use it again to illustrate behaviors. We have two behaviors for the **abs** function: either the input is positive or it is negative.

Behaviors allow us to specify the different cases for postconditions. We introduce them using the **behavior** keyword. Each behavior is named. For a given behavior, we have different assumptions about the input of the function, they are introduced with the clause **assumes** (note that since they characterize the input, the keyword `\old` cannot be used there). However, the properties expressed by this clause do not have to be verified before the call, they can be verified and in this case, the postcondition specified in our behavior applies. The postconditions of a particular behavior are introduced using **ensures**. Finally, we can ask WP to verify that behaviors are disjoint (to guarantee determinism) and complete (to guarantee that we cover all possible input).

Behaviors are disjoint if for any (valid) input of the function, it corresponds to the assumption (**assumes**) of a single behavior. Behaviors are complete if any (valid) input of the function corresponds to at least one behavior.

For example, for **abs** we can write the specification:

```
1  #include <limits.h>
2
3  /*@
4   requires val > INT_MIN;
5   assigns  \nothing;
6
7   ensures \result >= 0;
8
9   behavior pos:
10    assumes 0 <= val;
11    ensures \result == val;
12
13  behavior neg:
14    assumes val < 0;
15    ensures \result == -val;
16
17  complete behaviors;
18  disjoint behaviors;
19 */
20 int abs(int val){
```

3. Function contract

```
21  if(val < 0) return -val;
22  return val;
23 }
```

Note that declaring behaviors does not forbid to specify global postconditions. For example here, we have specified that for any behavior, the function must return a positive value.

Let us now slightly modify the assumptions of each behavior to illustrate the meaning of `complete` and `disjoint` :

- replace the assumption of `pos` with `val > 0`, in this case, behaviors are disjoint but incomplete (we miss `val == 0`),
- replace the assumption of `neg` with `val <= 0`, in this case, behaviors are complete but not disjoint (we have two assumptions corresponding to `val == 0`).



Even if `assigns` is a postcondition, indicating different assigns in each behavior is currently not well-handled by WP. If we need to specify this, we will:

- put our `assigns` before the behaviors (as we have done in our example) with all potentially modified non-local elements,
- add in postcondition of each behaviors the elements that are in fact not modified by indicating their new value to be equal to the `\old` one.

Behaviors are useful to simplify the writing of specifications when functions can have very different behaviors depending on their input. Without them, specification would be defined using implications expressing the same idea but harder to write and read (which would be error-prone). On the other hand, the translation of completeness and disjointedness would be necessarily written by hand which would be tedious and again error-prone.

3.3.1. Exercises

3.3.1.1. Distance

Take back the example about the computation of the distance between two integers. Considering that the written contract was:

```
1  #include <limits.h>
2
3  /*@
4   requires a < b ==> b - a <= INT_MAX ;
5   requires b <= a ==> a - b <= INT_MAX ;
6
7   ensures a < b ==> a + \result == b ;
8   ensures b <= a ==> a - \result == b ;
9  */
10 int distance(int a, int b){
11     if(a < b) return b - a ;
```

3. Function contract

```
12     else return a - b ;  
13 }
```

Re-write it using behaviors.

3.3.1.2. Reset on condition

Take back the example “reset on condition” from the previous section. Considering that the written contract was:

```
1  /*@  
2   requires \valid(a) && \valid_read(b) ;  
3   requires \separated(a, b) ;  
4  
5   assigns *a ;  
6  
7   ensures \old(*b) ==> *a == 0 ;  
8   ensures ! \old(*b) ==> *a == \old(*a) ;  
9   ensures *b == \old(*b);  
10 */  
11 void reset_1st_if_2nd_is_true(int* a, int const* b){  
12     if(*b) *a = 0 ;  
13 }
```

Re-write it using behaviors.

3.3.1.3. Days of the month

Take back the example “days of the month” from the first section. Considering that the written contract was:

```
1  /*@  
2   requires 1 <= m <= 12 ;  
3   ensures m \in { 2 } ==> \result == 28 ;  
4   ensures m \in { 1, 3, 5, 7, 8, 10, 12 } ==> \result == 31 ;  
5   ensures m \in { 4, 6, 9, 11 } ==> \result == 30 ;  
6 */  
7 int day_of(int m){  
8     int days[] = { 31, 28, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31 } ;  
9     return days[m-1] ;  
10 }
```

Re-write it using behaviors.

3.3.1.4. Max of pointer values, ordering

Take back the example “Max of pointed values” from the previous section, this time with the version that orders the values. Considering that the contract was:

3. Function contract

```
1  /*@
2   requires \valid(a) && \valid(b);
3   assigns  *a, *b ;
4   ensures  \old(*a) < \old(*b) ==> *a == \old(*b) && *b == \old(*a) ;
5   ensures  \old(*a) >= \old(*b) ==> *a == \old(*a) && *b == \old(*b) ;
6  */
7  void max_ptr(int* a, int* b){
8      if(*a < *b){
9          int tmp = *b ;
10         *b = *a ;
11         *a = tmp ;
12     }
13 }
```

Re-write it using behaviors.

3.3.1.5. Max of pointed values, returning the result

Take back the example “Max of pointed values” from the previous section, and more precisely, the version that returns the result. Considering that the contract was:

```
1  /*@
2   requires \valid_read(a) && \valid_read(b);
3   assigns  \nothing ;
4   ensures  *a < *b ==> \result == *b ;
5   ensures  *a >= *b ==> \result == *a ;
6   ensures  \result == *a || \result == *b ;
7  */
8  int max_ptr(int* a, int* b){
9      return (*a < *b) ? *b : *a ;
10 }
```

1. Rewrite it using behaviors
2. Modify the contract of 1. in order to make the behaviors non-disjoint, except this property, the contract should remain verified,
3. Modify the contract of 1. in order to make the behaviors incomplete, add a new behavior that makes the contract complete again,
4. Modify the function of 1. in order to accept `NULL` pointers for both `a` and `b`. If both of them are nul pointers, return `INT_MIN`, if one is a nul pointer, return the value of the other, else, return the maximum of them. Modify the contract accordingly by adding new behaviors. Be sure that they are disjoint and complete.

3.3.1.6. Order 3

Take back the example “Order 3 values” from the previous section. Considering that the contract was:

3. Function contract

```
1  /*@
2   requires \valid(a) && \valid(b) && \valid(c) ;
3   requires \separated(a, b, c);
4
5   assigns *a, *b, *c ;
6
7   ensures *a <= *b <= *c ;
8   ensures { *a, *b, *c } == \old({ *a, *b, *c }) ;
9
10  ensures \old(*a == *b == *c) ==> *a == *b == *c ;
11  ensures \old(*a == *b < *c || *a == *c < *b || *b == *c < *a) ==> *a == *b ;
12  ensures \old(*a == *b > *c || *a == *c > *b || *b == *c > *a) ==> *b == *c ;
13 */
14 void order_3(int* a, int* b, int* c){
15     if(*a > *b){ int tmp = *b ; *b = *a ; *a = tmp ; }
16     if(*a > *c){ int tmp = *c ; *c = *a ; *a = tmp ; }
17     if(*b > *c){ int tmp = *b ; *b = *c ; *c = tmp ; }
18 }
```

Rewrite it using behaviors. Note that you should have one general behaviors and 3 specific behaviors. Are these behaviors complete? Are they disjoint?

3.4. WP Modularity

For this last part, let us talk about function call composition, and have a closer look at WP. We will also have a look at the way we can split our programs in different files when we want to prove them using WP.

Our goal is to prove the `max_abs` function, that returns the maximum absolute value of two values:

```
6  int max_abs(int a, int b){
7     int abs_a = abs(a);
8     int abs_b = abs(b);
9
10    return max(abs_a, abs_b);
11 }
```

Let us start by (over-)separating the declarations and definitions of the different functions we need (and have previously proved) into header/source files, that are `abs` and `max`. We obtain, for `abs`:

File `abs.h` :

```
1  #ifndef _ABS
2  #define _ABS
3
4  #include <limits.h>
5
6  /*@
7   requires val > INT_MIN;
8   assigns \nothing;
9
10  behavior pos:
```

3. Function contract

```
11     assumes 0 <= val;
12     ensures \result == val;
13
14     behavior neg:
15         assumes val < 0;
16         ensures \result == -val;
17
18     complete behaviors;
19     disjoint behaviors;
20 */
21 int abs(int val);
22
23 #endif
```

File abs.c

```
1 #include "abs.h"
2
3 int abs(int val){
4     if(val < 0) return -val;
5     return val;
6 }
```

We can notice that we put our function contract in the header file. The goal is to be able to import the specification at the same time as the declaration when we need it in another file. Indeed, WP needs the contract of the function when it is called. First to prove that the precondition is verified (and thus that the call is legal), and second to get in return the postcondition that is useful to prove the right properties after the function call.

We can create a file using the same format for the `max` function. In both cases, we can open the source file (we do not need to specify header files in the command line) with Frama-C and notice that the specification is indeed associated to the function and that we prove it.

Now, we can prepare our files for the `max_abs` function with the header:


```
1 #ifndef _MAX_ABS
2 #define _MAX_ABS
3
4 int max_abs(int a, int b);
5
6 #endif
```

And its source file:

```
1 #include <limits.h>
2 #include "max_abs.h"
3 #include "abs.h"
4 #include "max.h"
5
6 int max_abs(int a, int b){
7     int abs_a = abs(a);
8     int abs_b = abs(b);
9
10    return max(abs_a, abs_b);
11 }
```

3. Function contract

We can open the source file in Frama-C. If we look at the side panel, we can see that the header files we have included in `abs_max` correctly appear and if we look at the function contracts for them, we can see some blue and green bullets:



These bullets indicate that, since we do not have the implementation of the function, the postcondition of the function is assumed to be true. It is an important strength of the deductive proof of programs compared to some other formal methods: functions are verified in isolation from each other.

When we are not currently performing the proof of a function, its specification is considered to be correct: we do not try to prove it when we are proving another function, we only verify that the precondition is correctly established when we call it. It provides very modular proofs and specifications that are therefore more reusable. Of course, if our proof relies on the specification of another function, it must be provable to ensure that the proof of the program is complete. But, we can also consider that we trust a function that comes from an external library that we do not want to prove (or for which we do not even have the source code).

The careful reader could specify and prove the `max_abs` function. A possible answer is provided there:

```
4  /*@
5   requires a > INT_MIN;
6   requires b > INT_MIN;
7
8   assigns \nothing;
9
10  ensures \result >= 0;
11  ensures \result >= a && \result >= -a && \result >= b && \result >= -b;
12  ensures \result == a || \result == -a || \result == b || \result == -b;
13 */
14 int max_abs(int a, int b);
```

3.4.1. Exercises

3.4.1.1. Days of the month

Specify the function `leap year` that returns true if the year received as an input is leap. Use this functions to complete the function `days_of` in order to return the number of days of the month received as an input, including the right behavior when the year is leap for February.

3. Function contract

```
1 int leap(int y){
2     return ((y % 4 == 0) && (y % 100 != 0)) || (y % 400 == 0) ;
3 }
4
5 int days_of(int m, int y){
6     int days[] = { 31, 28, 31, 30, 31, 30, 31, 31, 30, 31, 30, 31 } ;
7     int n = days[m-1] ;
8     // code
9 }
```

3.4.1.2. Alpha-numeric character

Write and specify the different functions used by `is_alpha_num`. Provide a contract for each of them and provide the contract of `is_alpha_num`.

```
1 int is_alpha_num(char c){
2     return
3         is_lower_alpha(c) ||
4         is_upper_alpha(c) ||
5         is_digit(c) ;
6 }
```

Declare an enumeration with values `LOWER`, `UPPER`, `DIGIT` and `OTHER`, and a function `character_kind` that returns, using the different functions `is_lower`, `is_upper`, `is_digit`, the kind of character received in input. Use behaviors to specify the contract of this function and be sure that they are disjoint and complete.

3.4.1.3. Order 3 values

Taking back the function `max_ptr` that orders two values, putting the maximum at the first location and the minimum at the second, write a function `min_ptr` that uses this function and produces the opposite operation. Use these functions to complete the four functions that orders 3 values. For each variant (increasing and decreasing), write it once using only `max_ptr` and once using only `min_ptr`. Write a precise contract for each of these functions and prove them.

```
1 void max_ptr(int* a, int* b){
2     if(*a < *b){
3         int tmp = *b ;
4         *b = *a ;
5         *a = tmp ;
6     }
7 }
8
9 void min_ptr(int* a, int* b){
10    // use max_ptr
11 }
12
13 void order_3_inc_max(int* a, int* b, int* c){
14    //in increasing order using max_ptr
```

3. Function contract

```
15 }
16
17 void order_3_inc_min(int* a, int* b, int* c){
18     //in increasing order using min_ptr
19 }
20
21 void order_3_dec_max(int* a, int* b, int* c){
22     //in decreasing order using max_ptr
23 }
24
25 void order_3_dec_min(int* a, int* b, int* c){
26     //in decreasing order using min_ptr
27 }
```

3. *Function contract*

During this part of the tutorial, we have studied how we can specify functions using contracts, composed of a pre and a postcondition, as well as some features ACSL provides to express those properties. We have also seen why it is important to be precise when we specify and how the introduction of behaviors can help us to write more understandable and safer specification.

However, we have not studied one important aspect: the proof of programs with loops. Before that, we should have a closer look at the way WP works.

4. Basic instructions and control structures

i

This part is more formal than what we have seen so far. If the reader wishes to concentrate on the usage of the tool, they can skip the introduction and the first two sections (about the basic instructions and the bonus training) of this chapter. If what we presented so far has been difficult for the reader from a formal point of view, it is well possible to reserve the introduction and the two sections for a later reading.

The sections on loops, however, are indispensable. We will highlight the more formal parts of these sections.

We will associate with every C programming construct

- the corresponding inference rule together,
- its governing rule from the weakest precondition calculus and
- examples that show its usage.

Not necessarily in that order and sometimes only with a loose connection to the tool. Since the first rules are quite simple, we will discuss them in a fairly theoretical manner. Later on, however, our presentation will rely more and more on the tool, in particular when we begin dealing with loops.

4.0.1. Inference rules

An inference rule is of the form

$$\frac{P_1 \quad \dots \quad P_n}{C}$$

and means that in order to assure that the conclusion C is true, first the truth of the premises P_1 , ..., and P_n has to be established. In case that the rule has no premises

$$\frac{}{C}$$

then nothing has to be assured in order to conclude the truth of C , and it is called an axiom.

On the other hand, in order to prove that a certain premise is true, it might be necessary to employ other inference rules which would lead to something like this:

$$\frac{\frac{}{P_1} \quad \frac{\frac{}{P_{n_1}} \quad \frac{}{P_{n_2}}}{P_n}}{C}}$$

This way, we obtain step by step the *deduction tree* of our reasoning. In our case, the premises and conclusions under consideration will in general be *Hoare triples*.

4. Basic instructions and control structures

4.0.2. Hoare triples

We are now returning to concept of a Hoare triple:

$$\{P\} \quad C \quad \{Q\}$$

In the beginning of this tutorial we have seen that this triple expresses the following: if the property P holds before the execution of C and if C terminates, then the property Q holds too. For example, if we take up again our (slightly modified) program for the computation of the absolute value:

```
1  /*@
2   ensures \result >= 0;
3   ensures (val >= 0 ==> \result == val ) && (val < 0 ==> \result == -val);
4  */
5  int abs(int val){
6      int res;
7      if(val < 0) res = - val;
8      else      res = val;
9
10     return res;
11 }
```

The rules of Hoare logic tell us that in order to show that our program satisfies its contract we have to verify the properties shown in the braces. (We have omitted one postcondition in order to simplify the presentation.)

```
1  int abs(int val){
2      int res;
3      // { P }
4      if(val < 0){
5          // { (val < 0) && P }
6          res = - val;
7          // { \at(val, Pre) >= 0 ==> res == val && \at(val, Pre) < 0 ==> res == -val }
8      } else {
9          // { !(val < 0) && P }
10         res = val;
11         // { \at(val, Pre) >= 0 ==> res == val && \at(val, Pre) < 0 ==> res == -val }
12     }
13     // { \at(val, Pre) >= 0 ==> res == val && \at(val, Pre) < 0 ==> res == -val }
14
15     return res;
16 }
```

Yet, Hoare logic does not tell us how we can automatically obtain the precondition P of our program `abs`. Dijkstra's *weakest-precondition calculus*, on the other hand, allows us to compute from a given postcondition Q and a code snippet C the minimal precondition P that ensures Q after the execution of C . We are thus in a position to determine for our example `abs` the desired property P .

In this chapter, we present the different cases of the function wp which, starting from a given postcondition and a program (or statement), computes the *weakest* precondition that allows us to establish the validity of the postcondition. We will use the following notation to define the computation that corresponds to one or several statements:

$$wp(Instruction(s), Post) := WeakestPrecondition$$

4. Basic instructions and control structures

The function *wp* will guarantee that the Hoare triple

$$\{ wp(C, Q) \} \quad C \quad \{ Q \}$$

really is a valid triple.

We will thereby often use ACSL assertions in order to represent the upcoming concepts:

```
1 //@ assert some_property ;
```

These assertions correspond in fact to possible intermediate steps for the properties in our Hoare triples. We can, for example, replace the properties of our function `abs` by corresponding ACSL assertions (we have omitted here the property `P` because it is just `true`):

```
1 int abs(int val){
2   int res;
3   if(val < 0){
4     //@ assert val < 0 ;
5     res = - val;
6     //@ assert \at(val, Pre) >= 0 ==> res == val && \at(val, Pre) < 0 ==> res == -val ;
7   } else {
8     //@ assert !(val < 0) ;
9     res = val;
10    //@ assert \at(val, Pre) >= 0 ==> res == val && \at(val, Pre) < 0 ==> res == -val ;
11  }
12  //@ assert \at(val, Pre) >= 0 ==> res == val && \at(val, Pre) < 0 ==> res == -val ;
13
14  return res;
15 }
```

4.1. Basic concepts

4.1.1. Assignment

Assignment is the most basic operation one can have in an imperative language (leaving aside the “do nothing” operation that is not particularly interesting). The weakest precondition calculus associates the following

$$wp(x = E, Post) := Post[x \leftarrow E]$$

Here the notation $P[x \leftarrow E]$ means “the property P where x is replaced by E ”. In our case this corresponds to “the postcondition $Post$ where x is replaced by E ”. The idea is that the postcondition of an assignment of E to x can only be true if replacing all occurrences of x in the formula by E leads to a property that is true. For example:

```
1 // { P }
2 x = 43 * c ;
3 // { x = 258 }
```

4. Basic instructions and control structures

$$P = wp(x = 43 * c, \{x = 258\}) = \{43 * c = 258\}$$

The function wp allows us to compute, as weakest precondition of the assignment provided our expected postcondition, the formula $\{43 * c = 258\}$, thus obtaining the following Hoare triple:

```
1 // { 43*c = 258 }
2 x = 43 * c ;
3 // { x = 258 }
```

In order to compute the precondition of the assignment we have replaced each occurrence of x in the postcondition by the assigned value $E = 43 * c$. If our program were of the form:

```
1 int c = 6 ;
2 // { 43*c = 258 }
3 x = 43 * c ;
4 // { x = 258 }
```

we could submit the formula " $43 * 6 = 258$ " to our automatic prover in order to determine whether it is really valid. The answer would of course be "yes" because the property is easy to verify. If we had, however, given the value 7 to the variable c the prover's reply would be "no" since the formula $43 * 7 = 258$ is not true.

Taking into account the weakest precondition calculus, we can now write the inference rule for the Hoare triple of an assignment as

$$\frac{}{\{Q[x \leftarrow E]\} \quad x = E \quad \{Q\}}$$

We note that there is no premise to verify. Does this mean that the triple is necessarily true? Yes. However, it does not mean that the precondition is respected by the program to which the assignment belongs or that the precondition is at all possible. Here the automatic provers come into play.

For example, we can ask Frama-C to verify the following line:

```
1 int a = 42;
2 //@ assert a == 42;
```

which is, of course, directly proven by Qed, since it is a simple applications of the assignment rule.

i

We remark that according to the C standard, an assignment is in fact an expression. This allows us, for example, to write `if((a = foo()) == 42)`. In Frama-C, an assignment will always be treated as a statement. Indeed, if an assignment occurs within a larger expression, then the Frama-C preprocessor, while building the abstract syntax tree, systematically performs a *normalization step* that produces a separate assignment



statement.

4.1.1.1. Assignment of pointed value

In C, thanks to (because of?) pointers, we can have programs with aliases, meaning that two pointers can point to the same memory location. Our weakest precondition calculus should consider these cases. For example, let us consider this simple Hoare triple:

```

1 //@ assert p = q ;
2 *p = 1 ;
3 //@ assert *p + *q == 2 ;

```

This Hoare triple is correct, since `p` and `q` are in alias, modifying `*p` also modifies `*q`, thus both these expression evaluate to 1 and the postcondition is true. However let us apply the weakest precondition calculus from the postcondition:

$$\begin{aligned}
 wp(*p = 1, *p + *q = 2) &= (*p + *q = 2)[*p \leftarrow 1] \\
 &= (1 + *q = 2)
 \end{aligned}$$

We get the weakest precondition: `1 + *q == 2`, and thus we could deduce that the weakest precondition is `*q == 1`, which is true, but does not allow us to conclude that the program is correct, since in our formula we do not have anything that models that `p == q ==> *q == 1`. In fact, here, we would like to be able to compute a weakest precondition like:

$$\begin{aligned}
 wp(*p = 1, *p + *q = 2) &= (1 + *q = 2 \vee q = p) \\
 &= (*q = 1 \vee q = p)
 \end{aligned}$$

For this, we have to take care of aliasing. A common way to do this is to consider that the memory is one particular variable (let us name this variable M) on which we can perform two operations: get the element at a particular location l in memory (which returns an expression) and set the element at a particular location l to a new value v (which returns the new memory).

We denote:

- $get(M, l)$ with the notation $M[l]$
- $set(M, l, v)$ with the notation $M[l \mapsto v]$

And basically, the get operation can be seen as follows:

$$\begin{aligned}
 M[l1 \mapsto v][l2] &= \text{if } l1 = l2 \text{ then } v \\
 &\quad \text{if } l1 \neq l2 \text{ then } M[l2]
 \end{aligned}$$

If there is no value associated to the location we use for a get, the value is undefined (thus, the memory is partial function). Of course, at the beginning of a function, the memory context can be populated with the memory locations for which a value is known to be defined.

4. Basic instructions and control structures

Now, we can change a little bit the weakest precondition calculus for assignment of pointed memory location. For this, we consider that we have an implicit variable M that models the memory, and we define the assignment of a memory location as an update of the memory such that now the corresponding pointer points to the written expression.

$$wp(*x = E, Q) := Q[M \leftarrow M[x \mapsto E]]$$

And evaluating a pointed value $*x$ in a formula now requires us to use the get operator to ask the right value. Thus we can for example compute the weakest precondition of our previous program:

$$wp(*p = 1, *p + *q = 2) = (*p + *q = 2)[M \leftarrow M[p \mapsto 1]] \quad (1)$$

$$= (M[p] + M[q] = 2)[M \leftarrow M[p \mapsto 1]] \quad (2)$$

$$= (M[p \mapsto 1][p] + M[p \mapsto 1][q] = 2) \quad (3)$$

$$= (1 + M[p \mapsto 1][q] = 2) \quad (4)$$

$$= (1 + (\text{if } q = p \text{ then } 1 \text{ else } M[q]) = 2) \quad (5)$$

$$= (\text{if } q = p \text{ then } 1 + 1 = 2 \text{ else } 1 + M[q] = 2) \quad (6)$$

$$= (q = p \vee M[q] = 1) \quad (7)$$

1. we have to apply the rule of assignment for pointers, but for this we need to introduce M ,
2. we replace pointer accesses in the formula by a call to *get* on M ,
3. we apply the replacement asked by the assignment rule,
4. we use the definition of the *get* operator for the expression about p ($M[p \mapsto 1][p] = 1$)
5. we use the definition of the *get* operator for the expression about q
($M[p \mapsto 1][q] = \text{if } q = p \text{ then } 1 \text{ else } M[q]$)
6. we perform some simplification to the formula ...
7. ... and finally conclude that either $M[q] = 1$ or $p = q$.

Then in our program, since we know that $p = q$, we can conclude that the program is correct.

The WP plugin does not exactly work like this. In particular, it depends on the memory model chosen for the proof that will make different assumption about the memory is organized. For the memory model we use, the typed memory model, in fact WP creates multiple variables for memory. However, let us have a look at the verification condition generated for the postcondition of the swap function:

4. Basic instructions and control structures

```

/*@ requires \valid(a) /\ \valid(b);
    ensures *\old(a) == \old(*b) /\ *\old(b) == \old(*a);
    assigns *a, *b;
*/

```

Information Messages (0) Console Properties Values Red Alarms WP Goals

Global All Results

Raw Obligation Proved Goal

No Script

Goal Post-condition:
 Let $x = \text{Mint_0}[a]$.
 Let $x_1 = \text{Mint_0}[b]$.
 Let $x_2 = \text{Mint_0}[a \leftarrow x_1][b \leftarrow x][a]$.
 Assume {
 Type: $\text{is_sint32}(x) \wedge \text{is_sint32}(x_1) \wedge \text{is_sint32}(x_2)$.
 (* Heap *)
 Have: $(\text{region}(a.\text{base}) \leq 0) \wedge (\text{region}(b.\text{base}) \leq 0) \wedge \text{linked}(\text{Malloc_0})$.
 (* Pre-condition *)
 Have: $\text{valid_rw}(\text{Malloc_0}, a, 1) \wedge \text{valid_rw}(\text{Malloc_0}, b, 1)$.
 }
 Prove: $x_2 = x_1$.

We can see, in the beginning of the verification condition, that a variable `Mint_0` representing a memory of values of integer types have been created, and that this memory is updated and accessed using the operators we previously introduced (see the definition of the variable `x_2`).

4.1.2. Composition of statements

For a statement to be valid, its precondition must allow us by means of executing the said statement to reach the desired postcondition. Now we would like to execute several statements one after another. Here the idea is that the postcondition of the first statement is compatible with the required precondition of the second statement and so on for the third statement.

The inference rule that corresponds to this idea utilizes the following Hoare triples:

$$\frac{\{P\} \quad S1 \quad \{R\} \quad \{R\} \quad S2 \quad \{Q\}}{\{P\} \quad S1; S2 \quad \{Q\}}$$

In order to verify the composed statement $S1; S2$ we rely on an intermediate property R that is at the same time the postcondition of $S1$ and the precondition of $S2$. (Please note that $S1$ and $S2$ are not necessarily simple statements; they themselves can be composed statements.) The problem is, however, that nothing indicates us how to determine the properties P and R .

The weakest-precondition calculus now says us that the intermediate property R can be computed as the weakest precondition of the second statement. The property P , on the other hand, then is computed as the weakest precondition of the first statement. In other words, the weakest precondition of the composed statement $S1; S2$ is determined as follows:

$$wp(S1; S2, Post) := wp(S1, wp(S2, Post))$$

The WP plugin of Frama-C performs all these computations for us. Thus, we do not have to write the intermediate properties as ACSL assertions between the lines of codes.

4. Basic instructions and control structures

```

1  int main(){
2      int a = 42;
3      int b = 37;
4
5      int c = a+b; // i:1
6      a -= c;      // i:2
7      b += a;      // i:3
8
9      //@assert b == 0 && c == 79;
10 }
```

4.1.2.1. Proof tree

When we have more than two statements, we can consider the last statement as second statement of our rule and all the preceding ones as first statement. This way we traverse step by step backwards the statements in our reasoning. With the previous program this looks like:

$$\frac{\frac{\{P\} \quad i_1; \quad \{Q_{-2}\} \quad \{Q_{-2}\} \quad i_2; \quad \{Q_{-1}\}}{\{P\} \quad i_{-1}; \quad i_{-2}; \quad \{Q_{-1}\}} \quad \{Q_{-1}\} \quad i_3; \quad \{Q\}}{\{P\} \quad i_{-1}; \quad i_{-2}; \quad i_{-3}; \quad \{Q\}}$$

The weakest-precondition calculus allows us to construct the property Q_{-1} starting from the property Q and statement i_3 which in turn enables us to derive the property Q_{-2} from the property Q_{-1} and statement i_2 . Finally, P can be determined from Q_{-2} and i_1 .

Now that we can verify programs that consist of several statements it is time to add some structure to them.

4.1.3. Conditional rule

For a conditional statement to be true, one must be able to reach the postcondition through both branches. Of course, for both branches, the same precondition (of the conditional statement) must hold. In addition, we have that in the if-branch the condition is true while in the else-branch it is false.

We therefore have, as in the case of composed statements, two facts to verify (in order to avoid confusion we are using here the syntax *if B then S1 else S2*):

$$\frac{\{P \wedge B\} \quad S1 \quad \{Q\} \quad \{P \wedge \neg B\} \quad S2 \quad \{Q\}}{\{P\} \quad \text{if } B \text{ then } S1 \text{ else } S2 \quad \{Q\}}$$

Our two premises are therefore that we can both in the if-branch and the else-branch reach the postcondition from the precondition.

The result of the weakest-precondition calculus for a conditional statement reads as follows:

$$wp(\text{if } B \text{ then } S1 \text{ else } S2, Post) := (B \Rightarrow wp(S1, Post)) \wedge (\neg B \Rightarrow wp(S2, Post))$$

4. Basic instructions and control structures

This means that the condition B has to imply the weakest precondition of $S1$ in order to safely arrive at the postcondition. Analogously, the negation of B must imply the weakest precondition of $S2$.

4.1.3.1. Empty `else`-branch

Following this definition, we obtain for the case of an empty else-branch the following rule by simply replacing the statement $S2$ by the empty statement `skip`.

$$\frac{\{P \wedge B\} \quad S1 \quad \{Q\} \quad \{P \wedge \neg B\} \quad \text{skip} \quad \{Q\}}{\{P\} \quad \text{if } B \text{ then } S1 \text{ else skip } \{Q\}}$$

The triple for `else` is:

$$\{P \wedge \neg B\} \quad \text{skip} \quad \{Q\}$$

which means that we need to ensure:

$$P \wedge \neg B \Rightarrow Q$$

In short, if the condition B of `if` is false, this means that the postcondition of the complete conditional statement is already established before entering the else-branch (since it does not do anything).

As an example, we consider the following code snippet where we reset a variable c to a default value in case it had not been properly initialized by the user.

```

1  int c;
2
3  // ... some code ...
4
5  if(c < 0 || c > 15){
6      c = 0;
7  }
8  //@ assert 0 <= c <= 15;
```

Let

$$\begin{aligned}
 & wp(\text{if } \neg(c \in [0; 15]) \text{ then } c := 0, \{c \in [0; 15]\}) \\
 &:= (\neg(c \in [0; 15]) \Rightarrow wp(c := 0, \{c \in [0; 15]\})) \wedge (c \in [0; 15] \Rightarrow wp(\text{skip}, \{c \in [0; 15]\})) \\
 &= (\neg(c \in [0; 15]) \Rightarrow 0 \in [0; 15]) \wedge (c \in [0; 15] \Rightarrow c \in [0; 15]) \\
 &= (\neg(c \in [0; 15]) \Rightarrow \text{true}) \wedge \text{true}
 \end{aligned}$$

The property can be verified: independent of the evaluation of $\neg(c \in [0; 15])$, the implication will hold.

4.1.4. Bonus Stage - Consequence rule

It can sometimes be useful to strengthen a postcondition or to weaken a precondition. The former will often be established by us to facilitate the work of the prover, the latter is more often verified by the tool as the result of computing the weakest precondition.

The inference rule of Hoare logic is the following:

$$\frac{P \Rightarrow WP \quad \{WP\} \quad c \quad \{SQ\} \quad SQ \Rightarrow Q}{\{P\} \quad c \quad \{Q\}}$$

(We remark that the premises here are not only Hoare triples but also formulas to verify.)

For example, if our postcondition is too complex, it may generate a weaker precondition that is, however, too complicated, thus making the work of provers more difficult. We can then create a simpler intermediate postcondition SQ , that is, however, stricter and implies the real postcondition. This is the part $SQ \Rightarrow Q$.

Conversely, the calculation of the precondition will usually generate a complicated and often weaker formula than the precondition we want to accept as input. In this case, it is our tool that will check the implication between what we want and what is necessary for our code to be valid. This is the part $P \Rightarrow WP$.

We can illustrate this with the following code. Note that here the code could be proved by WP without the weakening and strengthening of properties because the code is very simple, it is just to illustrate the rule of consequence.

```

1  /*@
2   requires P: 2 <= a <= 8;
3   ensures  Q: 0 <= \result <= 100 ;
4   assigns  \nothing ;
5  */
6  int constrained_times_10(int a){
7   //@ assert P_imply_WP: 2 <= a <= 8 ==> 1 <= a <= 9 ;
8   //@ assert WP:        1 <= a <= 9 ;
9
10  int res = a * 10;
11
12  //@ assert SQ:          10 <= res <= 90 ;
13  //@ assert SQ_imply_Q: 10 <= res <= 90 ==> 0 <= res <= 100 ;
14
15  return res;
16 }
```

(Note: We have omitted here the control of integer overflow.)

Here we want to have a result between 0 and 100. But we know that the code will not produce a result outside the bounds of 10 and 90. So we strengthen the postcondition with an assertion that at the end `res`, the result, is between 0 and 90. The calculation of the weakest precondition of this property together with the assignment `res = 10 * a` yields a weaker precondition `1 <= a <= 9` and we know that `2 <= a <= 8` gives us the desired guarantee.

When there are difficulties to carry out a proof on more complex code, then it is often helpful to write assertions that produce stronger, yet easier to verify, postconditions. Note that in the

4. Basic instructions and control structures

previous code, the lines `P_imply_WP` and `SQ_imply_Q` are never used because this is the default reasoning of WP. They are just here for illustrating the rule.

4.1.5. Bonus Stage - Constancy rule

Certain sequences of instructions may concern and involve different variables. Thus, we may initialize and manipulate a certain number of variables, begin to use some of them for a time, before using other variables. When this happens, we want our tool to be concerned only with variables that are susceptible to change in order to obtain the simplest possible properties.

The rule of inference that defines this reasoning is the following:

$$\frac{\{P\} \quad c \quad \{Q\}}{\{P \wedge R\} \quad c \quad \{Q \wedge R\}}$$

where c does not modify any variable in R . In other words: “To check the triple, let’s get rid of the parts of the formula that involve variables that are not influenced by c and prove the new triple.” However, we must be careful not to delete too much information, since this could mean that we are not able to prove our properties.

As an example, let us consider the following code (here again, we ignore potential integer overflows):

```
1  /*@
2   requires a > -99 ;
3   requires b > 100 ;
4   ensures  \result > 0 ;
5   assigns  \nothing ;
6  */
7  int foo(int a, int b){
8      if(a >= 0){
9          a++ ;
10     } else {
11         a += b ;
12     }
13     return a ;
14 }
```

If we look at the code of the `if` block, we notice that it does not use the variable `b`. Thus, we can completely omit the properties about `b` in order to prove that `a` will be strictly greater than 0 after the execution of the block:

```
1  /*@
2   requires a > -99 ;
3   requires b > 100 ;
4   ensures  \result > 0 ;
5   assigns  \nothing ;
6  */
7  int foo(int a, int b){
8      if(a >= 0){
9          //@ assert a >= 0; // and nothing about b
10         a++ ;
11     } else {
12         a += b ;
13     }
14     return a ;
15 }
```

4. Basic instructions and control structures

```
15 }
```

On the other hand, in the `else` block, even if `b` is not modified, formulating properties only about `a` would render a proof impossible for humans. The code would be:

```
1  /*@
2   requires a > -99 ;
3   requires b > 100 ;
4   ensures  \result > 0 ;
5   assigns  \nothing ;
6  */
7  int foo(int a, int b){
8      if(a >= 0){
9          /*@ assert a >= 0; // and nothing about b
10         a++;
11     } else {
12         /*@ assert a < 0 && a > -99 ; // and nothing about b
13         a += b ;
14     }
15     return a ;
16 }
```

In the `else` block, knowing that `a` lies between -99 and 0, but knowing nothing about `b`, we could hardly know if the operation `a += b` produces a result that is greater than 0.

The WP plug-in will, of course, prove the function without problems, since it produces by itself the properties that are necessary for the proof. In fact, the analysis which variables are necessary or not (and, consequently, the application of the constancy rule) is conducted directly by WP.

Let us finally remark that the constancy rule is an instance of the consequence rule

$$\frac{P \wedge R \Rightarrow P \quad \{P\} \quad c \quad \{Q\} \quad Q \Rightarrow Q \wedge R}{\{P \wedge R\} \quad c \quad \{Q \wedge R\}}$$

If the variables of R have not been modified by the operation (which, on the other hand, may modify the variables of P to produce Q), then the properties $P \wedge R \Rightarrow P$ and $Q \Rightarrow Q \wedge R$ hold.

4.1.6. Exercises

4.1.6.1. A serie of assignment

Compute by hand the weakest precondition of the following program:

```
1  /*@
2   requires -10 <= x <= 0 ;
3   requires 0 <= y <= 5 ;
4   ensures -10 <= \result <= 10 ;
5  */
6  int function(int x, int y){
7      int res ;
8      y += 10 ;
```

4. Basic instructions and control structures

```
9   x -= 5 ;
10  res = x + y ;
11  return res ;
12 }
```

Deduce that the program is correct with respect to its contract using the right rule.

4.1.6.2. Empty “then” branch in conditional

We previously shown that when a conditional structure has an empty “else” branch, that means that the conjunction of the precondition and the negation of the condition must be enough to verify the postcondition of the complete conditional structure. For both of the following question, we only need the inference rules and no WP calculus.

Show that when, instead of the “else” branch, the “then” branch is empty, the conjunction of the condition and the precondition must implies the postcondition of the “else” branch.

Show that when both branches are empty, the overall condition is just a skip operation.

4.1.6.3. Short circuit

C compilers implement short circuit for conditions. For example, that means that a code like this one (**without “else” block**) :

```
1  if(cond1 && cond2){
2    // code
3 }
```

can be written as:

```
1  if(cond1){
2    if(cond2){
3      // code
4    }
5 }
```

Show that on those two source code, the weakest precondition calculus generates an equivalent weakest precondition for equivalent for any code in the “then” block. Note that we assume the conditions to be pure expressions (without side-effects).

4.1.6.4. A larger program

Compute by hand the weakest precondition of the following program:

4. Basic instructions and control structures

```
1  /*@
2   requires -5 <= y <= 5 ;
3   requires -5 <= x <= 5 ;
4   ensures  -15 <= \result <= 25 ;
5  */
6  int function(int x, int y){
7      int res ;
8
9      if(x < 0){
10         x = 0 ;
11     }
12
13     if(y < 0){
14         x += 5 ;
15     } else {
16         x -= 5 ;
17     }
18
19     res = x - y ;
20
21     return res ;
22 }
```

Deduce that the program is correct with respect to its contract using the right rule.

4.2. Loops

Loops need a particular treatment in deductive verification of programs. These are the only control structures that will require important work from us. We cannot avoid this because without loops, it is difficult to write and prove interesting programs.

Before we look at the way we specify loop, we can answer to a rightful question: why are loops so complex?

4.2.1. Induction and invariant

The nature of loops makes their analysis complex. When we perform our reasoning, we need a rule to determine the precondition from a given sequence of instructions and a postcondition. Here, the problem is that we cannot *a priori* deduce how many times a loop iterates, and consequently, we cannot know how many times variables are modified.

We then proceed using an inductive reasoning. We have to find a property that is true before we start to execute the loop and that, if it is true at the beginning of an iteration, remains true at the end (and that is consequently true at the beginning of the next iteration). When the loop ends, we add the knowledge that the condition of the loop is false and that should allow us to deduce that the postcondition of the loop is verified.

This type of property is called a loop invariant. A loop invariant is a property that must be true before and after each loop iteration. And more precisely, each time the condition of the loop is checked. For example with the following loop:

4. Basic instructions and control structures

```
1  for(int i = 0 ; i < 10 ; ++i){ /* */ }
```

The property $0 \leq i \leq 10$ is a loop invariant. The property $-42 \leq i \leq 42$ is also an invariant (even if it is far less precise). The property $0 < i \leq 10$ is not an invariant because it is not true at the beginning of the execution of the loop. The property $0 \leq i < 10$ **is not a loop invariant**, it is not true at the end of the last iteration that sets the value of `i` to 10.

To verify an invariant I , WP then produces the following “reasoning”:

- verify that I is true at the beginning of the loop (establishment)
- verify that if I is true before an iteration, then I is true after (preservation).

4.2.1.1. Formal - Inference rule

Let us note the invariant I , the inference rule corresponding to loops is defined as follows:

$$\frac{\{I \wedge B\} c \{I\}}{\{I\} \text{ while}(B)\{c\} \{I \wedge \neg B\}}$$

And the weakest precondition calculus is the following:

$$wp(\text{while}(B)\{c\}, Post) := I \wedge ((B \wedge I) \Rightarrow wp(c, I)) \wedge ((\neg B \wedge I) \Rightarrow Post)$$

Let us detail this formula:

- (1) the first I corresponds to the establishment of the invariant, in layman’s terms, this is the “precondition” of the loop,
- the second part of the conjunction $((B \wedge I) \Rightarrow wp(c, I))$ corresponds to the verification of the operation performed by the body of the loop:
 - the precondition that we know of the loop body (let us note KWP , “Known WP”) is $(KWP = B \wedge I)$. That is the fact we have entered the loop (B is true), and that the invariant is verified at this moment (I , is true before we start the loop by (1), and we want to verify that it will be true at the end of the body of the loop in (2)),
 - (2) it remains to verify that KWP implies the actual precondition* of the body of the loop ($KWP \Rightarrow wp(c, Post)$). What we want at the end of the loop is the preservation of the invariant I (B is maybe not true anymore however), formally $KWP \Rightarrow wp(c, I)$, that is to say $(B \wedge I) \Rightarrow wp(c, I)$,
 - * it corresponds to the application of the consequence rule previously explained.
- finally, the last part $((\neg B \wedge I) \Rightarrow Post)$ expresses the fact that when the loop ends ($\neg B$), and the invariant I has been maintained, it must imply that the wanted postcondition of the loop is verified.

In this computation, we can notice that the wp function does not indicate any way to obtain the invariant I . We have to specify ourselves this property about our loops.

4. Basic instructions and control structures

4.2.1.2. Back to the WP plugin

There exist tools that can infer invariant properties (provided that these properties are simple, automatic tools remain limited). This is not the case for WP. We have to manually annotate our programs to specify the invariant of each loop. To find and write invariants for our loops will always be the hardest part of our work when we want to prove programs.

Indeed, when there are no loops, the weakest precondition calculus function can automatically provide the verifiable properties of our programs, this is not the case for loop invariant properties for which we do not have computation procedures. We have to find and express them correctly, and depending on the algorithm, they can be quite subtle and complex.

In order to specify a loop invariant, we add the following annotations before the loop:

```
1  int main(){
2    int i = 0;
3
4    /*@
5      loop invariant 0 <= i <= 30;
6    */
7    while(i < 30){
8      ++i;
9    }
10   //@assert i == 30;
11 }
```



REMINDER : The invariant is: $i \leq 30$!

Why? Because along the loop, i is comprised between 0 and **included** 30. 30 is indeed the value that allows us to leave the loop. Moreover, one of the properties required by the weakest precondition calculus is that when the loop condition is invalidated, by knowing the invariant, we can prove the postcondition (Formally $(\neg B \wedge I) \Rightarrow Post$).

The postcondition of our loop is $i = 30$ and must be implied by $\neg i < 30 \wedge 0 \leq i \leq 30$. Here, it is true since:

$$i \geq 30 \wedge 0 \leq i \leq 30 \Rightarrow i = 30$$

On the opposite, if we exclude the equality to 30, the postcondition would be unreachable.

Again, we can have a look at the list of verification conditions in “WP Goals”:

4. Basic instructions and control structures

```

int main(void)
{
  int __retres;
  int i;
  i = 0;
  /*@ loop invariant 0 ≤ i ≤ 30; */
  while (i < 30) {
    i ++;
  }
  /*@ assert i == 30; */ ;
  __retres = 0;
  return __retres;
}

```

Module	Goal	Model	Qed	Alt-Ergo	Coq	Why:
main	Invariant (preserved)	Typed	—	●		
main	Invariant (established)	Typed	●			
main	Assertion	Typed	●			

We notice that WP produces two different verification conditions: the establishment of the invariant and its preservation. WP produces exactly the reasoning we previously described to prove the assertion. In recent versions of Frama-C, Qed has become particularly aggressive and powerful, and the generated verification condition does not show those details (showing directly “True”). Using the option `-wp-no-simpl` at start, we can however see these details:

```

int main(void)
{
  int __retres;
  int i = 0;
  /*@ loop invariant 0 ≤ i ≤ 30; */
  while (i < 30) {
    i ++;
  }
  /*@ assert i == 30; */ ;
  __retres = 0;
  return __retres;
}

```

Information	Messages (0)	Console	Properties	Values	Red Alarms	WP Goals
<p>Global All Results</p> <p>Raw Obligation Proved Goal</p> <p><i>No Script</i></p> <hr/> <p>Goal Assertion:</p> <p>Assume {</p> <p> Type: is_sint32(i).</p> <p> (* Invariant *)</p> <p> Have: (0 ≤ i) /\ (i ≤ 30).</p> <p> (* Else *)</p> <p> Have: 30 ≤ i.</p> <p>}</p> <p>Prove: i = 30.</p> <hr/> <p>Prover Alt-Ergo: Valid (12ms) (18).</p>						

But is our specification precise enough?

4. Basic instructions and control structures

```
1 int main(){
2   int i = 0;
3   int h = 42;
4
5   /*@
6    loop invariant 0 <= i <= 30;
7   */
8   while(i < 30){
9     ++i;
10  }
11  //@assert i == 30;
12  //@assert h == 42;
13 }
```

And the result is:

```
int main(void)
{
  int __retres;
  int i;
  int h;
  i = 0;
  h = 42;
  /*@ loop invariant 0 ≤ i ≤ 30; */
  while (i < 30) {
    i ++;
  }
  /*@ assert i == 30; */ ;
  /*@ assert h == 42; */ ;
  __retres = 0;
  return __retres;
}
```

It seems not.

4.2.2. The assigns clause ... for loops

In fact, considering loops, WP **only** reasons about what is provided by the user to perform its reasoning. And here, the invariant does not specify anything about the way the value of `h` is modified (or not). We could specify the invariant of all program variables, but it would be a lot of work. ACSL simply allows to add `assigns` annotations for loops. Any other variable is considered to keep its old value. For example:

```
1 int main(){
2   int i = 0;
3   int h = 42;
4
5   /*@
6    loop invariant 0 <= i <= 30;
7    loop assigns i;
8   */
9   while(i < 30){
10    ++i;
11  }
12  //@assert i == 30;
13  //@assert h == 42;
14 }
```

4. Basic instructions and control structures

This time, we can establish the proof that the loop correctly behaves. However, we cannot prove that it terminates. The loop invariant alone does not give enough information to perform such a proof. For example, in our program, we could modify the loop, removing the loop body:

```
1  /*@
2    loop invariant 0 <= i <= 30;
3    loop assigns i;
4  */
5  while(i < 30){
6
7  }
```


The invariant is still verified, but we cannot prove that the loop ends: it is infinite.

4.2.3. Partial correctness and total correctness - Loop variant

In deductive verification, we find two types of correctness, the partial correctness and the total correctness. In the first case, the formulation of the correctness property is “if the precondition is valid, and **if** the computation terminates, then the postcondition is valid”. In the second case, “if the precondition is valid, **then** the computation terminates and the postcondition is valid”. By default, WP considers only partial correctness:

```
1  void foo(){
2    while(1){}
3    /*@ assert \false;
4  }
```

If we try to verify this code activating the verification of absence of RTE, we get this result:



```
void foo(void)
{
    while (1) {
    }
    /*@ assert \false; */ ;
    return;
}
```

Information Messages (0) Console Properties Values WP Goals

[kernel] Parsing infinite.c (with preprocessing)
[rte] annotating function bar
[rte] annotating function foo
[wp] [CFG] Goal foo_assert : Valid (Unreachable)
[wp] 0 goal scheduled
[wp] Proved goals: 0 / 0

The assertion “False” is proved! For a very simple reason: since the condition of the loop is “True” and no instruction of the loop body allows to leave the loop, it does not terminate. As we are proving the code with partial correctness, and as the execution does not terminate, we can prove anything about the code that follows the non terminating part of the code. However, if the termination is not trivially provable, the assertion will probably not be proved.

4. Basic instructions and control structures

i

Note that a (provably) unreachable assertion is always proved to be true:

```
void bar(void)
{
    goto End;
    /*@ assert \false; */ ;
    End: ;
    return;
}
```

Information	Messages (0)	Console	Properties	Values	WP Goals
[kernel] Parsing 3-3-goto_end.c (with preprocessing)					
[rte] annotating function bar					
[wp] Running WP plugin...					
[wp] [CFG] Goal bar_assert : Valid (Unreachable)					
[wp] 0 goal scheduled					
[wp] Proved goals: 0 / 0					

And this is also the case when we trivially know that an instruction produces a runtime error (for example dereferencing `NULL`), or inserting “False” in postcondition as we have already seen with `abs` and the parameter `INT_MIN`.

In order to prove the termination of a loop, we use the notion of loop variant. A loop variant is not a property but a value. It is an expression that involves the variables modified by the loop and that provides an upper bound to the number of iterations that remains to be executed by the loop before any iteration. Thus, this expression is greater or equals to 0, and strictly decreases at each loop iteration (this must also be verified by induction by WP).

If we take our previous example, we add the loop variant with this syntax:

```
1  int main(){
2      int i = 0;
3      int h = 42;
4
5      /*@
6          loop invariant 0 <= i <= 30;
7          loop assigns i;
8          loop variant 30 - i;
9      */
10     while(i < 30){
11         ++i;
12     }
13     //@assert i == 30;
14     //@assert h == 42;
15 }
```

Again, we can have a look at the generated verification conditions:

4. Basic instructions and control structures

```

int main(void)
{
    int __retres;
    int i;
    int h;
    i = 0;
    h = 42;
    /*@ loop invariant 0 ≤ i ≤ 30;
    loop assigns i;
    loop variant 30 - i; */
    while (i < 30) {
        i ++;
    }
    /*@ assert i == 30; */ ;
    /*@ assert h == 42; */ ;
    __retres = 0;
    return __retres;
}

```

Module	Goal	Model	Qed	Alt-Ergo	Coq
main	Invariant (preserved)	Typed	—	●	
main	Invariant (established)	Typed	●		
main	Assertion	Typed	●		
main	Assertion	Typed	●		
main	Loop assigns ...	Typed	●		
main	Loop variant at loop (decrease)	Typed	●		
main	Loop variant at loop (positive)	Typed	●		

The loop variant generates two verification conditions: verify that the value specified in the variant is positive, and prove that it strictly decreases during the execution of the loop. And if we delete the line of code that increments `i`, WP cannot prove anymore that `30 - i` strictly decreases.

We can also note that being able to give a loop variant does not necessarily induce that we can give the exact number of remaining iterations of the loop, as we do not always have a so precise knowledge of the behavior of the program. We can for example build an example like this one:

```

1  #include <stddef.h>
2
3  /*@
4   ensures min <= \result <= max;
5  */
6  size_t random_between(size_t min, size_t max);
7
8  void random_loop(size_t bound){
9      /*@
10     loop invariant 0 <= i <= bound ;
11     loop assigns i;
12     loop variant i;
13     */
14     for(size_t i = bound; i > 0; ){
15         i -= random_between(1, i);
16     }
17 }

```

Here, at each iteration, we decrease the value of the variable `i` by a value comprised between 1 and `i`. Thus, we can ensure that the value of `i` is positive and strictly decreases during each loop iteration, but we cannot say how many loop iterations remain to be executed.

4. Basic instructions and control structures

The loop variant is then only an upper bound on the number of iteration, not an expression of their exact number.

Note also that a loop variant only needs to be positive at the beginning of the execution of the block of the loop. Thus, in the following code:

```
1  int i = 5 ;
2  while(i >= 0){
3      i -= 2 ;
4  }
```

Even if `i` can be negative when the loop exits, it is still a variant since we do not start the execution of the block of the loop again.

4.2.4. Create a link between postcondition and invariant

Let us consider the following specified program. Our goal is to prove that this function returns the old value of `a` plus 10.

```
1  /*@
2   ensures \result == \old(a) + 10;
3  */
4  int add_ten(int a){
5      /*@
6       loop invariant 0 <= i <= 10;
7       loop assigns i, a;
8       loop variant 10 - i;
9      */
10     for (int i = 0; i < 10; ++i)
11         ++a;
12
13     return a;
14 }
```

The weakest precondition calculus does not allow to deduce information that is not part of the loop invariant. In a code like:

```
1  /*@
2   ensures \result == \old(a) + 10;
3  */
4  int add_ten(int a){
5      ++a;
6      ++a;
7      ++a;
8      //...
9      return a;
10 }
```

By reading the instructions backward from the postcondition, we always keep all knowledge about `a`. On the opposite, as we previously mentioned, outside the loop, WP only considers the information provided by the invariant. Consequently, our `add_10` function cannot be

4. Basic instructions and control structures

proved: the invariant does not say anything about `a`. To create a link between the postcondition and the invariant, we have to add this knowledge. See, for example:

```
1  /*@
2   ensures \result == \old(a) + 10;
3  */
4  int add_ten(int a){
5      /*@
6       loop invariant 0 <= i <= 10;
7       loop invariant a == \at(a, Pre) + i; //< ADDED
8       loop assigns i, a;
9       loop variant 10 - i;
10     */
11     for (int i = 0; i < 10; ++i)
12         ++a;
13
14     return a;
15 }
```

i

This need can appear as a very strong constraint. This is not really the case. There exists strongly automated analysis that can compute loop invariant properties. For example, without a specification, an abstract interpretation would easily compute `0 <= i <= 10` and `old(a) <= a <= \old(a)+10`. However, it is often more difficult to compute the relations that exist between the different variables of a program, for example the equality expressed by the invariant we have added and that is absolutely necessary to prove the postcondition of the function.

4.2.5. Early termination of loop

A loop invariant must be true each time the condition of the loop is checked. In fact, that also means that it must be true before an iteration, and after each **complete** iteration. Let us illustrate on an example this important idea.

```
1  int main(){
2      int i = 0;
3      int h = 42;
4
5      /*@
6       loop invariant 0 <= i <= 30;
7       loop assigns i;
8       loop variant 30 - i;
9     */
10     while(i < 30){
11         ++i;
12
13         if(i == 30) break ;
14     }
15     //@assert i == 30;
16     //@assert h == 42;
17 }
```

4. Basic instructions and control structures

In this function, when the loop reaches the index 30, we break the loop before checking the condition again. While the invariant is verified, let us show that we can now further constrain it.

```
1  int main(){
2      int i = 0;
3      int h = 42;
4
5      /*@
6          loop invariant 0 <= i <= 29;
7          loop assigns i;
8          loop variant 30 - i;
9      */
10     while(i < 30){
11         ++i;
12
13         if(i == 30) break ;
14     }
15     //@assert i == 30;
16     //@assert h == 42;
17 }
```

Here we can see that we have excluded 30 from the range of values of `i` and everything is still verified by WP. This is particularly interesting because it does not apply only to the invariant, none of the loop properties need to be verified in this last iteration. For example, we can write this code that is still verified:

```
1  int main(){
2      int i = 0;
3      int h = 42;
4
5      /*@
6          loop invariant 0 <= i <= 29;
7          loop assigns i;
8          loop variant 30 - i;
9      */
10     while(i < 30){
11         ++i;
12
13         if(i == 30){
14             i = 42 ;
15             h = 84 ;
16             break ;
17         }
18     }
19     //@assert i == 42;
20     //@assert h == 84;
21 }
```

We can see that we can write the variable `h` even if it is not listed in the `loop assigns` clause, and that we can give the value 42 to `i` which does not respect the invariant, and also makes the expression of the variant negative. In fact, everything happens as if we had written:

```
1  int main(){
2      int i = 0;
3      int h = 42;
4
```

4. Basic instructions and control structures

```
5  /*@
6     loop invariant 0 <= i <= 29;
7     loop assigns i;
8     loop variant 30 - i;
9  */
10 while(i < 29){
11     i++ ;
12 }
13
14 if(i < 30){
15     ++i;
16
17     if(i == 30){
18         i = 42 ;
19         h = 84 ;
20     }
21 }
22 //@assert i == 42;
23 //@assert h == 84;
24 }
```

This is an interesting scheme. It basically corresponds to any algorithm that searches, using a loop, a particular condition respected by an element in a given data-structure and stops when it finds it to perform some operations that are thus not really part of the loop. From a verification point of view, it allows us to simplify the contract of the loop: we know that the (potentially complex) operations performed just before we stop do not need to be considered when designing the invariant.

4.2.6. Exercises

4.2.6.1. Loop invariant

Write a suitable invariant for the following loop and prove it using the command.

```
1  frama-c -wp your-file.c
```

```
2  int x = 0 ;
3
4  while(x > -10){
5      -- x ;
6  }
```

Is the property $-100 \leq x \leq 100$ a correct invariant? Explain why.

4.2.6.2. Loop variant

Write a suitable invariant and variant for the follow loop and prove it using the command:

```
1  frama-c -wp your-file.c
```

4. Basic instructions and control structures

```
2  int x = -20 ;
3
4  while(x < 0){
5      x += 4 ;
6  }
```

If your variant does not precisely state the number of remaining iteration, add a variable that records exactly the number of remaining iterations and use it as a variant. You might need to add an invariant.

4.2.6.3. Loop assigns

Write a suitable loop assigns clause for this loop such that the assertion on line 9 is proved as well as the assigns clause. Let us ignore runtime errors in this proof.

```
2  int h = 42 ;
3  int x = 0 ;
4  int e = 0 ;
5  while(e < 10){
6      ++ e ;
7      x += e * 2 ;
8  }
9  //@ assert h == 42 ;
```

Once the proof succeeds, completely remove the assigns clause and find another way to ensure that the assertion is verified using annotations (note that you can add a C label in the code). What do you deduce about the notion of `loop assigns` clause?

4.2.6.4. Early termination

Write a suitable contract for this loop such that the assertions on lines 9 and 10 are proved as well as the contract of the loop.

```
1  void foo(){
2      int i ;
3      int x = 0 ;
4      for(i = 0 ; i < 20 ; ++i){
5          if(i == 19){
6              x++ ;
7              break ;
8          }
9      }
10     //@ assert x == 1 ;
11     //@ assert i == 19 ;
```

4.3. More examples on loops

4.3.1. Example with read-only arrays

The array is the most common data structure when we are working with loops. It is then a good example base to exercise with loops, and these examples allow to rapidly show interesting invariant and will allow us to introduce some important ACSL constructs.

We can for example use the search function that allows to find a value in an array:

```

1  #include <stddef.h>
2
3  /*@
4   requires \valid_read(array + (0 .. length-1));
5
6   assigns  \nothing;
7
8   behavior in:
9     assumes \exists size_t off ; 0 <= off < length && array[off] == element;
10    ensures array <= \result < array+length && *\result == element;
11
12   behavior notin:
13     assumes \forall size_t off ; 0 <= off < length ==> array[off] != element;
14    ensures \result == NULL;
15
16   disjoint behaviors;
17   complete behaviors;
18 */
19 int* search(int* array, size_t length, int element){
20   /*@
21    loop invariant 0 <= i <= length;
22    loop invariant \forall size_t j; 0 <= j < i ==> array[j] != element;
23    loop assigns i;
24    loop variant length-i;
25   */
26   for(size_t i = 0; i < length; i++)
27     if(array[i] == element) return &array[i];
28   return NULL;
29 }
```

There are enough ideas in this example to introduce some important syntax.

First, as we previously presented, the `\valid_read` predicate (as well as `\valid`) allows us to specify not only the validity of a readable address but also to state that a range of contiguous addresses is valid. It is expressed using this syntax:

```

1  //@ requires \valid_read(a + (0 .. length-1));
```

This precondition states that all addresses `a+0`, `a+1`, ..., `a+length-1` must be valid readable locations.

We also introduced two notations that are used almost all the time in ACSL, the keywords `\forall` (\forall) and `\exists` (\exists), the universal logic quantifiers. The first one allows to state that for any element, some property is true, the second one allows to say that there exists some element such that the property is true. If we comment a little bit the corresponding lines in our specification, we can read them this way:

4. Basic instructions and control structures

```
1 /*@
2 // for all "off" of type "size_t", IF "off" is comprised between 0 and "length"
3 //                                     THEN the cell "off" in "a" is different from "element"
4 \forall size_t off ; 0 <= off < length ==> a[off] != element;
5
6 // there exists "off" of type "size_t", such that "off" is comprised between 0 and "length"
7 //                                     AND the cell "off" in "a" equals to "element"
8 \exists size_t off ; 0 <= off < length && a[off] == element;
9 */
```

If we want to summarize the use of these keyword, we would say that on a range of values, a property is true, either about at least one of them or about all of them. A common scheme is to constrain this set using another property (here: `0 <= off < length`) and to prove the actual interesting property on this smaller set. **But using `exists` and `forall` is fundamentally different.**

With `\forall type a ; p(a) ==> q(a)`, the constraint `p` is followed by an implication. For any element where a first property `p` is verified, we have to also verify the second property `q`. If we use a conjunction, as we do for “exists” (which we will later explain), that would mean that all elements verify both `p` and `q`. In our previous example, it is clearly not the case as not all integers are comprised between 0 and `length`. Sometimes, it could be what we want to express, but it would then not correspond anymore to the idea of constraining a set for which we want to verify some other property.

With `\exists type a ; p(a) && q(a)`, the constraint `p` is followed by a conjunction. We say there exists an element such that it satisfies the property `p` at the same time it also satisfies `q`. If we use an implication, as we do for “forall”, such an expression will always be true if `p` is not a tautology! Why? Is there an “a” such that `p(a)` implies `q(a)`? Let us take any “a” such that `p(a)` is false, then the implication is true.

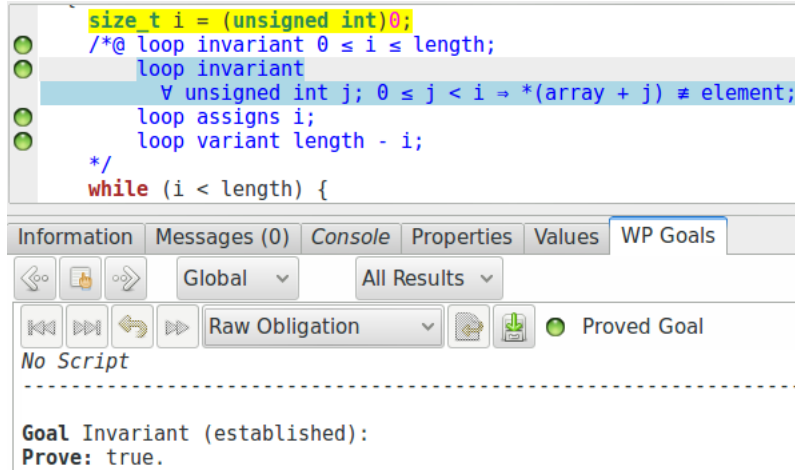
This part of the invariant deserves a particular attention:

```
1 // @ loop invariant \forall size_t j; 0 <= j < i ==> array[j] != element;
```

Indeed, it defines the treatment performed by our loop, it indicates to WP what happens in the loop (or more precisely: what we learn) along the execution. Here, this formula indicates that at each iteration of the loop, we know that for each memory location between 0 and the next location to visit (`i` excluded), the memory location contains a value different from the element we are looking for.

The verification condition associated to the preservation of this invariant is a bit complex and it is not really interesting to precisely look at it, on the contrary, the proof that the invariant is established before executing the loop is interesting:

4. Basic instructions and control structures



We note that this property, while quite complex, is proved easily by Qed. If we look at the parts of the programs on which the proof relies, we can see that the instruction `i = 0` is highlighted and this is, indeed, the last instruction executed on `i` before we start the loop. And consequently, if we replace the value of `i` by 0 inside the formula of the invariant, we get:

```
1  /*@ loop invariant \forall size_t j; 0 <= j < 0 ==> array[j] != element;
```

“For all j , greater or equal to 0 and strictly lower than 0”, this part of the formula is necessarily false, our implication is then necessarily true.

4.3.2. Examples with mutable arrays

Let us present two examples with mutation of arrays. One with a mutation of all memory locations, the other with selective modifications.

4.3.2.1. Reset

Let us have a look at the function that resets an array of integers to 0.

```
1  #include <stddef.h>
2
3  /*@
4   requires \valid(array + (0 .. length-1));
5   assigns array[0 .. length-1];
6   ensures \forall size_t i; 0 <= i < length ==> array[i] == 0;
7  */
8  void reset(int* array, size_t length){
9      /*@
10     loop invariant 0 <= i <= length;
11     loop invariant \forall size_t j; 0 <= j < i ==> array[j] == 0;
12     loop assigns i, array[0 .. length-1];
13     loop variant length-i;
14  */
15     for(size_t i = 0; i < length; ++i)
16         array[i] = 0;
```

4. Basic instructions and control structures

```
17 }
```

We can see that we use a construct for the invariant that is quite similar to what we have written for the previous example: we put an invariant to constraint the value of `i`, and we put another invariant that explains what the loop allowed to learn so far (all visited elements are now 0). Let us just highlight the `assigns` and `loop assigns` clauses: again, we can use the notation `n .. m` to indicate which parts of the array are modified.

4.3.2.2. Search and replace

The last example we will detail to illustrate the proof of functions with loops is the algorithm “search and replace”. This algorithm selectively modifies values in a range of memory locations. It is generally harder to guide the tool in such a case, because on one hand we must keep track of what is modified and what is not, and on the other hand, the induction relies on this fact.

As an example, the first specification we can write for this function could be this one:

```
1  #include <stddef.h>
2
3  /*@
4   requires \valid(array + (0 .. length-1));
5   assigns array[0 .. length-1];
6
7   ensures \forall size_t i; 0 <= i < length && \old(array[i]) == old
8           ==> array[i] == new;
9   ensures \forall size_t i; 0 <= i < length && \old(array[i]) != old
10          ==> array[i] == \old(array[i]);
11 */
12 void search_and_replace(int* array, size_t length, int old, int new){
13   /*@
14    loop invariant 0 <= i <= length;
15    loop invariant \forall size_t j; 0 <= j < i && \at(array[j], Pre) == old
16                  ==> array[j] == new;
17    loop invariant \forall size_t j; 0 <= j < i && \at(array[j], Pre) != old
18                  ==> array[j] == \at(array[j], Pre);
19    loop assigns i, array[0 .. length-1];
20    loop variant length-i;
21   */
22   for(size_t i = 0; i < length; ++i){
23     if(array[i] == old) array[i] = new;
24   }
25 }
```

We use the logic function `\at(v, Label)` that gives us the value of the variable `v` at the program point `Label`. If we look at the usage of this function here, we see that in the invariant we try to establish a relation between the old values of the array and the potentially new values:

```
1  /*@
2   loop invariant \forall size_t j; 0 <= j < i && \at(array[j], Pre) == old
3           ==> array[j] == new;
4   loop invariant \forall size_t j; 0 <= j < i && \at(array[j], Pre) != old
5           ==> array[j] == \at(array[j], Pre);
6  */
```

4. Basic instructions and control structures

For every memory location that contained the value to be replaced, it now must contain the new value. All other values must remain unchanged. While we previously relied on the `assigns` clauses for this kind of properties, here, we do not have this possibility because while the ACSL language provides some facilities that we could use to write a very precise `assigns` clause, WP would not perfectly exploit it. Thus we have to use an invariant and keep an approximation of the assigned memory location.

If we try to prove this invariant with WP, it fails. In such a case, the simpler method is to add different assertions that expresses the different intermediate properties using assertions which we expect to be easily proved and which imply the invariant. Here, we can easily notice that WP does not succeed in maintaining the knowledge that we have not modified the remaining part of the array yet:

```
1  for(size_t i = 0; i < length; ++i){
2      //@assert array[i] == \at(array[i], Pre); // proof failure
3      if(array[i] == old) array[i] = new;
4  }
```

We can add this information as an invariant:

```
13  /*@
14      loop invariant 0 <= i <= length;
15      loop invariant \forall size_t j; 0 <= j < i && \at(array[j], Pre) == old
16                      ==> array[j] == new;
17      loop invariant \forall size_t j; 0 <= j < i && \at(array[j], Pre) != old
18                      ==> array[j] == \at(array[j], Pre);
19      loop invariant \forall size_t j; i <= j < length
20                      ==> array[j] == \at(array[j], Pre);
21      loop assigns i, array[0 .. length-1];
22      loop variant length-i;
23  */
24  for(size_t i = 0; i < length; ++i){
25      if(array[i] == old) array[i] = new;
26  }
```

And this time the proof succeeds.

4.3.3. Exercises

For all these exercises, use the following command line to start verification:

```
1  frama-c-gui -wp -wp-rte -warn-unsigned-overflow your-file.c
```

4.3.3.1. Non-mutating: Forall, Exists, ...

Currently, function pointers are not supported by WP. Let us consider that we have a function:

4. Basic instructions and control structures

```
3  /*@
4   assigns \nothing ;
5   ensures \result <=> // some property about value
6  */
7  int pred(int value){
8      // your code
9  }
```

Write your own code and corresponding postcondition and then write the following functions together with their contract and prove their correctness. Note that, as `pred` is a C function, it cannot be used in the specification of the functions you have to write, thus the property that you have considered for this predicate should be inlined in the contract of the different functions.

- `forall_pred` returns true if and only if `pred` is true for all elements ;
- `exists_pred` returns true if and only if `pred` is true for at least one element ;
- `none_pred` returns true if and only if `pred` is false for all elements ;
- `some_not_pred` returns true if and only if `pred` is false for at least one element.

The two last functions should be written by just calling the two first ones.

4.3.3.2. Non-mutating: Equality of ranges

Write, specify and prove the function `equal` that returns true if and only if two ranges of values equal. Write, using the `equal` function, the code of `different` that returns true if two ranges are different, your postcondition should contain an existential quantifier.

```
3  int equal(const int* a_1, const int* a_2, size_t n){
4  }
5  }
6  }
7  int different(const int* a_1, const int* a_2, size_t n){
8  }
9  }
```

4.3.3.3. Binary Search

The following function searches for the position of a particular value in an array, assuming that this array is sorted. First, let us consider that the length of the array is an int and use int values to deal with indexes. One could note that there are two behavior: either the value exists in the array (and the result is the index of this value) or not (and the result is -1).

```
1  #include <stddef.h>
2
3  /*@
```

4. Basic instructions and control structures

```
4   requires Sorted:
5     \forall integer i, j ; 0 <= i <= j < len ==> arr[i] <= arr[j] ;
6   */
7   int bsearch(int* arr, int len, int value){
8     if(len == 0) return -1 ;
9
10    int low = 0 ;
11    int up = len-1 ;
12
13    while(low <= up){
14      int mid = low + (up - low)/2 ;
15      if (arr[mid] > value) up = mid-1 ;
16      else if(arr[mid] < value) low = mid+1 ;
17      else return mid ;
18    }
19    return -1 ;
20 }
```

This function is a bit tricky to prove, so let us provide some hints. First, this time the length is provided as an `int`, so we need to further restrict this value to make it coherent. Second, one of the invariant properties should state the bounds of `low` and `up`, but note that for both of them one of the bounds is not needed. Finally, the second invariant property should state that if some index of the array stores the value, this index must be correctly bounded.

Harder: Modify this function in order to receive `len` as a `size_t`. You will have to slightly modify the algorithm and the proof. As an hint, we advise making `up` an excluded bound for the search.

4.3.3.4. Mutating: Addition of vectors

Write, specify and prove the function that adds two vectors into a third one. Put arbitrary constraints to solve the integer overflow. Consider that the resulting vector is entirely separated from the input vectors, however the same vector should be usable for both input vectors.

```
3 void add_vectors(int* v_res, const int* v1, const int* v2, size_t len){
4
5 }
```

4.3.3.5. Mutating: Reverse

Write, specify and prove the function that reverses a vector in place. Take care of the unmodified part of the vector at some iteration of the loop. Use the previously proved `swap` function.

```
3 void swap(int* a, int* b);
4
5 void reverse(int* array, size_t len){
6
7 }
```

4. Basic instructions and control structures

4.3.3.6. Mutating: Copy

Write, specify and prove the function `copy` that copies a range of values into another array, starting from the first cell of the array. First consider (and specify) that the two ranges are entirely separated.

```
3 void copy(int const* src, int* dst, size_t len){  
4  
5 }
```

Harder: The true copy and copy backward functions

In fact, a strong separation is not necessary. The actual precondition must guarantee that if the two arrays overlap, the beginning of the destination range must not be in the source range:

```
1 //@ requires \separated(&src[0 .. len-1], dst) ;
```

In essence, by copying the element in that order, we can shift elements from the end of a particular range to the beginning. However, that means that we have to be more precise in our contract: we do not guarantee anymore an equality with the values of the source array but with the *old* values of the source array. And we have also to be more precise in our invariant, first by also specifying the relation in regard to the previous state of the memory, and second by adding an invariant that shows that the source array is not modified from the `i`th we are visiting to the end.

Finally, it is also possible to write a function that copies the elements from the end to the beginning, in this case, again, arrays can overlap but the condition is not exactly the same. Write, specify and prove the function `copy_backward` that copies elements in the reverse order.

4.4. Function calls

4.4.1. Calling a function

4.4.1.1. Formal - Weakest precondition calculus

When a function is called, the contract of this function is used to determine the precondition of the call. But one has to consider two important facts to express the weakest precondition calculus.

First, the postcondition of the called function f is not necessarily directly the precondition that was computed for the instructions that follow the call to f . For example, if we have a program: `x = f() ; c` and $wp(c, Q) = 0 \leq x \leq 10$, whereas the postcondition of the function `f` is $1 \leq x \leq 9$, we have to express some weakening between the actual precondition of `c` and the computed one. For this, we refer to the section 4.1.4, the idea is simply to verify that the postcondition of the function implies the computed precondition.

4. Basic instructions and control structures

Second, in C, a function can have side-effects. Thus, the values of the variables referenced in input are not necessarily the same as they were after the call to the function, and the contract may express some property about those values before and after the call. So, if we have labels in the postcondition, we must correctly replace them.

In order to define the weakest-precondition calculus of function calls, let us introduce some notation to make things clearer. For this, consider this example:

```

1  /*@ requires \valid(x) && *x >= 0 ;
2     assigns *x ;
3     ensures *x == \old(*x)+1 ; */
4  void inc(int* x);
5
6  void foo(int* a){
7      L1:
8      inc(a) ;
9      L2:
10 }

```

The weakest precondition of the function call asks us to consider the contract of the function that is called (here, in `foo`, when we call the `inc` function). Of course, before the call to the function we have to verify its precondition, so it is part of the weakest precondition. But, we also have to consider the postcondition of the function, else that would mean that we do not consider its effect.

Thus, it is important to notice that in the precondition, the considered memory state is the one where we compute the weakest precondition, whereas for the postcondition it is not the case, the considered memory state is the one that follows the call, while we need to explicitly mention the old state to speak about the values before the call. For example, considering the contract of `inc` when we call it in `foo`, `*x` in the precondition is `*a` at `L1`, while `*x` in the postcondition is `*a` at `L2`. Consequently the pre and the postcondition must be considered slightly differently when it comes to mutable memory location. Note that for the value of the parameter `x` itself, there is no such consideration: this value cannot be modified by the call.

Now, let us define the weakest precondition of a function call. For this, we denote:

- \vec{v} a vector of values v_1, \dots, v_n and v_i the i^{th} value,
- \vec{t} the arguments provided to the function when we call it,
- \vec{x} the parameters in the function definition,
- \vec{a} the assigned values (seen from the outside, once instantiated),
- $here(x)$ a value in postcondition
- $old(x)$ a value in precondition

We name **f:Pre** the precondition of the function, and **f:Post** the postcondition:

$$\begin{aligned}
 wp(f(\vec{t}), Q) := & \text{f:Pre}[x_i \leftarrow t_i] \\
 & \wedge \forall \vec{v}, \quad (\text{f:Post}[x_i \leftarrow t_i, here(a_j) \leftarrow v_j, old(a_j) \leftarrow a_j] \Rightarrow Q[here(a_j) \leftarrow v_j])
 \end{aligned}$$

4. Basic instructions and control structures

We can detail a little bit the reasoning for each part of this formula.

First, note that in both pre and postcondition, each named parameter x_i is replaced with the corresponding argument ($[x_i \leftarrow t_i]$), as we said before we do not have to consider memory states there because those values cannot be changed by the function call. For example in the contract of `inc`, each `x` would be replaced by the argument `a`.

Then, in the part of the formula that corresponds to the postcondition, we can see that we introduce a $\forall \vec{v}$. The goal is here to model the fact that the function can write any value in each memory location that is assigned. So, for each of the assigned location a_j (that is for our call to `inc`, `*(&a)`), we generate a value v_j that is its value after the call. But, if we want to check that the postcondition gives us the right result, we cannot accept *any value* for each assigned location, we just want the ones *that allows to satisfy the postcondition*.

So these values are used to transform the postcondition of the function and verify that it implies the postcondition in input of the weakest precondition. This is done by replacing, for each assigned location a_j , its value *here* with the value v_j that it is supposed to get after the call ($here(a_j) \leftarrow v_j$). Finally, we have to replace each *old* value by its value before the call, and for each $old(a_j)$, it is simply a_j ($old(a_j) \leftarrow a_j$).

4.4.1.2. Formal - Example

Let us illustrate this on an example by applying the weakest precondition calculus to this short code, assuming the contract we previously proposed for the `swap` function.

```

1  int a = 4 ;
2  int b = 2 ;
3
4  swap(&a, &b) ;
5
6  //@ assert a == 2 && b == 4 ;

```

We now compute the weakest precondition:

$$\begin{aligned}
wp(a = 4; b = 2; swap(\&a, \&b), a = 2 \wedge b = 4) = \\
wp(a = 4, wp(b = 2; swap(\&a, \&b), a = 2 \wedge b = 4)) = \\
wp(a = 4, wp(b = 2, wp(swap(\&a, \&b), a = 2 \wedge b = 4)))
\end{aligned}$$

Let us first consider separately:

$$wp(swap(\&a, \&b), a = 2 \wedge b = 4)$$

From this `assigns` clause, we know that the assigned values are $\&a = a$ and $\&b = b$. (Let us shorten *here* with H and *old* with O).

4. Basic instructions and control structures

$$\begin{aligned}
& \text{swap:Pre}[x \leftarrow \&a, y \leftarrow \&b] \\
& \wedge \forall v_a, v_b, (\text{swap:Post} \quad [x \leftarrow \&a, y \leftarrow \&b, \\
& \quad H(*(\&a)) \leftarrow v_a, H(*(\&b)) \leftarrow v_b, \\
& \quad O(*(\&a)) \leftarrow *(\&a), O(*(\&b)) \leftarrow *(\&b)]) \\
& \Rightarrow (H(a) = 2 \wedge H(b) = 4)[H(a) \leftarrow v_a, H(b) \leftarrow v_b]
\end{aligned}$$

For the precondition, we get :

$$\text{valid}(\&a) \wedge \text{valid}(\&b)$$

For the postcondition part, let us first write the expression from which we start before any term replacement (and without the syntax for the replacement for the sake of conciseness):

$$H(*x) = O(*y) \wedge H(*y) = O(*x) \Rightarrow H(a) = 2 \wedge H(b) = 4$$

First we replace the pointers ($x \leftarrow \&a, y \leftarrow \&b$) :

$$H(*(\&a)) = O(*(\&b)) \wedge H(*(\&b)) = O(*(\&a)) \Rightarrow H(a) = 2 \wedge H(b) = 4$$

Then, the *here* values, with the quantified v_i s ($H(a) \leftarrow v_a, H(b) \leftarrow v_b$):

$$v_a = O(*(\&b)) \wedge v_b = O(*(\&a)) \Rightarrow v_a = 2 \wedge v_b = 4$$

And the *old* values, with the value before call ($O(*(\&a)) \leftarrow *(\&a), O(*(\&b)) \leftarrow *(\&b)$):

$$v_a = *(\&b) \wedge v_b = *(\&a) \Rightarrow v_a = 2 \wedge v_b = 4$$

We can now simplify this formula to:

$$v_a = b \wedge v_b = a \Rightarrow v_a = 2 \wedge v_b = 4$$

So, $wp(\text{swap}(\&a, \&b), a = 2 \wedge b = 4)$ is:

$$P : \text{valid}(\&a) \wedge \text{valid}(\&b) \wedge \forall v_a, v_b, \quad v_a = b \wedge v_b = a \Rightarrow v_a = 2 \wedge v_b = 4$$

Let us immediately simplify the formula by noticing that validity properties are trivially true here (since the variable are allocated on the stack just before):

$$P : \forall v_a, v_b, \quad v_a = b \wedge v_b = a \Rightarrow v_a = 2 \wedge v_b = 4$$

Let us now compute $wp(a = 4, wp(b = 2, P))$, by first replacing b with 2 by the assignment rule:

$$\forall v_a, v_b, \quad v_a = 2 \wedge v_b = a \Rightarrow v_a = 2 \wedge v_b = 4$$

and then replacing a with 4 by the same rule:

$$\forall v_a, v_b, \quad v_a = 2 \wedge v_b = 4 \Rightarrow v_a = 2 \wedge v_b = 4$$

This last property is trivially true, thus the program is verified.

4.4.1.3. What should we keep in mind?

Functions are absolutely necessary to modular programming, and the weakest precondition calculus is fully compatible with this idea, allowing to reason about each function locally and compose proofs just as we compose function calls.

So as a reminder, we should just keep in mind the following general scheme:

4. Basic instructions and control structures

```
1  /*@
2    requires foo_R ;
3    assigns ... ;
4    ensures foo_E ;
5  */
6  type foo(parameters...){
7    // Here we suppose that foo_R holds
8
9
10   // Here we must prove that bar_R holds
11   bar(some parameters ...) ;
12   // Here we assume that bar_E holds
13
14
15   // Here we must prove that foo_E holds
16   return ... ;
17 }
```

Note that for the last statement, with weakest precondition calculus, the idea is more to show that our precondition is strong enough to ensure that the code leads to our postcondition. However, first, this vision is simpler to understand, and second the WP plugin does not actually perform a strict weakest precondition calculus but an highly optimized one that does not follows exactly the same rules.

4.4.2. Recursive functions

For now, WP does not check function termination. Of course, if a function is only composed of loops that terminate (that have a verified variant) and calls to functions that terminate, it terminates. However, one particular case requires more reasoning: recursive and mutually recursive functions. Currently termination of such functions is not supported with WP.

That basically means that using a function that does not terminate we can prove anything. For example:

```
1  /*@
2    assigns \nothing ;
3    ensures \false ;
4  */
5  void trick(){
6    trick() ;
7  }
8
9  int main(){
10   trick();
11   //@ assert \false ;
12 }
```

4. Basic instructions and control structures

```
int main(void)
{
    int __retres;
    trick();
    /*@ assert \false; */ ;
    __retres = 0;
    return __retres;
}

/*@ ensures \false;
   assigns \nothing; */
void trick(void)
{
    trick();
    return;
}
```

We can see that the function and the assertion are proved. And indeed the proof is correct: we consider partial correctness and we face a function that does not terminate: anything that follows a call to this function would be true.

Thus, the question is: what could we do in such a case? Again, we could use some kind of variant to bound the number of recursive calls. In ACSL, this is the role of the `decreases` clause:

```
1  /*@
2    decreases n ;
3  */
4  void ends(int n){
5      if(n > 0) ends(n-1);
6  }
```

This clause expresses exactly the same idea as a `loop variant`. The expression considered by a `decreases` clause is a positive expression that strictly decreases when the function is called again. However, it is still not supported by WP. Thus, for the moment, a recursive function cannot be totally proved with WP.

4. *Basic instructions and control structures*

In this part, we have seen how assignment and control structure are translated to a logic view of our program. We have spent quite a lot of time on loops because they represent the main difficulty we have to face when we want to specify and prove a program by deductive verification. The loop annotations allow us to express as precisely as possible their behavior.

In the next part of this tutorial, we will see more precisely the logic constructs provided by ACSL. They are important because they give us a way to write more abstract specification, that are easier to understand and to prove.

5. ACSL - Properties

From the beginning of this tutorial, we have used different predicates and logic functions provided by ACSL: `\valid`, `\valid_read`, `\separated`, `\old` and `\at`. There are others built-in predicates but we will not present them all, the reader can refer to [the documentation \(ACSL implementation\)](#) [↗](#) (note that everything is not necessarily supported by WP).

ACSL allows us to do something more than “just” specify our code using existing predicates and functions. We can define our own predicates, functions, relations, etc. Doing this, we can have more abstract specifications. It also allows us to factor specifications (for example defining what is a valid array), which have two pleasant consequences: our specifications are more readable and more understandable, and we can reuse existing proofs to ease the proof of new programs.

5.1. Some logical types

ACSL provides different logic types that allow us to write properties in a more abstract, mathematical world. Among the types that can be useful, some are dedicated to numbers, and allow to express properties or functions without having to think about constraints due to the size of the representation of primitive C types in memory. These types are `integer` and `real`, which respectively represent mathematical integers and reals (that are modeled to be as close to the reality we can, but this notion cannot be perfectly handled).

From now, we will often use integers instead of classical C `int`s. The reason is simply that a lot of properties and definitions are true regardless the size of the machine integer we have as input.

On the other hand, we will not talk about the differences that exist between `real` and `float/double`. It would require to speak about precise numerical calculus, and about proofs of programs that rely on such calculus which could deserve an entire dedicated tutorial.

5.2. Predicates

A predicate is a property about different objects that can be true or false. To sum up, we are writing predicates from the beginning of this tutorial in precondition, postcondition, assertion and loop invariant. ACSL allows us to name these predicates, as we could do for a boolean function in C, for example. An important difference, however, is that predicates (as well as logic functions that we will see later) must be pure. For example, they cannot produce side effects by modifying a pointed value.

5. ACSL - Properties

These predicates can receive some parameters. Moreover, they can also receive some C labels that will allow us to establish relations between different program points.

5.2.1. Syntax

Predicates are introduced using ACSL annotations. The syntax is the following:

```
1  /*@
2  predicate named_predicate { Lb10, ..., Lb1N }(type0 arg0, ..., typeN argN) =
3  //a logic relations between all these things
4  */
```

For example, we can define the predicate that checks whether an integer in memory is changed between two particular program points:

```
1  /*@
2  predicate unchanged{L0, L1}(int* i) =
3  \at(*i, L0) == \at(*i, L1);
4  */
```



Keep in mind that passing a value to a predicate is done, as it is done in C, by value. We cannot write this predicate by directly passing `i` in parameter:

```
1  /*@
2  predicate unchanged{L0, L1}(int i) =
3  \at(i, L0) == \at(i, L1);
4  */
```

Since `i` is just a copy of the received variable.

We can verify this code using our predicate:

```
6  int main(void){
7      int i = 13;
8      int j = 37;
9
10     Begin:
11     i = 23;
12
13     //@assert ! unchanged{Begin, Here}(&i);
14     //@assert  unchanged{Begin, Here}(&j);
15 }
```

We can also have a look at the verification conditions generated by WP and notice that, even it is slightly (syntactically) modified, the predicate is not unrolled by WP. The provers will determine themselves whether they need to use the definition of the predicate to establish the proof.

5. ACSL - Properties

```

int main(void)
{
    int __retres;
    int i = 13;
    int j = 37;
    Begin: i = 23;
    /*@ assert ~unchanged{Begin, Here}(&i); */ ;
    /*@ assert unchanged{Begin, Here}(&j); */ ;
    __retres = 0;
    return __retres;
}

```

Information Messages (0) Console Properties Values Red Alarms WP Goals

Global All Results

Raw Obligation Binary Proved Goal

No Script

Goal Assertion:
 Let a = global(L_i_26).
 Assume {
 (* Heap *)
 Have: linked(Malloc_0).
 (* Initializer *)
 Init: Mint_0[a] = 13.
 (* Initializer *)
 Init: Mint_0[global(L_j_27)] = 37.
 }
 Prove: !P_unchanged(Mint_0[a <- 23], Mint_0, a).

As we said earlier, one important use of predicates (and logic functions) is to make our specifications more readable and to factor it. An example can be to write a predicate that expresses the validity of an array in reading or writing. It allows us to avoid writing the complete expression every time we need it and to make it readable quickly:

```

3  /*@
4   predicate valid_range_rw(int* t, integer n) =
5     n >= 0 && \valid(t + (0 .. n-1));
6
7   predicate valid_range_r(int* t, integer n) =
8     n >= 0 && \valid_read(t + (0 .. n-1));
9  */
10
11 /*@
12  requires 0 < length;
13  requires valid_range_r(array, length);
14  //...
15 */
16 int* search(int* array, size_t length, int element);

```

In this specification, we do not give an explicit label to predicates for their definition, nor for their use. For the definition, Frama-C automatically creates an implicit label. At predicate use, the given label is implicitly **Here**. The fact we do not explicitly define the label in the definition of a predicate does not forbid to explicitly give a label when we use it.

Of course, predicates can be defined in header files in order to produce a utility library for specification for example.

5. ACSL - Properties

5.2.1.1. Predicate overloading

It is possible to overload predicates as long as the types of the parameters are different or the number of parameters changes. For example, we can redefine the `valid_range_r` as a predicate that takes in parameters both the beginning and the end of the range to consider. Then, we can write an overloaded version that uses the previous one for the particular case of ranges that starts at 0:

```
3  /*@
4   predicate valid_range_r(int* t, integer beg, integer end) =
5       end >= beg && \valid_read(t + (beg .. end-1)) ;
6
7   predicate valid_range_r(int* t, integer n) =
8       valid_range_r(t, 0, n) ;
9  */
10
11 /*@
12  requires 0 < length;
13  requires valid_range_r(array, length);
14  //...
15 */
16 int* search(int* array, size_t length, int element);
```

5.2.2. Abstraction

An other important use of predicates is to define the logical state of our data structures when programs start to be more complex. Our data structures must usually respect an invariant (again) that each manipulation function must maintain in order to ensure that the data structure will always remain coherent and usable through future calls.

It allows us to ease the reading of specifications. For example, we can define the specification required to ensure the safety of a fixed size stack. It could be done as illustrated here (note that we do not provide the definition of the predicates as it is not the purpose of our example, the careful reader could consider this as an exercise):

```
1  #include <stddef.h>
2  #define MAX_SIZE 42
3
4  struct stack_int{
5      size_t top;
6      int    data[MAX_SIZE];
7  };
8
9  /*@
10   predicate valid_stack_int(struct stack_int* s) = \true ; // to define
11   predicate empty_stack_int(struct stack_int* s) = \true ; // to define
12   predicate full_stack_int(struct stack_int* s) = \true ; // to define
13  */
14
15 /*@
16  requires \valid(s);
17  assigns *s;
18  ensures valid_stack_int(s) && empty_stack_int(s);
19  */
20 void initialize(struct stack_int* s);
21
```

5. ACSL - Properties

```
22  /*@
23   requires valid_stack_int(s) && !full_stack_int(s);
24   assigns *s;
25   ensures valid_stack_int(s);
26  */
27  void push(struct stack_int* s, int value);
28
29  /*@
30   requires valid_stack_int(s) && !empty_stack_int(s);
31   assigns \nothing;
32  */
33  int top(struct stack_int* s);
34
35  /*@
36   requires valid_stack_int(s) && !empty_stack_int(s);
37   assigns *s;
38   ensures valid_stack_int(s);
39  */
40  void pop(struct stack_int* s);
41
42  /*@
43   requires valid_stack_int(s);
44   assigns \nothing;
45   ensures \result == 1 <==> empty_stack_int(s);
46  */
47  int is_empty(struct stack_int* s);
48
49
50  /*@
51   requires valid_stack_int(s);
52   assigns \nothing;
53   ensures \result == 1 <==> full_stack_int(s);
54  */
55  int is_full(struct stack_int* s);
```

Here, the specification does not express functional properties. For example, we do not specify that when we perform the push of a value, and then we ask for the top of the stack, we get the same value. But we already have enough details to ensure that, even if we cannot prove that we always get the right result (behaviors such as “if I push v , top returns v ”), we can still guarantee that we do not produce runtime errors (if we provide correct predicates for the stack, and prove that the implementation of our functions ensures that no runtime errors can occur).

5.2.3. Exercises

5.2.3.1. Days of the month

Taking back the solution of the exercise 3.4.1.1 about days of the month, write a predicate to express that a year is leap and adapt the contracts using it.

5.2.3.2. Alpha-numeric character

Taking back the solution of the exercise 3.4.1.2 about alpha numeric characters, write predicates to express that a character is an upper letter, lower letter, and a digit. Adapt the contracts of the different functions using them.

5. ACSL - Properties

5.2.3.3. Max of 3 values

The following function returns the max of 3 input values:

```
1 int max_of(int* a, int* b, int* c){
2     if(*a >= *b && *a >= *c) return *a ;
3     if(*b >= *a && *b >= *c) return *b ;
4     return *c ;
5 }
```

Write a predicate to express that a value is one of three pointed values at a given memory state:

```
1 /*@
2  predicate one_of{L}(int value, int *a, int *b, int *c) =
3      // ...
4  */
```

Use the set notation. Write a contract to the function and prove that it is verified.

5.2.3.4. Binary Search

Taking back the solution of the exercise 4.3.3.3 about the binary search function with unsigned types, write a predicate that expresses that an array is sorted on a range of values starting at `begin` and ending at `end` (excluded). Overload this predicate in order to make `begin` optional with a default value of 0. Define a predicate that checks if an element is in a range of values of an array starting at index `begin` and ending at `end` (excluded), again overload this predicate to make the first bound optional.

Use those two predicates to simplify the contract of the function. Note that both behaviors `assumes` clause should be modified.

5.2.3.5. Search and replace

Taking back the example 4.3.2.2, about the search and replace function, write predicates that express that in some range of an array starting at index `begin` and ending at `end` (excluded), values

- remain unchanged between two labels,
- are replaced with some new value when it equals to some old value, then left unchanged

Overload both predicates to make the first bound optional. Use the obtained predicates to simplify the contract and loop invariant of the function.

5.3. Logic functions

Logic functions are meant to describe functions that can only be used in specifications. It allows us, first, to factor those specifications and, second, to define some operations on **integer** or **real** with the guarantee that they cannot overflow since they involve mathematical types.

Like predicates, they can receive different labels and values in parameter.

5.3.1. Syntax

To define a logic function, the syntax is the following:

```

1  /*@
2    logic return_type my_function\{ Label0, ..., LabelN \}( type0 arg0, ..., typeN argN ) =
3      formula using the arguments ;
4  */

```

We can for example define a mathematical [linear function](#) φ using a logic function:

```

1  /*@
2    logic integer ax_b(integer a, integer x, integer b) =
3      a * x + b;
4  */

```

And it can be used to prove the source code of the following function:

```

6  /*@
7    assigns \nothing ;
8    ensures \result == ax_b(3,x,4);
9  */
10 int function(int x){
11   return 3*x + 4;
12 }

```

```

/*@ logic Z ax_b(Z a, Z x, Z b) = a * x + b;

*/
/*@ ensures \result == ax_b(3, \old(x), 4);
   assigns \nothing; */
int function(int x)
{
  int __retres;
  /*@ assert rte: signed_overflow: (int)(3 * x) + 4 ≤ 2147483647; */
  /*@ assert rte: signed_overflow: -2147483648 ≤ 3 * x; */
  /*@ assert rte: signed_overflow: 3 * x ≤ 2147483647; */
  __retres = 3 * x + 4;
  return __retres;
}

```

5. ACSL - Properties

This code is indeed proved but some runtime-errors seems to be possible. We can again define some mathematical logic function that will provide the bounds of the linear function according to the machine type we use (from a logic point of view). It allows us to then add our bounds checking in the precondition of the function.

```
8  /*@
9   logic integer limit_int_min_ax_b(integer a, integer b) =
10     (a == 0) ? (b > 0) ? INT_MIN : INT_MIN-b :
11     (a < 0) ? (INT_MAX-b)/a :
12             (INT_MIN-b)/a ;
13
14   logic integer limit_int_max_ax_b(integer a, integer b) =
15     (a == 0) ? (b > 0) ? INT_MAX-b : INT_MAX :
16     (a < 0) ? (INT_MIN-b)/a :
17             (INT_MAX-b)/a ;
18 */
19
20 /*@
21   requires INT_MIN <= 3 * x ;
22   requires limit_int_min_ax_b(3,4) < x < limit_int_max_ax_b(3,4);
23   assigns \nothing ;
24   ensures \result == ax_b(3,x,4);
25 */
26 int function(int x){
27   return 3*x + 4;
28 }
```

i

Note that, as in specifications, computations are done using mathematical integers. We then do not need to care about some overflow risk with the computation of `INT_MIN-b` or `INT_MAX-b`.

Once this specification is provided, everything is fine. Of course, we could manually provide these bounds every time we create a linear logic function. But, by creating these bound computation functions, we directly get a way to compute them automatically which is quite comfortable.

Note that we also give a lower bound to the computation of `3 * x`, indeed while the bound provided for `x` by our logic function is defined for the complete computation, it does not say anything about the value obtained in the intermediate computation. For example here, the fact that `3 * x + 4` is not lower than `INT_MIN` does not guarantee that this is the case for `3 * x`.

5.3.2. Recursive functions and limits of logic functions

Logic functions (as well as predicates) can be recursively defined. However, such an approach will rapidly show some limits in their use for program proof. Indeed, when the automatic solver reasons on such logic properties, if such a function is met, it is necessary to evaluate it. SMT solvers are not meant to be efficient for this task, thus it is generally costly, producing too long proof resolution and eventually timeouts.

We can have a concrete example with the factorial function, using logic and using C language:

5. ACSL - Properties

```
1  /*@
2   logic integer factorial(integer n) = (n <= 0) ? 1 : n * factorial(n-1);
3  */
4
5  /*@
6   assigns \nothing ;
7   ensures \result == factorial(n) ;
8  */
9  int facto(int n){
10   if(n < 2) return 1 ;
11
12   int res = 1 ;
13   /*@
14    loop invariant 2 <= i <= n+1 ;
15    loop invariant res == factorial(i-1) ;
16    loop assigns i, res ;
17    loop variant n - i ;
18   */
19   for(int i = 2 ; i <= n ; i++){
20     res = res * i ;
21   }
22   return res ;
23 }
```

Without checking overflows, this function is easy and fast to prove. If we add runtime error checking, we see that there is a possibility of overflow on the multiplication.

On `int`, the maximum value for which we can compute factorial is 12. If we go further, it overflows. We can then add this precondition:

```
5  /*@
6   requires n <= 12 ;
7   assigns \nothing ;
8   ensures \result == factorial(n) ;
9  */
10 int facto(int n){
```

If we ask for a proof on this input, Alt-ergo will probably fail, whereas Z3 can compute the proof in less than a second. The reason is that in this case, the heuristics that are used by Z3 consider that it is a good idea to spend a bit more time on the evaluation of the function. We can for example change the maximum value of `n` to see how the different provers behave. With an `n` fixed to 9, Alt-ergo produces a proof in less than 10 seconds, whereas with a value of 10, even a minute is not enough.

Logic functions can then be defined recursively but without some more help, we are rapidly limited by the fact that provers need to perform evaluation or to “reason” by induction, two tasks for which they are not efficient. This can limit our possibilities for program proofs, but we will see later that we can get rid of these problems.

5.3.3. Exercises

5.3.3.1. Distance

Specify and prove the following program:

5. ACSL - Properties

```
1 int distance(int a, int b){
2     if(a < b) return b - a ;
3     else return a - b ;
4 }
```

For this, define two logic functions `abs` and `distance`. Use these functions to write the specification of the function.

5.3.3.2. Square

Write the body of the `square` function. Specify and prove the program. Use a `square` logic function.

```
1 int abs(int x){
2     return (x < 0) ? -x : x ;
3 }
4
5 unsigned square(int x){
6
7 }
```

Take care of the types of the variables and do not over-constrain the input of the function. Furthermore, when verifying the absence of runtime errors, do not forget to provide the option `-warn-unsigned-overflow`.

5.3.3.3. Iota

Here is a possible implementation of the iota function:

```
1 #include <limits.h>
2 #include <stddef.h>
3
4 void iota(int* array, size_t len, int value){
5     if(len){
6         array[0] = value ;
7
8         for(size_t i = 1 ; i < len ; i++){
9             array[i] = array[i-1]+1 ;
10        }
11    }
12 }
```

Write a logic function that returns the input value increased by one. Prove that after the execution of `iota`, the first value of the array is the input value and that each value of the array corresponds to the value that precedes it increased by one (using the previously defined logic function).

5.3.3.4. Vector add

In the following program, the `vec_add` function adds the second vector in input into the first one. Write a contract for the function `show_the_difference` that expresses, for each value of the vector `v1` the difference between the pre and the postcondition. For this, define a logic function `diff` that returns the difference between the value of a memory location at a label `L1` and the value at a label `L2`.

```

1  #include <stddef.h>
2  #include <limits.h>
3
4  /*@
5   predicate unchanged{L1, L2}(int* ptr, integer a, integer b) =
6     \forall integer i ; a <= i < b ==> \at(ptr[i], L1) == \at(ptr[i], L2) ;
7  */
8
9  /*@
10 requires \valid(v1 + (0 .. len-1));
11 requires \valid_read(v2 + (0 .. len-1));
12 requires \separated(v1 + (0 .. len-1), v2 + (0 .. len-1));
13 requires
14   \forall integer i ; 0 <= i < len ==> INT_MIN <= v1[i]+v2[i] <= INT_MAX ;
15
16 assigns v1[0 .. len-1];
17
18 ensures
19   \forall integer i ; 0 <= i < len ==> v1[i] == \old(v1[i]) + v2[i] ;
20 ensures
21   \forall integer i ; 0 <= i < len ==> v2[i] == \old(v2[i]) ;
22 */
23 void vec_add(int* v1, const int* v2, size_t len){
24   /*@
25     loop invariant 0 <= i <= len ;
26     loop invariant
27       \forall integer j ; 0 <= j < i ==> v1[j] == \at(v1[j], Pre) + v2[j] ;
28     loop invariant unchanged{Pre, Here}(v1, i, len) ;
29     loop assigns i, v1[0 .. len-1] ;
30     loop variant len-i ;
31   */
32   for(size_t i = 0 ; i < len ; ++i){
33     v1[i] += v2[i] ;
34   }
35 }
36
37 void show_the_difference(int* v1, const int* v2, size_t len){
38   vec_add(v1, v2, len);
39 }

```

Re-express the `unchanged` predicate using the logic function you have defined.

5.3.3.5. The sum of the N first integers

The following function compute the sum of the N first integers. Write a recursive logic functions that returns the sum of the N first integers and write a specification for the C function expressing that it computes the same value as provided by the logic function.

5. ACSL - Properties

```
1 int sum_n(int n){
2   if(n < 1) return 0 ;
3
4   int res = 0 ;
5   for(int i = 1 ; i <= n ; i++){
6     res += i ;
7   }
8   return res ;
9 }
```

Try to verify the absence of runtime errors. The integer overflow is not so simple to get rid of. However, write a precondition that should be enough to prove the function (remember that the sum of the N first integers can be expressed with a really simple formula ...). It will certainly not be enough to directly prove the absence of overflow, but we will see how to provide such an information in the next section.

5.4. Lemmas

Lemmas are general properties about predicates or functions. These properties can be proved in isolation of the rest of the proof of a program by automatic or (more often) interactive provers. Once this proof is done, the information that it states can be safely used to simplify the reasoning in other, more complex proofs, without having to prove it again. For example, if we state a lemma L that says $P \Rightarrow Q$ in any case, if at some point in another proof, we have to prove Q and we know P , we can directly conclude by using the lemma L without having to perform again the reasoning that brings us from P to Q .

In the previous section, we said that recursive function can make proof harder for SMT solvers. In such a case, lemmas can help us. We can write by ourselves the proofs that require inductive reasoning for some properties that we state as lemmas, and these lemmas can be used efficiently by SMT solvers to perform the other proofs, related to our programs.

5.4.1. Syntax

Again, we introduce lemmas using ACSL annotations. The syntax is following:

```
1 /*@
2   lemma name_of_the_lemma { Label0, ..., LabelN }:
3     property ;
4 */
```

This time, the properties we want to express do not depend on received parameters (except for labels). So we express these properties for universally quantified variables. For example, we can state this lemma, which is true, even if it is trivial:

5. ACSL - Properties

```
1  /*@
2    lemma lt_plus_lt:
3      \forall integer i, j ; i < j ==> i+1 < j+1;
4  */
```

This proof can be performed using WP. The property is, of course, proved using only Qed.

5.4.2. Example: properties of linear functions

We can come back to our linear functions and express some interesting properties about them:

```
20 /*@
21 lemma ax_b_monotonic_neg:
22   \forall integer a, b, i, j ;
23     a < 0 ==> i <= j ==> ax_b(a, i, b) >= ax_b(a, j, b);
24 lemma ax_b_monotonic_pos:
25   \forall integer a, b, i, j ;
26     a > 0 ==> i <= j ==> ax_b(a, i, b) <= ax_b(a, j, b);
27 lemma ax_b_monotonic_nul:
28   \forall integer a, b, i, j ;
29     a == 0 ==> ax_b(a, i, b) == ax_b(a, j, b);
30 */
```

For these proofs, Alt-ergo, will probably not be able to discharge all generated verification conditions. In this case, Z3 will certainly perform it. We can then write the following example code:

```
32 /*@
33   requires INT_MIN <= a*x <= INT_MAX ;
34   requires limit_int_min_ax_b(a,4) < x < limit_int_max_ax_b(a,4);
35   assigns \nothing ;
36   ensures \result == ax_b(a,x,4);
37 */
38 int function(int a, int x){
39   return a*x + 4;
40 }
41
42 /*@
43   requires INT_MIN <= a*x <= INT_MAX ;
44   requires INT_MIN <= a*y <= INT_MAX ;
45   requires a > 0;
46   requires limit_int_min_ax_b(a,4) < x < limit_int_max_ax_b(a,4) ;
47   requires limit_int_min_ax_b(a,4) < y < limit_int_max_ax_b(a,4) ;
48   assigns \nothing ;
49 */
50 void foo(int a, int x, int y){
51   int fmin, fmax;
52   if(x < y){
53     fmin = function(a,x);
54     fmax = function(a,y);
55   } else {
56     fmin = function(a,y);
57     fmax = function(a,x);
58   }
59   //@assert fmin <= fmax;
60 }
```

5. ACSL - Properties

If we do not give the lemmas provided earlier, Alt-ergo will not be able to prove the proof that `fmin` is lower or equals to `fmax`. With the lemmas it is however very easy for it since the property is the simply an instance of the lemma `ax_monotonic_pos`, the proof is then trivial as our lemma is considered to be true when are not currently proving it. Note that on this generalized version of `function`, Z3 will be probably more efficient to prove the absence of runtime errors.

5.4.3. Example: arrays and labels

Later in this tutorial, we will see some kind of definitions for which it is sometimes hard to reason about for SMT solvers when some mutations happen in memory. Thus, we will often need lemmas to state relations about the content of the memory between labels.

For now, let us illustrate with a first simple example. Consider the two following predicates:

```
1  /*@
2  predicate sorted(int* array, integer begin, integer end) =
3    \forall integer i, j ; begin <= i <= j < end ==> array[i] <= array[j] ;
4
5  predicate unchanged{L1, L2}(int *array, integer begin, integer end) =
6    \forall integer i ; begin <= i < end ==>
7      \at(array[i], L1) == \at(array[i], L2) ;
8  */
```

One could for example want to state that when an array is sorted, and some mutations happen in memory (creating a new memory state), but the content of the array remains unchanged, then the array is still sorted. This can be done with the following lemma:

```
10 /*@
11 lemma unchanged_sorted{L1, L2}:
12   \forall int* array, integer b, integer e ;
13     sorted{L1}(array, b, e) ==>
14       unchanged{L1, L2}(array, b, e) ==>
15         sorted{L2}(array, b, e) ;
16 */
```

We state this lemma for two labels `L1` and `L2`, and express that if any range in any array is sorted at `L1`, and unchanged from `L1` to `L2`, then it is still sorted at `L2`.

Note that this lemma is easily proved by SMT solvers. We will see later some examples where it is not so easy to get a proof.

5.4.4. Exercises

5.4.4.1. Multiplication property

Write a lemma that state that for three integers x , y and z , if x is greater or equals to 0, if z is greater or equals to y , then $x * z$ is greater or equals to $y * z$.

5. ACSL - Properties

This lemma will not be proved by SMT solvers, however if you ask a proof with Coq, the default tactic will probably discharge this verification condition automatically.

5.4.4.2. Locally sorted to globally sorted

The following program contains a function that requires an array to be sorted in the sense that each element is lower or equals to the element that follows it and calls the binary search function.

```
59 /*@
60   predicate element_level_sorted(int* array, integer fst, integer end) =
61     \forall integer i ; fst <= i < end-1 ==> array[i] <= array[i+1] ;
62 */
63 /*@
64   //lemma element_level_sorted_implies_sorted:
65   // ...
66 */
67
68 /*@
69   requires \valid_read(arr + (0 .. len-1));
70   requires element_level_sorted(arr, 0, len) ;
71   requires in_array(value, arr, len);
72
73   assigns \nothing ;
74
75   ensures 0 <= \result < len ;
76   ensures arr[\result] == value ;
77 */
78 unsigned bsearch_callee(int* arr, size_t len, int value){
79   return bsearch(arr, len, value);
80 }
```

Take back your proved binary search function from the exercise 5.2.3.4. You might notice that the precondition of the binary search function is stronger than what we know in precondition of the `bsearch_callee`. However, our precondition implies that the array is globally sorted. Write a lemma that states that if an array is `element_level_sorted` then it is `sorted`. This lemma will probably not be proved by SMT solvers, all remaining properties should be.

We provide a solution and the corresponding Coq proof on the GitHub repository of this book.

5.4.4.3. Sum of the N first integers

Take back your solution to the exercise 5.3.3.5 about the sum of the N first integers. Write a lemma that states that the result of the call to the logic function is $n * (n + 1) / 2$. This lemma will not be proved by SMT solvers.

We provide a solution and the corresponding Coq proof on the GitHub repository of this book.

5.4.4.4. Shift transitivity

The following program is composed of two functions. The first one is the `shift_array` function that shifts the elements of an array with a given offset (named `shift`). The second performs two successive shifts on the same array.

```

1  #include <stddef.h>
2  #include <limits.h>
3
4  /*@
5   // predicate shifted{L1, L2}(int* arr, integer fst, integer last, integer shift) =
6   // ...
7
8   // predicate unchanged{L1, L2}(int *a, integer begin, integer end) =
9   // ...
10
11  // lemma shift_ptr{...}:
12  // ...
13
14  // lemma shift_transitivity{...}:
15  // ...
16 */
17
18 void shift_array(int* array, size_t len, size_t shift){
19     for(size_t i = len ; i > 0 ; --i){
20         array[i+shift-1] = array[i-1] ;
21     }
22 }
23
24 /*@
25  requires \valid(array+(0 .. len+s1+s2-1)) ;
26  requires s1+s2 + len <= UINT_MAX ;
27  assigns array[s1 .. s1+s2+len-1];
28  ensures shifted{Pre, Post}(array, 0, len, s1+s2) ;
29 */
30 void double_shift(int* array, size_t len, size_t s1, size_t s2){
31     shift_array(array, len, s1) ;
32     shift_array(array+s1, len, s2) ;
33 }

```

Complete the predicates `shifted` and `unchanged` (the last should use the former). Complete the contract of the `shift_array` function and prove it using WP.

Express two lemmas about the `shifted` property.

The first one `shift_ptr` should state that shifting a range `fst+s1` to `last+s1` of an `array` with an offset `s2` is equivalent to shifting the range `fst` to `last` of the memory location `array+s1` with an offset `s2`.

The second one should state that when the elements of an array are shifted a first time with an offset `s1` and then a second time with an offset `s2`, then the complete shift corresponds to a single shift with an offset `s1+s2`.

The lemma `shift_ptr` will not be proved by SMT solvers, we provide a solution and the corresponding Coq proof on the GitHub repository of this book. All remaining properties should be proved automatically.

5.4.4.5. Shift sorted range

The following program is composed of two functions. The function `shift_and_search` shifts the element of an array and then performs a binary search.

5. ACSL - Properties

```
1 size_t bsearch(int* arr, size_t beg, size_t end, int value);
2
3 /*@
4 // lemma shifted_still_sorted{...}:
5 // ...
6 */
7
8 /*@
9 requires sorted(array, 0, len) ;
10 requires \valid(array + (0 .. len));
11 requires in_array(value, array, 0, len) ;
12
13 assigns array[1 .. len] ;
14
15 ensures 1 <= \result <= len ;
16 */
17 unsigned shift_and_search(int* array, size_t len, int value){
18   shift_array(array, len, 1);
19   return bsearch(array, 1, len+1, value);
20 }
```

Take back your proved binary search function from the exercise 5.2.3.4, modify the binary search function, its contract and its proof in order to be able to search in any range.

Use the `shift_array` function proved in the previous exercise.

Complete the contract of the function `shift_and_search`. You might notice that the precondition that requires the array to be sorted is not validated. Complete the lemma `shifted_still_sorted` that should state that if a range is sorted at some label and then shifted, the resulting range is still sorted.

The lemma does not have to be proved by SMT solvers. We provide a solution and the corresponding Coq proof on the GitHub repository of this book.

5. ACSL - Properties

In this part of the tutorial, we have seen different ACSL constructs that allow us to factor our specifications and to express general properties that can be used by our solver to make their task easier.

All techniques we have talk about are safe, since they do not *a priori* allow us to write false or contradictory specifications. At least if the specification only use such logic constructions and if every lemma, precondition (at call site), every postcondition, assertion, variant and invariant are correctly proved, the code is correct.

However, sometimes, such constructions are not enough to express all properties we want to express about our programs. The next constructions we will see give us some new possibilities about it, but it will be necessary to be really careful using them since an error would allow us to introduction false assumptions or silently modify the program we are verifying.

6. ACSL - Logic definitions and code

In this part of the tutorial, we will present three important notions of ACSL:

- inductive definitions,
- axiomatic definitions,
- ghost code.

In some cases, these notions are absolutely needed to ease the process of specification and, more importantly, proof. On one hand they force some properties to be more abstract when an explicit modeling would involve too much computation during proof. On the other hand, they force some properties to be explicitly modeled when they are harder to reason about when they are implicit.

Using these notions, we expose ourselves to the possibility to make our proof irrelevant if we make mistakes writing specification with it. Inductive predicates and axiomatic definitions involve the risk to introduce “false” in our assumptions, or to write imprecise definitions. Ghost code opens the risk to silently modify the verified program ... making us prove another program, which is not the one we want to prove.

6.1. Inductive definitions

Inductive predicates gives a way to state properties whose verification requires to reason by induction, that is to say: for a property $p(input)$, it is true for some base cases (for example, 0 is an even natural number), and knowing that $p(input)$ is true, it is also true for a *bigger input*, provided that some conditions relating *input* and *bigger input* are verified (for example, knowing that n is even, we define that $n + 2$ is also even). Thus, inductive predicates are generally useful to define properties recursively.

6.1.1. Syntax

For now, let us introduce the syntax of inductive predicates:

```
1  /*@
2   inductive property{ Label0, ..., LabelN }(type0 a0, ..., typeN aN) {
3   case c_1{Lq_0, ..., Lq_X}: p_1 ;
4   ...
5   case c_m{Lr_0, ..., Lr_Y}: p_km ;
6   }
7  */
```

6. ACSL - Logic definitions and code

Where all `c_i` are names and all `p_i` are logic formulas where `property` appears has a conclusion. Basically, `property` is true for some parameters and some labels, if it corresponds to one of the cases of the inductive definition.

We can have a look at the simple property we just talked about: how to define that a natural number is even using induction. We can translate the sentence: “0 is an even natural number, and if n is an even natural number, $n + 2$ is an even natural number”:

```
1 /*@
2   inductive even_natural{L}(integer n) {
3     case even_nul{L}:
4       even_natural(0);
5     case even_not_nul_natural{L}:
6       \forall integer n ;
7         even_natural(n) ==> even_natural(n+2);
8   }
9 */
```

Which perfectly describes the two cases:

- 0 is even (base case),
- if a natural n is even, $n + 2$ is also even.

However, this definition is not completely satisfying. For example, we cannot deduce that an odd number is not even. If we try to prove that 1 is even, we will have to check if -1 is even, and then -3, -5, etc, infinitely. Moreover, we generally prefer to define the induction cases the opposite way: defining the condition under which the wanted conclusion is true. For example, here, how can we verify that some n is natural and even? We first check whether it is 0, if not, we check if n is greater than 0 and then we verify that $n - 2$ is even:

```
1 /*@
2   inductive even_natural{L}(integer n) {
3     case even_nul{L}:
4       even_natural(0) ;
5     case even_not_nul_natural{L}:
6       \forall integer n ; n > 0 ==> even_natural(n-2) ==>
7         even_natural(n) ;
8   }
9 */
```

Here, we distinguish two cases:

- 0 is even,
- a natural n is even if it is greater than 0 and $n - 2$ is an even natural.

Taking the second case, we recursively decrease the number until we reach 0, and then the number is even, since 0 is even, or -1, and then there is no case in the inductive that corresponds to this value, so we could deduce that the property is false (however, we will see later that in the case of WP, we need the Coq proof assistant).

6. ACSL - Logic definitions and code

```
11 void foo(){
12     int a = 42 ;
13     //@ assert even_natural(a);
14 }
```

Of course, defining that some natural number is even inductively is not a good idea, since we can simply define it using modulo. We generally use them to define complex recursive properties.

6.1.1.1. Consistency

Inductive definitions bring the risk to introduce inconsistencies. Indeed, the property specified in the different cases are considered to be always true, so if we introduce a property that allows to prove `false`, we will be able to prove everything. While we will give more details about axioms in the Section 6.2, let us give two hints to avoid building such a bad definition.

First, we can make sure that inductive predicates are well founded. This can be done by syntactically restricting what we allow in an inductive definition, by making sure that each case has the form:

```
1 /*@
2   \forall y1,...,ym ; h1 ==> ... ==> h1 ==> P(t1,...,tn) ;
3 */
```

where the predicate `P` can only appear positively (so not negated with `!` - \neg) in the different hypotheses `hx`. Basically, it ensures that we cannot build both positive and negative occurrences of `P` for some parameters which would be incoherent.

This is for example verified by our previously defined predicate `even_natural`. While a definition like:

```
1 /*@
2   inductive even_natural{L}(integer n) {
3     case even_nul{L}:
4       even_natural(0) ;
5     case even_not_nul_natural{L}:
6       \forall integer n ; n > 0 ==> even_natural(n-2) ==>
7       // negative occurrence of even_natural
8       !even_natural(n-1) ==>
9       even_natural(n) ;
10  }
11 */
```

does not respect this constraint as the property `even_natural` appears negatively on line 8.

Second, it is possible to ask the generation of a Coq version of a proof obligation that needs the predicate `P`. For example, we can write a function:

6. ACSL - Logic definitions and code

```

1  /*@
2    requires P( params ... ) ;
3    ensures  BAD: \false ;
4  */ void function(params){
5
6  }

```

For example for our definition of `even_natural`, we would write:

```

13 /*@
14   requires even_natural(n) ;
15   ensures  \false ;
16 */ void function(int n){
17
18 }

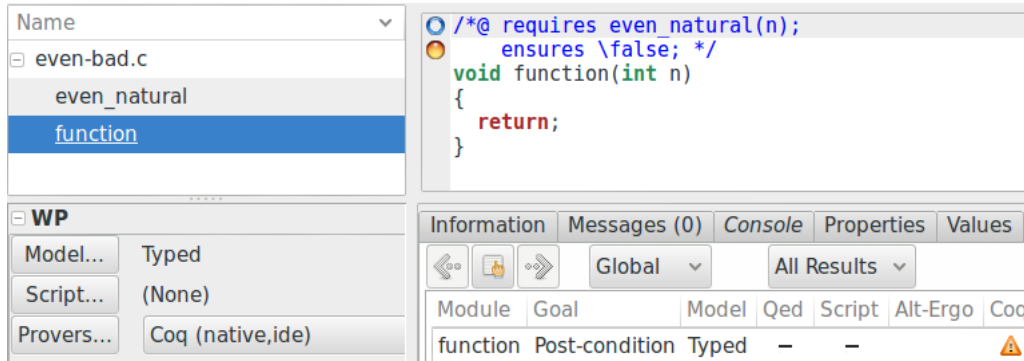
```

As Coq is stricter than Frama-C and WP on this kind of definition, if the inductive predicate is too strange (if it is not well founded), it will be rejected with an error. And indeed, with the bad definition of `even_natural` we just proposed, Coq complains when we try to prove the `ensures \false`, because there exists a non positive occurrence of `P_even_natural`, that encodes the `even_natural` we wrote in ACSL.

```

1  frama-c-gui -wp -wp-prop=BAD -wp-prover=native:coq file.c

```



```

Information Messages (0) Console Properties Values Red Alarms WP
[kernel] Parsing even-bad.c (with preprocessing)
[wp] Warning: Missing RTE guards
[wp] 1 goal scheduled
[wp] Proved goals: 0 / 1
[wp] [Coq] 'Axiomatic.v' compilation failed.

--- Coqc (stderr) :

File "/home/allan/.frama-c-wp/typed/Axiomatic.v", line 18, characters 0-250:
Error: Non strictly positive occurrence of "P_even_natural" in
"forall i : int,
0 < i ->
P_even_natural (i - 2) -> ~ P_even_natural (i - 1) -> P_even_natural i".

[wp] [Coq] Goal typed_function_post : Failed
Compilation of 'Axiomatic.v' failed.

```

6. ACSL - Logic definitions and code

However, that does not mean that we cannot write a inconsistent inductive definition. For example, the following definition is well founded, while it allows us to prove false:

```
1 /*@
2   inductive bad {
3     case always: \false ;
4   }
5 */
```

Thus we must be careful when defining inductive definitions to not introduce a definition that can lead to a proof of false.

6.1.2. Recursive predicate definitions

Inductive predicates are often useful to express recursive properties since it prevents the provers to unroll the recursion when it is possible.

For example, if we want to define that an array only contains 0s, we could write it as follows:

```
3 /*@
4   inductive zeroed{L}(int* a, integer b, integer e){
5     case zeroed_empty{L}:
6       \forall int* a, integer b, e; b >= e ==> zeroed{L}(a,b,e);
7     case zeroed_range{L}:
8       \forall int* a, integer b, e; b < e ==>
9         zeroed{L}(a,b,e-1) && a[e-1] == 0 ==> zeroed{L}(a,b,e);
10    }
11 */
```

And we can again prove that our reset to 0 is correct with this new definition:

```
14 /*@
15   requires \valid(array + (0 .. length-1));
16   assigns array[0 .. length-1];
17   ensures zeroed(array,0,length);
18 */
19 void reset(int* array, size_t length){
20   /*@
21     loop invariant 0 <= i <= length;
22     loop invariant zeroed(array,0,i);
23     loop assigns i, array[0 .. length-1];
24     loop variant length-i;
25   */
26   for(size_t i = 0; i < length; ++i)
27     array[i] = 0;
28 }
```

Depending on the Frama-C or automatic solvers versions, the proof of the preservation of the invariant could fail. A reason for this fail is the fact that the prover forgets that cells preceding the one we are currently processing are actually still set to 0. We can add a lemma in our knowledge base, stating that if a set of values of an array did not change between two program points, the first one being a point where the property "zeroed" is verified, then the property is still verified at the second program point.

6. ACSL - Logic definitions and code

```
13 /*@
14 predicate same_elems{L1,L2}(int* a, integer b, integer e) =
15   \forall integer i; b <= i < e ==> \at(a[i],L1) == \at(a[i],L2);
16
17 lemma no_changes{L1,L2}:
18   \forall int* a, integer b, e;
19     same_elems{L1,L2}(a,b,e) ==> zeroed{L1}(a,b,e) ==> zeroed{L2}(a,b,e);
20 */
```

Then we can add an assertion to specify what did not change between the beginning of the loop block (pointed by the label `L` in the code) and the end (which is `Here` since we state the property at the end):

```
35 for(size_t i = 0; i < length; ++i){
36   L:
37     array[i] = 0;
38     //@ assert same_elems{L,Here}(array,0,i);
39 }
```

Note that in this new version of the code, the property stated by our lemma is not proved using automatic solver, that cannot reason by induction. If the reader is curious, the (quite simple) Coq proof can be found in 6.4.

In this case study, using an inductive definition is not efficient: our property can be perfectly expressed using the basic constructs of the first order logic as we did before. Inductive definitions are meant to be used to write definitions that are not easy to express using the basic formalism provided by ACSL. It is here used to illustrate their use with a simple example.

6.1.3. Example: sort

Let us prove a simple selection sort:

```
3 size_t min_idx_in(int* a, size_t beg, size_t end){
4   size_t min_i = beg;
5   for(size_t i = beg+1; i < end; ++i)
6     if(a[i] < a[min_i]) min_i = i;
7   return min_i;
8 }
9
10 void swap(int* p, int* q){
11   int tmp = *p; *p = *q; *q = tmp;
12 }
13
14 void sort(int* a, size_t beg, size_t end){
15   for(size_t i = beg; i < end; ++i){
16     size_t imin = min_idx_in(a, i, end);
17     swap(&a[i], &a[imin]);
18   }
19 }
```

The reader can exercise by specifying and proving the search of the minimum and the swap function. We hide there a correct version of these specification and code (Answers 6.4), we will

6. ACSL - Logic definitions and code

focus on the specification and the proof of the sort function that is a interesting use case for inductive definitions.

Indeed, a common error we could do, trying to prove a sort function, would be to write this specification (which is correct!):

```
6  /*@
7   predicate sorted(int* a, integer b, integer e) =
8     \forall integer i, j; b <= i <= j < e ==> a[i] <= a[j];
9  */
10
11 /*@
12  requires \valid(a + (beg .. end-1));
13  requires beg < end;
14  assigns  a[beg .. end-1];
15  ensures sorted(a, beg, end);
16 */
17 void sort(int* a, size_t beg, size_t end){
18   /* @ // add invariant */
19   for(size_t i = beg ; i < end ; ++i){
20     size_t imin = min_idx_in(a, i, end);
21     swap(&a[i], &a[imin]);
22   }
23 }
```

This specification is correct. But if we recall correctly the part of the tutorial about specifications, we have to *precisely* express what we expect of the program. With this specification, we do not prove all required properties expected for a sort function. For example, this function correctly answers to the specification:

```
8  /*@
9   requires \valid(a + (beg .. end-1));
10  requires beg < end;
11
12  assigns  a[beg .. end-1];
13
14  ensures sorted(a, beg, end);
15 */
16 void fail_sort(int* a, size_t beg, size_t end){
17   /*@
18   loop invariant beg <= i <= end;
19   loop invariant \forall integer j; beg <= j < i ==> a[j] == 0;
20   loop assigns i, a[beg .. end-1];
21   loop variant end-i;
22   */
23   for(size_t i = beg ; i < end ; ++i)
24     a[i] = 0;
25 }
```

Our specification does not express the fact that all elements initially present in the array must still be in the array after executing the sort function. That is to say: the sort function produces a sorted permutation of the original array.

Defining the notion of permutation can be done using an inductive definition. While we will see later a version of this property that is more general, let us for now limit us to a notion of permutation that is more specific to our current need. We can limit us to a few cases. First, the array is a permutation of itself, then swapping two values of the array produces a new

6. ACSL - Logic definitions and code

permutation if we do not change anything else. And finally, if we create the permutation p_2 of p_1 , and then the permutation p_3 of p_2 , then by transitivity p_3 is a permutation of p_1 .

The corresponding inductive definition is the following:

```

37  /*@
38  predicate swap_in_array{L1,L2}(int* a, integer b, integer e, integer i, integer j) =
39      b <= i < e && b <= j < e &&
40      \at(a[i], L1) == \at(a[j], L2) &&
41      \at(a[j], L1) == \at(a[i], L2) &&
42      \forall integer k; b <= k < e && k != j && k != i ==>
43          \at(a[k], L1) == \at(a[k], L2);
44
45  inductive permutation{L1,L2}(int* a, integer b, integer e){
46  case reflexive{L1}:
47      \forall int* a, integer b,e ; permutation{L1,L1}(a, b, e);
48  case swap{L1,L2}:
49      \forall int* a, integer b,e,i,j ;
50          swap_in_array{L1,L2}(a,b,e,i,j) ==> permutation{L1,L2}(a, b, e);
51  case transitive{L1,L2,L3}:
52      \forall int* a, integer b,e ;
53          permutation{L1,L2}(a, b, e) && permutation{L2,L3}(a, b, e) ==>
54              permutation{L1,L3}(a, b, e);
55  }
56  */

```

We can then specify that our sort function produces the sorted permutation of the original array and we can then prove it by providing the invariant of the function:

```

64  /*@
65  requires beg < end && \valid(a + (beg .. end-1));
66  assigns a[beg .. end-1];
67  ensures sorted(a, beg, end);
68  ensures permutation{Pre, Post}(a,beg,end);
69  */
70  void sort(int* a, size_t beg, size_t end){
71      /*@
72      loop invariant beg <= i <= end;
73      loop invariant sorted(a, beg, i) && permutation{Pre, Here}(a, beg, end);
74      loop invariant \forall integer j,k; beg <= j < i ==> i <= k < end ==> a[j] <= a[k];
75      loop assigns i, a[beg .. end-1];
76      loop variant end-i;
77      */
78      for(size_t i = beg ; i < end ; ++i){
79          /*@ ghost begin: ;
80          size_t imin = min_idx_in(a, i, end);
81          swap(&a[i], &a[imin]);
82          /*@ assert swap_in_array{begin,Here}(a,beg,end,i,imin);
83      }
84  }

```

This time, our property is precisely defined, the proof is relatively easy to produce, only requiring to add an assertion in the block of the loop to state that it only performs a swap of values in the array (and then giving the transition to the next permutation). To define this swap notion, we use a particular annotation (at line 16), introduced using the keyword `ghost`. Here, the goal is to introduce a label in the code that in fact does not exist from the program point of view, and is only visible from a specification point of view. We present the “ghost” features in the final section of this chapter, for now let us focus on axiomatic definitions.

6.1.4. Exercises

6.1.4.1. Sum of the N first integers

Take back your solution to the exercise 5.4.4.3 about the sum of the N first integers. Rewrite the logic function using an inductive predicate that states that some integer equals the sum of the N first integers.

```

1  #include <limits.h>
2
3  /*@
4   inductive is_sum_n(integer n, integer res) {
5     // ...
6   }
7  */
8
9  /*@
10 requires n*(n+1) <= 2*INT_MAX ;
11 assigns \nothing ;
12 // ensures ... ;
13 */
14 int sum_n(int n){
15     if(n < 1) return 0 ;
16
17     int res = 0 ;
18     /*@
19     loop invariant 1 <= i <= n+1 ;
20     // loop invariant ... ;
21     loop assigns i, res ;
22     loop variant n - i ;
23     */
24     for(int i = 1 ; i <= n ; i++){
25         res += i ;
26     }
27     return res ;
28 }

```

Adapt the contract of the function and the lemma(s). Note that lemma(s) could certainly not be proved by SMT solvers. We provide a solution and corresponding Coq proofs on the GitHub repository of this book.

6.1.4.2. Greatest Common Divisor

Write an inductive predicate that states that some integer is the greatest common divisor of two others. The goal of the exercise is to prove that the function `gcd` computes the greatest common divisor. Thus, we do not have to specify all the cases for the predicate. Indeed, a close look at the loop shows us that after the first iteration `a` is greater or equals to `b` and it is maintained by the loop. Thus, we consider two cases for the inductive predicate:

- `b` is 0,
- if some `d` is the GCD of `b` and `a % b`, then it is the GCD of `a` and `b`

6. ACSL - Logic definitions and code

```
1  #include <limits.h>
2
3  /*@ inductive is_gcd(integer a, integer b, integer div) {
4      case gcd_zero: // ...
5      case gcd_succ: // ...
6  }
7  */
8
9  /*@
10     requires a >= 0 && b >= 0 ;
11     assigns \nothing ;
12     // ensures ... ;
13 */
14 int gcd(int a, int b){
15     /*@
16         // loop invariant \forall integer t ; ... ;
17     */
18     while (b != 0){
19         int t = b;
20         b = a % b;
21         a = t;
22     }
23     return a;
24 }
```

Express the postcondition of the function, and complete the invariant to prove that it is verified. Note that the invariant should make use of the inductive case `gcd_succ`.

6.1.4.3. Power function

In this exercise, we do not consider RTEs.

Write an inductive predicate that states that some integer `r` equals to x^n . Consider the two cases: either n is 0 or it is greater and then it should be related to the value x^{n-1} .

```
1  /*@
2     inductive is_power(integer x, integer n, integer r) {
3     case zero: // ...
4     case N: // ...
5     }
6  */
```

First prove the naive version of the power function:

```
13 /*@
14     requires n >= 0 ;
15     // assigns ...
16     // ensures ...
17 */
18 int power(int x, int n){
19     int r = 1 ;
20     /*@
21         loop invariant 1 <= i <= n+1 ;
22         // loop invariant ...
23     */
24     for(int i = 1 ; i <= n ; ++i){
25         r *= x ;
26     }
```

6. ACSL - Logic definitions and code

```
26   }
27   return r ;
28 }
```

Now, let us prove a faster version of the power function:

```
30  /*@
31   requires n >= 0 ;
32   // assigns ...
33   // ensures ...
34  */
35  int fast_power(int x, int n){
36    int r = 1 ;
37    int p = x ;
38    /*@
39     loop invariant \forall integer v ; // ...
40    */
41    while(n > 0){
42      if(n % 2 == 1) r = r * p ;
43      p *= p ;
44      n /= 2 ;
45    }
46    //@ assert is_power(p, n, 1) ;
47
48    return r ;
49  }
```

In this version, we exploit two properties about the power operator:

- $(x^2)^n = x^{2n}$
- $x * (x^2)^n = x^{2n+1}$

that allows us to divide n by 2 at each step of the loop instead of decreasing it by one (which makes the algorithm $O(\log n)$ instead of $O(n)$). Express the two previous properties in lemmas:

```
8  /*@
9   lemma power_even: ...
10  lemma power_odd: ...
11  */
```

First express the lemma `power_even`, the SMT solver might be able to combine the use of this lemma and the inductive predicate to deduce `power_odd`. The Coq proof of the `power_even` lemma is provided on the GitHub repository of this book, as well as the proof of the `power_odd` in case the SMT solver does not prove it.

Finally, complete the contract and loop invariant of the `fast_power` function. For this notice that at the beginning of the loop $x^{old(n)} = p^n$, and that each iteration uses the previous properties to update r , in the sense that we have $x^{old(n)} = r * p^n$ during all the loop, until we have $n = 0$ and thus $p^n = 1$.

6.1.4.4. Permutation

Take back the definitions of the `shifted` and `unchanged` predicates from the exercise 5.4.4.4. Use the shift predicate to express the rotate predicate that expresses that some elements of an array are rotated to the left in the sense that all elements are shifted of one element to the left except the last one that is put in the first cell of the range. Use this predicate to prove the rotate function:

```

12  /*@
13   predicate rotate{L1, L2}(int* arr, integer fst, integer last) =
14     // ...
15  */
16
17  /*@
18   assigns arr[beg .. end-1] ;
19   ensures rotate{Pre, Post}(arr, beg, end) ;
20  */
21  void rotate(int* arr, size_t beg, size_t end){
22      int last = arr[end-1] ;
23      for(size_t i = end-1 ; i > beg ; --i){
24          arr[i] = arr[i-1] ;
25      }
26      arr[beg] = last ;
27  }

```

Express a new version of the notion of permutation with an inductive predicate that considers four cases:

- permutation is reflexive,
- if the left part of the range (until an index of the range) is rotated between two labels, we still have a permutation,
- if the right part of the range (from an index of the range) is rotated between two labels, we still have a permutation,
- permutation is transitive.

```

30  /*@
31   inductive permutation{L1, L2}(int* arr, integer fst, integer last){
32     case reflexive{L1}: // ...
33     case rotate_left{L1,L2}: // ...
34     case rotate_right{L1,L2}: // ...
35     case transitive{L1,L2,L3}: // ...
36   }
37  */

```

Complete the contract of `two_rotates` that successively rotates the first half of the array and then the second half and prove that it maintains the permutation.

```

43  void two_rotates(int* arr, size_t beg, size_t end){
44      rotate(arr, beg, beg+(end-beg)/2) ;
45      //@ assert permutation{Pre, Here}(arr, beg, end) ;
46      rotate(arr, beg+(end-beg)/2, end) ;
47  }

```

6.2. Axiomatic definitions

Axioms are properties we state to be true no matter the situation. It is a good way to state complex properties that will allow the proof process to be more efficient by abstracting their complexity. Of course, as any property expressed as an axiom is assumed to be true, we have to be very careful when we use them to defined properties: if we introduce a false property in our assumptions, “false” becomes “true” and we can then prove anything.

6.2.1. Syntax

Axiomatic definitions are introduced using this syntax:

```

1  /*@
2    axiomatic Name_of_the_axiomatic_definition {
3      // here we can define or declare functions and predicates
4
5      axiom axiom_name { Label0, ..., LabelN }:
6        // property ;
7
8      axiom other_axiom_name { Label0, ..., LabelM }:
9        // property ;
10
11     // ... we can put as many axioms we need
12   }
13  */

```

For example, we can write this axiomatic block:

```

1  /*@
2    axiomatic lt_plus_lt{
3      axiom always_true_lt_plus_lt:
4        \forall integer i, j; i < j ==> i+1 < j+1 ;
5    }
6  */

```

And we can see that in Frama-C, this property is actually assumed to be true¹:

```

/*@
axiomatic lt_plus_lt {
  axiom always_true_lt_plus_lt:  $\forall \mathbb{Z} i, \mathbb{Z} j; i < j \Rightarrow i + 1 < j + 1;$ 
}
*/

```

¹In section 6.4, we present an *extremely* useful axiom.

6.2.1.1. Link with lemmas

Lemmas and axioms allows to express the same kinds of properties. Namely, properties expressed about quantified variables (and possibly global variables, but it is quite rare since it is often difficult to find a global property about such variables being both true and interesting). Apart this first common point, we can also notice that when we are not considering the definition of the lemma itself, lemmas are assumed to be true by WP exactly as axioms are.

The only difference between lemmas and axioms from a proof point of view is that we must provide a proof that each lemma is true, whereas an axiom is always assumed to be true.

6.2.2. Recursive function or predicate definitions

Axiomatic definitions of recursive functions and predicates are particularly useful since they will prevent provers from unrolling the recursion when it is possible.

The idea is then not to define directly the function or the predicate but to declare it and then to define axioms that specify its behavior. If we come back to the factorial function, we can define it axiomatically as follows:

```

1  /*@
2    axiomatic Factorial{
3      logic integer factorial(integer n);
4
5      axiom factorial_0:
6        \forall integer i; i <= 0 ==> 1 == factorial(i) ;
7
8      axiom factorial_n:
9        \forall integer i; i > 0 ==> i * factorial(i-1) == factorial(i) ;
10   }
11  */

```

In this axiomatic definition, our function has no body. Its behavior is only defined by the axioms we have stated about it. Except this, nothing changes, in particular the logic function can be used in our specification just as before.

A small subtlety that we must take care of is the fact that if some axioms state properties about the content of some pointed memory cells, we have to specify considered memory blocks using the **reads** notation in the declaration. If we omit such a specification, the predicate or function will be considered to be stated about the received pointers and not about pointer memory blocks. So, if the code modifies the content of an array for which we had proven that the predicate or function gives some result, this result will not be considered to be potentially different.

For example, if we take the inductive property we stated for "zeroed" in the previous chapter, we can write it using an axiomatic definition, and it will be written like this:

```

1  /*@
2    axiomatic A_all_zeros{
3      predicate zeroed[L](int* a, integer b, integer e) reads a[b .. e-1];
4

```

6. ACSL - Logic definitions and code

```
5   axiom zeroed_empty{L}:
6     \forall int* a, integer b, e; b >= e ==> zeroed{L}(a,b,e);
7
8   axiom zeroed_range{L}:
9     \forall int* a, integer b, e; b < e ==>
10      zeroed{L}(a,b,e-1) && a[e-1] == 0 ==> zeroed{L}(a,b,e);
11 }
12 */
```

Notice the `reads[b .. e-1]` that specifies the memory location on which the predicate depends. While it is not necessary to specify what are the memory locations read in an inductive definition, we have to specify such an information for axiomatically defined properties.

6.2.3. Consistency

By adding axioms to our knowledge base, we can produce more complex proofs since some part of these proofs, expressed by axioms, do not need to be proved anymore (they are already specified to be true) shortening the proof process. However, using axiomatic definitions, **we must be extremely careful**. Indeed, even a small error could introduce false in the knowledge base, making our whole reasoning futile. Our reasoning would still be correct, but relying on false knowledge, it would only learn incorrect things.

The simplest example is the following:

```
1  /*@
2    axiomatic False{
3      axiom false_is_true: \false;
4    }
5  */
6
7  int main(){
8    // Examples of proved properties
9
10   /*@ assert \false;
11   /*@ assert \forall integer x; x > x;
12   /*@ assert \forall integer x,y,z ; x == y == z == 42;
13   return *(int*) 0;
14 }
```

And everything is proved, comprising the fact that the dereferencing of 0 is valid:

```
int main(void)
{
  int __retres;
  /*@ assert \false; */ ;
  /*@ assert \forall Z x; x > x; */ ;
  /*@ assert \forall Z x, Z y, Z z; x == y == z == 42; */ ;
  /*@ assert rte: mem_access: \valid_read((int *)0); */
  __retres = *((int *)0);
  return __retres;
}
```

Of course, this example is extreme, we would not write such an axiom. The problem is in fact that it is really easy to write an axiomatic definition that is subtly false when we express more complex properties, or adding assumptions about the global state of the system.

6. ACSL - Logic definitions and code

When we start to create axiomatic definitions, it is worth adding assertions or postconditions requiring a proof of false that we expect to fail to ensure that the definition is not inconsistent. However, it is often not enough! If the subtlety that creates the inconsistency is hard enough to find, provers could need a lot of information other than the axiomatic definition itself to be able to find and use the inconsistency, we then need to always be careful!

More specifically, specifying the values read by a function or a predicate is important for the consistency of an axiomatic definition. Indeed, as previously explained, if we do not specify what is read when a pointer is received, an update of a value in the array does not invalidate a property known about the content of the array. In such a case, the proof is performed but since the axiom does not talk about the content of the array, we do not prove anything.

For example, in the function that resets an array to 0, if we modify the axiomatic definition, removing the specification of the values that are read by the predicate (`reads a[b .. e-1]`), the proof will still be performed, but will not prove anything about the content of the arrays. For example, the following function:

```
16 /*@
17   requires length > 10 ;
18   requires \valid(array + (0 .. length-1));
19   requires zeroed(array,0,length);
20   assigns  array[0 .. length-1];
21   ensures  zeroed(array,0,length);
22 */
23 void bad_function(int* array, size_t length){
24   array[5] = 42 ;
25 }
```

is proved to be correct, while we obviously changed a value in the array and the value is not 0 anymore.

Note that unlike inductive definitions, where Coq provides us a way to control that what we write in ACSL is relatively well defined, we do not have such a mechanism for axiomatic definitions. Basically, even with Coq such a definition is translated into a list of axioms that are thus assumed.

6.2.4. Example: counting occurrences of a value

In this example, we want to prove that an algorithm actually counts the occurrences of a value in an array. We start by axiomatically defining what is the number of occurrences of a value in an array:

```
3 /*@
4   axiomatic Occurrences_Axiomatic{
5     logic integer l_occurrences_of{L}(int value, int* in, integer from, integer to)
6       reads in[from .. to-1];
7
8     axiom occurrences_empty_range{L}:
9       \forall int v, int* in, integer from, to;
10        from >= to ==> l_occurrences_of{L}(v, in, from, to) == 0;
11
12     axiom occurrences_positive_range_with_element{L}:
```

6. ACSL - Logic definitions and code

```

13     \forall int v, int* in, integer from, to;
14     (from < to && in[to-1] == v) ==>
15         l_occurrences_of(v,in,from,to) == 1+l_occurrences_of(v,in,from,to-1);
16
17     axiom occurrences_positive_range_without_element{L}:
18     \forall int v, int* in, integer from, to;
19     (from < to && in[to-1] != v) ==>
20         l_occurrences_of(v,in,from,to) == l_occurrences_of(v,in,from,to-1);
21 }
22 */

```

We have three different cases:

- the considered range of values is empty: the number of occurrences is 0,
- the considered range of values is not empty and the last element is the one we are searching for: the number of occurrences is the number of occurrences in the current range without the last element, plus 1,
- the considered range of values is not empty and the last element is not the one we are searching for: the number of occurrences is the number of occurrences in the current range without the last element.

Then, we can write the C function that computes the number of occurrences of a value in an array and prove it:

```

24 /*@
25   requires \valid_read(in+(0 .. length));
26   assigns  \nothing;
27   ensures  \result == l_occurrences_of(value, in, 0, length);
28 */
29 size_t occurrences_of(int value, int* in, size_t length){
30     size_t result = 0;
31
32     /*@
33       loop invariant 0 <= result <= i <= length;
34       loop invariant result == l_occurrences_of(value, in, 0, i);
35       loop assigns i, result;
36       loop variant length-i;
37     */
38     for(size_t i = 0; i < length; ++i)
39         result += (in[i] == value)? 1 : 0;
40
41     return result;
42 }

```

An alternative way to specify, in this code, that **result** is at most **i**, is to express a more general lemma about the number of occurrences of a value in an array, since we know that it is comprised between 0 and the size of maximum considered range of values:

```

1 /*@
2 lemma l_occurrences_of_range{L}:
3     \forall int v, int* array, integer from, to:
4         from <= to ==> 0 <= l_occurrences_of(v, a, from, to) <= to-from;
5 */

```

6. ACSL - Logic definitions and code

An automatic solver cannot discharge this lemma. It would be necessary to prove it interactively using Coq, for example. By expressing generic manually proved lemmas, we can often add useful tools to provers to manipulate more efficiently our axiomatic definitions, without directly adding new axioms that would augment the chances to introduce errors. Here, we still have to realize the proof of the lemma to get a complete proof.

6.2.5. Example: The `strlen` function

In this section, let us prove the C `strlen` function:

```
1  #include <stddef.h>
2
3  size_t strlen(char const *s){
4      size_t i = 0 ;
5      while(s[i] != '\0'){
6          ++i;
7      }
8      return i ;
9  }
```

First, we have to provide a suitable contract. Let us suppose that we have a logic function `strlen` that returns the length of a string, that is to say, what we expect of our C function:

```
1  /*@
2   logic integer strlen(char const* s) = // let's see later
3  */
```

Basically, we want to receive a valid string in input, and we want to compute a value that equals to the result of our logic function `strlen` applied to this string, of course this function does not assign anything. Defining what is a valid string is not that simple. Indeed, previously in this tutorial, we only worked with arrays, receiving in input both the array and the size of the array, however here, and as it is common in C, we suppose that the string ends with a character `'\0'`. That means that we basically need the `strlen` function to define what is a valid string. Let us first use this definition (note that we use the `\valid_read` variant of pointer validity since we do not expect the function to modify the string) and provide a contract for `strlen`:

```
10 /*@
11  predicate valid_read_string(char * s) =
12      \valid_read(s + (0 .. strlen(s))) ;
13  */
14
15 /*@
16  requires valid_read_string(s) ;
17  assigns \nothing ;
18  ensures \result == strlen(s) ;
19  */
20 size_t strlen(char const *s)
```

6. ACSL - Logic definitions and code

Defining the logic function `strlen` is a bit tricky. Indeed we want to compute the length of a string by finding the character `'\0'`, we expect to find it but in fact, we can easily imagine that we receive a string of infinite length. In this case, we would like to return an error value, but it is basically impossible to guarantee that the computation terminates, thus a logic function cannot be used to express this property.

Thus, let us define this function axiomatically. First, let us define what is read by the function, which is: any memory cell from the pointer to an infinite range of address. Then we consider two cases: the string is finite, or it is not, that leads to two axioms: `strlen` returns a positive value that corresponds to the index of the first `'\0'` character, and returns a negative value if no such value exists.

```
4  /*@
5  axiom StrLen {
6    logic integer strlen(char * s) reads s[0 .. ] ;
7
8    axiom pos_or_nul{L}:
9      \forall char* s, integer i ;
10      (0 <= i && (\forall integer j ; 0 <= j < i ==> s[j] != '\0') && s[i] == '\0') ==>
11        strlen(s) == i ;
12
13    axiom no_end{L}:
14      \forall char* s ;
15      (\forall integer i ; 0 <= i ==> s[i] != '\0') ==> strlen(s) < 0 ;
```

And now, we can be more precise for our definition of `\valid_read_string`, a valid string is a string such that it is valid from the first index to `strlen` of the string and, such that this value is greater than 0 (since an infinite string is not a valid string):

```
27 /*@
28 predicate valid_read_string(char * s) =
29   strlen(s) >= 0 && \valid_read(s + (0 .. strlen(s))) ;
30 */
```

With this definition we can now go further and provide a suitable invariant to the loop of the `strlen` function. It is quite simple: `i` ranges between 0 and `strlen(s)`, for all values met before the iteration `i`, they are not `'\0'`. This loop assigns `i` and the variant corresponds to the distance between `i` and `strlen(s)`. However, if we try to produce the proof of correctness of the function, it fails. And to get more information we can try a verification asking RTE with the verification that unsigned integers do not overflow:

6. ACSL - Logic definitions and code

```

/*@ requires valid_read_string(s);
    ensures \result == strlen(\old(s));
    assigns \nothing;
*/
size_t strlen(char const *s)
{
    size_t i = (unsigned int)0;
    /*@ loop invariant 0 ≤ i ≤ strlen(s);
        loop invariant ∀ ℤ j; 0 ≤ j < i → *(s + j) ≠ '\000';
        loop assigns i;
        loop variant strlen(s) - i;
    */
    while (1) {
        /*@ assert rte: mem_access: \valid_read(s + i); */
        if (! ((int)*(s + i) != '\000')) {
            break;
        }
        /*@ assert rte: unsigned_overflow: 0 ≤ i + (unsigned int)1; */
        /*@ assert rte: unsigned_overflow: i + (unsigned int)1 ≤ 4294967295;
        */
        i += (size_t)1;
    }
    return i;
}

```

We see that the prover fails to discharge the verification condition related to the range of `i`, and that `i` can exceed the maximum of an unsigned int. One could try to provide a limit to the value of `strlen(s)` in precondition:

```

33  requires valid_read_string(s) && strlen(s) <= UINT_MAX ;

```

However, it is not enough, and the reason is that while we have defined that the value of `strlen(s)` is defined to be the index of the first `'\0'` in the array, the converse is not true: knowing that the value of `strlen(s)` is positive is not enough to deduce that the value at the corresponding index is `'\0'`. Thus we extend the axiomatic definition with another proposition that gives us this fact (we also add one for the values that precede the `strlen(s)` index even if here, it is not necessary):

```

17  axiom index_of_strlen{L}:
18      \forall char* s ;
19      strlen(s) >= 0 ==> s[strlen(s)] == '\0' ;
20
21  axiom before_strlen{L}:
22      \forall char* s ;
23      strlen(s) >= 0 ==> (\forall integer i ; 0 <= i < strlen(s) ==> s[i] != '\0') ;

```

And this time the proof succeeds. Frama-C provides its own standard library headers, and they include an axiomatic definition for the `strlen` logic function. It can be found in the installation directory of Frama-C, under the directory `libc`, the file is named `__fc_string_axiomatic.h`. Note that this definition include more axioms in order to be able to deduce more properties about `strlen`.

6.2.6. Exercises

6.2.6.1. Occurrence counting

The following program cannot be proved with the axiomatic definition we previously defined about occurrences counting:

```

14  /*@
15   requires \valid_read(in+(0 .. length));
16   assigns  \nothing;
17   ensures  \result == l_occurrences_of(value, in, 0, length);
18  */
19  size_t occurrences_of(int value, int* in, size_t length){
20      size_t result = 0;
21
22      for(size_t i = length; i > 0 ; --i)
23          result += (in[i-1] == value) ? 1 : 0;
24
25      return result;
26  }

```

Re-express the axiomatic definition in a form that allows to prove the program.

6.2.6.2. Greatest Common Divisor

Express the logic function that allows to compute the greatest common divisor as an axiomatic definition, write the contract of the `gcd` function and prove it:

```

1  #include <limits.h>
2
3  /*@
4   axiomatic GCD {
5       // ...
6   }
7  */
8
9  /*@
10   requires a >= 0 && b >= 0 ;
11   // assigns ...
12   // ensures ...
13  */
14  int gcd(int a, int b){
15      while (b != 0){
16          int t = b;
17          b = a % b;
18          a = t;
19      }
20      return a;
21  }

```

6.2.6.3. Sum of the N first integers

Express the logic function that allows to compute the sum of the N first integers as an axiomatic definition. Write the contract of the following `sum_n` function and prove it:

6. ACSL - Logic definitions and code

```
1  #include <limits.h>
2
3  /*@ axiomatic Sum_n {
4      // ...
5  }
6  */
7
8  /*@ lemma sum_n_value: // ... */
9
10 /*@
11   requires n >= 0 ;
12   // requires ...
13   // assigns ...
14   // ensures ...
15 */
16 int sum_n(int n){
17     if(n < 1) return 0 ;
18
19     int res = 0 ;
20     /*@ loop invariant 1 <= i <= n+1 ;
21        // ...
22     */
23     for(int i = 1 ; i <= n ; i++){
24         res += i ;
25     }
26     return res ;
27 }
```

6.2.6.4. Permutation

Take back the example about selection sort (section 6.1.3). Re-express the permutation predicate using an axiomatic definition. Take care of the `reads` clause (in particular, note that the predicate relates two memory labels).

```
1  /*@
2   axiomatic Permutation {
3       // ...
4   }
5  */
6
7  /*@
8   predicate sorted(int* a, integer b, integer e) =
9       \forall integer i, j; b <= i <= j < e ==> a[i] <= a[j];
10 */
11
12 /*@
13   requires beg < end && \valid(a + (beg .. end-1));
14   assigns a[beg .. end-1];
15   ensures sorted(a, beg, end);
16   ensures permutation{Pre, Post}(a,beg,end);
17 */
18 void sort(int* a, size_t beg, size_t end){
19     /*@
20     loop invariant beg <= i <= end;
21     loop invariant sorted(a, beg, i) && permutation{Pre, Here}(a, beg, end);
22     loop invariant \forall integer j,k; beg <= j < i ==> i <= k < end ==> a[j] <= a[k];
23     loop assigns i, a[beg .. end-1];
24     loop variant end-i;
25     */
26     for(size_t i = beg ; i < end ; ++i){
27         /*@ ghost begin: ;
28         size_t imin = min_idx_in(a, i, end);
```

```

29     swap(&a[i], &a[imin]);
30     //@ assert swap_in_array{begin,Here}(a,beg,end,i,imin);
31 }
32 }

```

6.3. Ghost code

The previous techniques we have seen in this chapter are meant to make the specification more abstract. The role of ghost code is the opposite, here, we find in fact a way to enrich our specification with information expressed as concrete C code. The idea is to add variables and source code that is not part of the actual program to verify and is thus only visible for the verification tool. It is used to make explicit some logic properties that, else, would only be known implicitly.

6.3.1. Syntax

Ghost code is added using annotations that contain C code introduced using the **ghost** keyword:

```

1  /*@
2     ghost
3     // C code
4  */

```

The only rules we have to respect in such a code, is that we cannot write a memory block that is not itself defined in ghost code, and that the code must close any block it would open. Apart from this, any computation can be inserted provided that it modifies *only* ghost variables.

Here are some examples of ghost code:

```

1  //@ ghost int ghost_glob_var = 0;
2
3  void foo(int a){
4     //@ ghost int ghost_loc_var = a;
5
6     //@ ghost Ghost_label: ;
7     a = 28 ;
8
9     //@ ghost if(a < 0){ ghost_loc_var = 0; }
10
11    //@ assert ghost_loc_var == \at(a, Ghost_label) == \at(a, Pre);
12 }

```

While for this chapter, it will not be necessary, as we will see later, it is sometimes useful to write some contracts or assertions in ghost code. As it must be specified in code that is already in C comments, we can use a specific syntax to provide ghost contracts or assertions. We open ghost annotations with the syntax `/@` and close them with `@/`.

For example :

```

1 void foo(unsigned n){
2   /*@ ghost
3     unsigned i ;
4
5     /*
6       loop invariant 0 <= i <= n ;
7       loop assigns i ;
8       loop variant n - i ;
9     */
10    for(i = 0 ; i < n ; ++i){
11
12    }
13    /* assert i == n ; */
14  */
15 }

```

6.3.2. Ghost code validity

We must again be careful using ghost code. Indeed, the tool does not perform any verification to ensure that we do not write in the memory of the program by mistake. This problem being, in fact, an undecidable problem, this analysis would require a proof by itself. For example, this code is allowed as input of Frama-C even if it explicitly modifies the state of the program we want to verify:

```

1 int a;
2
3 void foo(){
4   /*@ ghost a = 42;
5 }

```

Thus, it is easy to modify the behavior of the code we want to verify by adding ghost code and in this case, if we prove a property about the program, we would prove it about the modified program, not about the actual one. One must thus take care of modifying only ghost memory locations when writing ghost code.

Furthermore, as we do not have restriction on the kind of C code we can write, we can write loops. As we said in section 4.2.3, there are two kinds of correctness: partial and total correctness, the second one allowing to prove that a program terminates. In the case of ghost code, it is absolutely necessary to prove total correctness as an infinite loop in the ghost code would allow us to prove anything about the program.

For example, the following program is proved because of the non terminating ghost code:

```

1 /*@ ensures \false ; */
2 void foo(void){
3   /*@ ghost
4     while(1){}
5   */
6 }

```

6.3.3. Make a logical state explicit

The goal of ghost code is to make explicit some information that would else be implicit. For example, in the verification of the sort algorithm, we used it to create a label in the program that is not visible by the compiler but that we can use in the verification. The fact that the value was swapped between the two labels was implicitly provided by the contract of the function, adding this ghost label allows us to write an explicit assertion of this fact.

Let us take a more complex example where we more clearly create explicit knowledge about the program. Here, we want to prove that the following function returns the value of the maximal sum of subarrays of a given array. A subarray of an array **a** is a contiguous subset of values of **a**. For example, for an array `{ 0 , 3 , -1 , 4 }`, some subarrays can be `{}`, `{ 0 }`, `{ 3 , -1 }`, `{ 0 , 3 , -1 , 4 }`, ... Note that as we allow empty arrays for subarrays, the sum is at least 0. In the previous array, the subarray with the maximal sum is `{ 3 , -1 , 4 }`, the function would then return 6.

```

3  int max_subarray(int *a, size_t len) {
4      int max = 0;
5      int cur = 0;
6      for(size_t i = 0; i < len; i++) {
7          cur += a[i];
8          if (cur < 0) cur = 0;
9          if (cur > max) max = cur;
10     }
11     return max;
12 }
```

In order to specify the previous function, we need an axiomatic definition about sum. This is not too complex, the careful reader can express the needed axioms as an exercise:

```

3  /*@
4      axiomatic Sum_array{
5          logic integer sum(int* array, integer begin, integer end) reads array[begin .. (end-1)];
```

Some correct axioms are available at: 6.4

The specification of the function is the following:

```

14 /*@
15     requires \valid(a+(0..len-1));
16     assigns \nothing;
17     ensures \forall integer l, h; 0 <= l <= h <= len ==> sum(a,l,h) <= \result;
18     ensures \exists integer l, h; 0 <= l <= h <= len && sum(a,l,h) == \result;
19 */
```

For any bounds, the value returned by the function must be greater or equal to the sum of the elements between these bounds, and there must exist some bounds such that the returned value is exactly the sum of the elements between these bounds. About this specification, when we want to add the loop invariant, we realize that we miss some information. We want to express

6. ACSL - Logic definitions and code

what are the values `max` and `cur` and what are the relations between them, but we cannot do it!

Basically, in order to prove our postcondition, we need to know that there exists some bounds `low` and `high` such that the computed sum corresponds to these bounds. However, in our code, we do not have anything that expresses it from a logic point of view, and we cannot *a priori* make the link between this logic formalization. We then use ghost code to record these bounds and express the loop invariant.

We first need two variables to record the bounds of the maximum sum range, let us call them `low` and `high`. Every time we find a range where the sum is greater than the current one, we update our ghost variables. These bounds then correspond to the sum currently stored by `max`. That induces that we need other bounds: the ones that corresponds to the sum stored by the variable `cur` from which we build the bounds corresponding to `max`. For these bounds, we only add a single ghost variable: the current low bound `cur_low`, the high bound being the variable `i` of the loop.

```

14  /*@
15   requires \valid(a+(0..len-1));
16   assigns \nothing;
17   ensures \forall integer l, h; 0 <= l <= h <= len ==> sum(a,l,h) <= \result;
18   ensures \exists integer l, h; 0 <= l <= h <= len && sum(a,l,h) == \result;
19  */
20  int max_subarray(int *a, size_t len) {
21      int max = 0;
22      int cur = 0;
23      /*@ ghost size_t cur_low = 0;
24       /*@ ghost size_t low = 0;
25       /*@ ghost size_t high = 0;
26
27      /*@
28       loop invariant BOUNDS: low <= high <= i <= len && cur_low <= i;
29
30       loop invariant REL :   cur == sum(a,cur_low,i) <= max == sum(a,low,high);
31       loop invariant POST:   \forall integer l; 0 <= l <= i ==> sum(a,l,i) <= cur;
32       loop invariant POST:   \forall integer l, h; 0 <= l <= h <= i ==> sum(a,l,h) <= max;
33
34       loop assigns i, cur, max, cur_low, low, high;
35       loop variant len - i;
36  */
37      for(size_t i = 0; i < len; i++) {
38          cur += a[i];
39          if (cur < 0) {
40              cur = 0;
41              /*@ ghost cur_low = i+1; */
42          }
43          if (cur > max) {
44              max = cur;
45              /*@ ghost low = cur_low; */
46              /*@ ghost high = i+1; */
47          }
48      }
49      return max;
50  }

```

The invariant `BOUNDS` expresses how the different bounds are ordered during the computation. The invariant `REL` expresses what the variables `cur` and `max` mean depending on the bounds. Finally, the invariant `POST` allows us to create a link between the invariant and the postcondition of the function.

6. ACSL - Logic definitions and code

The reader can verify that this function is indeed correctly proved without RTE verification. If we add RTE verification, the overflow on the variable `cur`, that is the sum, seems to be possible (and it is indeed the case).

Here, we do not try to fix this because it is not the topic of this example. The way we can prove the absence of RTE here strongly depends on the context where we use this function. A possibility is to strongly restrict the contract, forcing some properties about values and the size of the array. For example, we could limit the maximal size of the array and bound each value of the different cells. Another possibility would be to add an error value in case of overflow (-1 for example), and to specify that when an overflow is produced, this value is returned.

6.3.4. Exercises

6.3.4.1. Two times

This program computes `2 * x` using a loop. Use a ghost variable `i` to express as an invariant that the value of `r` is `i * 2` and complete the proof.

```
1  /*@
2   requires x >= 0 ;
3   assigns \nothing ;
4   ensures \result == 2 * x ;
5  */
6  int times_2(int x){
7      int r = 0 ;
8      /*@
9       loop invariant 0 <= x ;
10      loop invariant r == // ...
11      loop invariant // ...
12     */
13     while(x > 0){
14         r += 2 ;
15         x -- ;
16     }
17     return r;
18 }
```

6.3.4.2. Playing with arrays

In this function, we receive an array, and we have a loop where we do nothing except that we have indicated that it assigns the content of the array. However, we would like to prove in postcondition that the array has not been modified.

```
1  /*@
2   requires \valid(a + (0 .. 9)) ;
3   assigns a[0 .. 9] ;
4   ensures \forall integer j ; 0 <= j < 10 ==> a[j] == \old(a[j]) ;
5  */
6  void foo(int a[10]){
7      //@ ghost int g[10] ;
8      /*@ ghost
9       ...
10     */
```

```

11
12  /*@
13   loop invariant 0 <= i <= 10 ;
14   loop invariant // ...
15   loop assigns i, a[0 .. 9] ;
16   loop variant 10 - i ;
17  */
18  for(int i = 0; i < 10; i++);
19 }

```

Without modifying the `assigns` clause of the loop and without using the keyword `\at`, prove that the function does not modify the array. For this, complete the ghost code and the loop invariant, by assuring that the array `g` represents the old value of `a`.

6.4. Hidden content

6.4.1. Coq Proof of the `no_changes` lemma

```

1  Inductive P_zeroed : farray addr Z -> addr -> Z -> Z -> Prop :=
2    | Q_zeroed_empty: forall (i_1 i : Z), forall (t : farray addr Z),
3      forall (a : addr), ((i <= i_1)%Z) -> ((P_zeroed t a i_1%Z i)%Z)
4    | Q_zeroed_range: forall (i_1 i : Z), forall (t : farray addr Z),
5      forall (a : addr), let x := (i%Z - 1%Z)%Z in
6        (((t.[ shift_sint32 a x ]) = 0)%Z) -> ((i_1 < i)%Z) ->
7          ((P_zeroed t a i_1%Z x)) -> ((P_zeroed t a i_1%Z i)%Z).
8
9  Definition P_same_elems (Mint_0 : farray addr Z) (Mint_1 : farray addr Z)
10    (a : addr) (b : Z) (e : Z) : Prop :=
11    forall (i : Z), let a_1 := (shift_sint32 a i)%Z in ((b <= i)%Z) ->
12      ((i < e)%Z) -> (((Mint_1.[ a_1 ]) = (Mint_0.[ a_1 ]))%Z).
13
14  (* The property to prove *)
15  Goal
16    forall (i_1 i : Z),
17    forall (t_1 t : farray addr Z),
18    forall (a : addr),
19      ((P_zeroed t a i_1%Z i)%Z) ->
20      ((P_same_elems t_1 t a i_1%Z i)%Z) ->
21      ((P_zeroed t_1 a i_1%Z i)%Z).
22
23  Proof.
24    (* We introduce our variable and the main hypothesis *)
25    intros b e Mi Mi' arr H.
26    (* We reason by induction on our first (inductive) hypothesis *)
27    induction H ; intros Same.
28    (* Base case, by using the first case of the inductive predicate *)
29    + constructor 1.
30      (* The only premise to prove is a trivial relation between the bounds *)
31      omega.
32    + unfold x in * ; clear x.
33      (* Induction case, by using the second case of the inductive predicate *)
34      constructor 2.
35      (* We have three premises *)
36      - (* First: the first cell in new memory must be zero, we replace 0 with
37        the cell in old memory *)
38        rewrite <- H ; symmetry.
39        (* And show that the cells are the same *)
40        apply Same ; omega.
41      - (* Second, we have to prove a trivial relation about the bounds *)
42        omega.

```

6. ACSL - Logic definitions and code

```
43 - (* Third we use our induction hypothesis to show that the property
44    holds on the first part of the array *)
45   apply IHP_zeroed.
46   intros i' ; intros.
47   apply Same ; omega.
48 Qed.
```

6.4.2. Specified sort functions

```
3  /*@
4   requires \valid_read(a + (beg .. end-1));
5   requires beg < end;
6
7   assigns \nothing;
8
9   ensures \forall integer i; beg <= i < end ==> a[\result] <= a[i];
10  ensures beg <= \result < end;
11 */
12 size_t min_idx_in(int* a, size_t beg, size_t end){
13     size_t min_i = beg;
14
15     /*@
16      loop invariant beg <= min_i < i <= end;
17      loop invariant \forall integer j; beg <= j < i ==> a[min_i] <= a[j];
18      loop assigns min_i, i;
19      loop variant end-i;
20     */
21     for(size_t i = beg+1; i < end; ++i){
22         if(a[i] < a[min_i]) min_i = i;
23     }
24     return min_i;
25 }
26
27 /*@
28  requires \valid(p) && \valid(q);
29  assigns *p, *q;
30  ensures *p == \old(*q) && *q == \old(*p);
31 */
32 void swap(int* p, int* q){
33     int tmp = *p; *p = *q; *q = tmp;
34 }
```

6.4.3. An important axiom

Currently, our automatic solvers are not powerful enough to compute *the Answer to the Ultimate Question of Life, the Universe, and Everything*. We can help them by stating it as an axiom! Now, we just have to understand the question to determine in which case this result can be useful ...

```
1  /*@
2   axiomatic Ax_answer_to_the_ultimate_question_of_life_the_universe_and_everything {
3       logic integer the_ultimate_question_of_life_the_universe_and_everything{L} ;
4
5       axiom answer{L}:
6           the_ultimate_question_of_life_the_universe_and_everything{L} = 42;
7   }
```

6. ACSL - Logic definitions and code

```
8  */
```

6.4.4. Sum axioms

```
3  /*@
4   axiomat Sum_array{
5     logic integer sum(int* array, integer begin, integer end) reads array[begin .. (end-1)];
6
7     axiom empty:
8       \forall int* a, integer b, e; b >= e ==> sum(a,b,e) == 0;
9     axiom range:
10      \forall int* a, integer b, e; b < e ==> sum(a,b,e) == sum(a,b,e-1)+a[e-1];
11   }
12  */
```

6. ACSL - Logic definitions and code

In this part, we have covered some advanced constructions of the ACSL language that allow to express and prove more complex properties about programs.

Badly used, these features can make our analyses incorrect, we then need to be careful manipulating them and not hesitate to check them again and again, or eventually express properties to verify about them to assure that we are not introducing an incoherence in our program or our assumptions.

7. Proof methodologies

Now that we have presented most of the important features of ACSL for program proof, let us have a more general view of program proof with Frama-C and WP. We will present some approaches that can be used depending on the target of verification, the expected level of confidence and the kind of formalization (and features of ACSL) we use.

7.1. Absence of runtime errors: Minimal contracts

We have seen that program proof allows to verify two main aspects of program correctness: first that programs do not contain runtime errors, second that programs correctly implement their specification. However, it is sometimes hard to guarantee the later, and the former is already an interesting step for program correctness.

Indeed, runtime errors often cause the presence of so-called “undefined behaviors” in C programs. These undefined behaviors are often vectors of security breaches and thus, guaranteeing their absence already protects us of a lot of attack vectors. The absence of runtime errors can be verified thanks to an approach called minimal contracts.

7.1.1. Principle

The minimal contracts approach is guided by the RTE plugin. Basically, the idea is to generate the assertions related to absence of runtime errors for all the functions of a module or a project and to write the minimal set of (correct) contracts that are sufficient to prove that the runtime errors cannot happen. Most of the time, we need far less lines of specification than what is usually required to prove the functional correctness of the program.

Let us take a simple example with the absolute value function.

```
1 int abs(int x){
2   return (x < 0) ? -x : x ;
3 }
```

Here, we can generate the assertions required to prove the absence of runtime errors, which generate this program:

```
1 /* Generated by Frama-C */
2 int abs(int x)
3 {
4   int tmp;
```

7. Proof methodologies

```
5  if (x < 0)
6      /*@ assert rte: signed_overflow: -2147483647 ≤ x; */
7      tmp = - x;
8  else tmp = x;
9  return tmp;
10 }
```

Thus, we only need to specify as a requirement that the value of `x` must be greater than `INT_MIN` :

```
1  /*@
2      requires x > INT_MIN ;
3  */
4  int abs(int x){
5      return (x < 0) ? -x : x ;
6  }
```

This condition is enough to prove that no runtime error can happen in the function.

As we will see later, the function is however generally used in a particular context. So this contract will likely not be enough. For example, we often have global variables in our program and here we do not specify what is assigned by the function. Most of the time the “assigns” clause cannot be ignored (which is expected in a language where everything is mutable by default). Moreover, if one take the absolute value of an integer, it is surely because they need a positive value. In reality, the minimal contract of the absolute value function is more likely the following:

```
1  /*@
2      requires x > INT_MIN ;
3      assigns \nothing ;
4      ensures \result >= 0 ;
5  */
6  int abs(int x){
7      return (x < 0) ? -x : x ;
8  }
```

But this addition should only be guided by the verification of the context(s) where the function is used once we have proved the absence of runtime errors in the function itself.

7.1.2. Example: the search function

Now that we have the principle in mind, let us work on more complex examples, in particular with an example that involves a loop.

```
1  #include <stddef.h>
2
3  int* search(int* array, size_t length, int element){
4      for(size_t i = 0; i < length; i++)
5          if(array[i] == element) return &array[i];
6      return NULL;
7  }
```

7. Proof methodologies

Once we have generated the assertions related to runtime errors, we have the following program:

```
1  /* Generated by Frama-C */
2  #include "stddef.h"
3  int *search(int *array, size_t length, int element)
4  {
5      int *__retres;
6      {
7          size_t i = (unsigned int)0;
8          while (i < length) {
9              /*@ assert rte: mem_access: \valid_read(array + i); */
10             if (*(array + i) == element) {
11                 __retres = array + i;
12                 goto return_label;
13             }
14             i += (size_t)1;
15         }
16     }
17     __retres = (int *)0;
18     return_label: return __retres;
19 }
```

We have to prove that any cell visited by the program can be read, thus we need to express as a precondition that the array is `\valid_read` on the corresponding range of indices. However, this is not enough to complete the proof since we have a loop in this program, so we have to provide a suitable invariant. We also probably want to prove that the loop terminates.

Thus, we get the following minimally specified function:

```
1  #include <stddef.h>
2
3  /*@
4   requires \valid_read(array + (0 .. length-1)) ;
5  */
6  int* search(int* array, size_t length, int element){
7      /*@
8       loop invariant 0 <= i <= length ;
9       loop assigns i ;
10      loop variant length - i ;
11      */
12      for(size_t i = 0; i < length; i++)
13          if(array[i] == element) return &array[i];
14      return NULL;
15 }
```

This contract can be compared with the contract provided for the search function in section 4.3.1, and we can see that it is much more simple.

Now let us imagine that the function is used in the following program:

```
1  void foo(int* array, size_t length){
2      int* p = search(array, length, 0) ;
3      if(p){
4          *p += 1 ;
5      }
6  }
```

7. Proof methodologies

We again have to provide a suitable contract for the function, again by having a look at the assertion that RTE asks us to verify:

```
26 void foo(int *array, size_t length)
27 {
28     int *p = search(array, length, 0);
29     if (p)
30         /*@ assert rte: mem_access: \valid(p); */
31         /*@ assert rte: mem_access: \valid_read(p); */
32         /*@ assert rte: signed_overflow: *p + 1 ≤ 2147483647; */
33         (*p) ++;
34     return;
35 }
```

Thus we have to verify that:

- the pointer we received from `search` is valid,
- `*p + 1` does not overflow,
- we respect the contract of the `search` function.

In addition to the contract of `foo`, we have to provide some more information in the contract of `search`. Indeed, we will not be able to prove that the pointer is valid when it is not null if the function does not guarantee that the pointer is in the range of our array in this case. Furthermore, we will not be able to prove that `*p` is less than `INT_MAX` if the function can modify it.

This leads us to this complete annotated program:

```
1  #include <stddef.h>
2  #include <limits.h>
3
4  /*@
5   requires \valid_read(array + (0 .. length-1)) ;
6   assigns \nothing ;
7   ensures \result == NULL ||
8           (\exists integer i ; 0 ≤ i < length && array+i == \result) ;
9  */
10 int* search(int* array, size_t length, int element){
11     /*@
12      loop invariant 0 ≤ i ≤ length ;
13      loop assigns i ;
14      loop variant length - i ;
15     */
16     for(size_t i = 0; i < length; i++){
17         if(array[i] == element) return &array[i];
18     }
19     return NULL;
20 }
21
22 /*@
23 requires \forall integer i ; 0 ≤ i < length ==> array[i] < INT_MAX ;
24 requires \valid(array + (0 .. length-1)) ;
25 */
26 void foo(int *array, size_t length){
27     int *p = search(array, length, 0);
28     if (p){
29         *p += 1 ;
30     }
31 }
```

7.1.3. Advantages and limitations

The evident advantage of this approach is the fact that it can guarantee the absence of runtime errors in any function of a module or a program in (relative) isolation of the other functions. Furthermore, this absence of runtime errors is guaranteed for any use of the function as long as the precondition is verified when it is called. That allows to gain some confidence into a system with a relatively low-cost approach.

However, as we have seen, when we use a function it can change the knowledge we need about its behavior, requiring to make the contract richer and richer. Thus we can progressively reach a state where we basically proved the functional correctness of the function.

Furthermore, even proving the absence of runtime errors is sometimes not trivial as we have for example seen with functions like the factorial or the sum of N integers, that require to give quite a lot of information to SMT solvers, in order to prove that we cannot meet a integer overflow.

Finally, sometimes the minimal contracts of a function or a module basically is the full functional specification, and thus formally verifying the absence of runtime errors requires a full functional verification. This is commonly the case when we have to deal with complex data structures where the properties that are required for the absence of runtime errors depend on the functional behavior of the function, maintaining some non-trivial invariant about the data structure.

7.1.4. Exercises

7.1.4.1. Some simple examples

Prove the absence of runtime errors in the following program using a minimal contracts approach:

```
1 void max_ptr(int* a, int* b){
2     if(*a < *b){
3         int tmp = *b ;
4         *b = *a ;
5         *a = tmp ;
6     }
7 }
8
9 void min_ptr(int* a, int* b){
10     max_ptr(b, a);
11 }
12
13 void order_3_inc_min(int* a, int* b, int* c){
14     min_ptr(a, b) ;
15     min_ptr(a, c) ;
16     min_ptr(b, c) ;
17 }
18
19 void incr_a_by_b(int* a, int const* b){
20     *a += *b;
21 }
```

7. Proof methodologies

7.1.4.2. Reverse

Prove the absence of runtime errors for the following `reverse` function and its dependency using a minimal contracts approach. Note that the `swap` function should also be specified with minimal contracts only. Do not forget to add the option `-warn-unsigned-overflow`.

```
1  #include <stddef.h>
2
3  void swap(int* a, int* b){
4      int tmp = *a;
5      *a = *b;
6      *b = tmp;
7  }
8
9  void reverse(int* array, size_t len){
10     for(size_t i = 0 ; i < len/2 ; ++i){
11         swap(array+i, array+len-i-1) ;
12     }
13 }
```

7.1.4.3. Binary search

Prove the absence of runtime errors for the following `bsearch` function using a minimal contracts approach. Do not forget to add the option `-warn-unsigned-overflow`.

```
1  #include <stddef.h>
2
3  size_t bsearch(int* arr, size_t len, int value){
4      if(len == 0) return -1 ;
5
6      size_t low = 0 ;
7      size_t up = len ;
8
9      while(low < up){
10         size_t mid = low + (up - low)/2 ;
11         if (arr[mid] > value) up = mid ;
12         else if(arr[mid] < value) low = mid+1 ;
13         else return mid ;
14     }
15     return -1 ;
16 }
```

7.1.4.4. Sort

Prove the absence of runtime errors for the following `sort` function and its dependencies using a minimal contracts approach. Note that these dependencies should also be specified with minimal contracts only. Do not forget to add the option `-warn-unsigned-overflow`.

```
1  #include <stddef.h>
2
3  size_t min_idx_in(int* a, size_t beg, size_t end){
4      size_t min_i = beg;
5      for(size_t i = beg+1; i < end; ++i){
6          if(a[i] < a[min_i]) min_i = i;
7      }
8      return min_i;
9  }
10
11 void swap(int* p, int* q){
12     int tmp = *p; *p = *q; *q = tmp;
13 }
14
15 void sort(int* a, size_t beg, size_t end){
16     for(size_t i = beg ; i < end ; ++i){
17         size_t imin = min_idx_in(a, i, end);
18         swap(&a[i], &a[imin]);
19     }
20 }
```

7.2. Guiding assertions and triggering of lemmas

There are different levels of automation for the verification of programs. From tools that are completely automatic, like for example abstract interpreters that do not require any help from the human (or at least not so much), to interactive tools, like proof assistants where we write the proof mostly by hand and the tools is just there to check that we do it right.

Tools like WP (any many others like Why3, Spark, ...) tend to maximize automation. However, the more the properties we want to prove are complex, the harder it will be to get automatically the proof. Thus, we often need to help the tools in order to achieve the verification. This is done by providing more annotations to help the verification condition generation process. Adding a loop invariant is for example a way to be able to produce an inductive reasoning about a loop while automatic provers are generally bad at this kind of task.

This kind of technique has been called “auto-active” verification. This word is the contraction of “automatic” and “interactive”. It is automatic in the sense that most of the proof is performed by automatic tools, but it is also somewhat interactive since as users, we manually provide information to the tools.

In this section, we will see in more details how we can use assertions to guide the proofs. By adding assertions, we create some base of knowledge (properties that are known to be true) that are collected by the verification condition generator during the WP computation process and that are then given to the automatic solvers that consequently have more information and thus can potentially prove more complex properties.

7.2.1. Proof context

In order to understand what is exactly the benefit of adding assertions in annotations of a program, let us first have a closer look at the verification condition generated by WP from the

7. Proof methodologies

annotated source code and how assertions are taken into account. For this, we consider the following predicate (one can recognize Pythagoras' theorem):

```
1  /*@
2   predicate rectangle{L}(integer c1, integer c2, integer h) =
3     c1 * c1 + c2 * c2 == h * h ;
4  */
```

Let us first consider this example:

```
6  /*@
7   requires \separated(x, y, z);
8   requires 3 <= *x <= 5 ;
9   requires 4 <= *y <= 5 ;
10  requires *z <= 5 ;
11  requires *x+2 == *y+1 == *z ;
12 */
13 void example_1(int* x, int* y, int* z){
14   //@ assert rectangle(*x, *y, *z);
15   //@ assert rectangle(2* (*x), 2* (*y), 2* (*z));
16 }
```

Here, we have specified a precondition that is complex enough so that WP cannot directly guess the values in input of the function. In fact, the values are exactly: `*x == 3`, `*y == 4` and `*z == 5`. Now, if we have a look at the verification condition generated for our first assertion, we can see this (be sure to select the view “Full Context” or “Raw Obligation” - they are not exactly the same but quite similar, the former is just slightly pretty printed):

The screenshot shows a software verification tool interface. At the top, a C function `example_1` is displayed with two assertions: `//@ assert rectangle(*x, *y, *z);` and `//@ assert rectangle(2 * *x, 2 * *y, 2 * *z);`. Below the code, a toolbar contains tabs for `Information`, `Messages (0)`, `Console`, `Properties`, `Values`, `Red Alarms`, and `WP Goals`. The `WP Goals` tab is active, showing a `Raw Obligation` view. The goal text is as follows:

```
Goal Assertion:
Let x_1 = Mint_0[y].
Let x_2 = Mint_0[x].
Let x_3 = Mint_0[z].
Assume {
  Type: is_sint32(x_2) /\ is_sint32(x_1) /\ is_sint32(x_3).
  (* Heap *)
  Have: (region(x.base) <= 0) /\ (region(y.base) <= 0) /\
        (region(z.base) <= 0).
  (* Pre-condition *)
  Have: (y != x) /\ (z != x) /\ (z != y) /\ (x_1 = (1 + x_2)) /\
        (x_3 = (1 + x_1)) /\ (3 <= x_2) /\ (4 <= x_1) /\ (x_2 <= 5) /\
        (x_1 <= 5) /\ (x_3 <= 5).
}
Prove: P_rectangle(x_2, x_1, x_3).
```

That is to say the different constraints we have stated as a precondition to the function (note that the values are not exactly the same, and some more properties have been specified). What is more interesting is having a look at the verification condition generated for the second

7. Proof methodologies

assertion (note that we have edited all the remaining screenshots in this section in order to focus on what is important, the other properties can just be ignored in our case):

```

void example_1(int *x, int *y, int *z)
{
  /*@ assert rectangle(*x, *y, *z); */ ;
  /*@ assert rectangle(2 * *x, 2 * *y, 2 * *z); */ ;
  return;
}

```

Information Messages (1) Console Properties Values Red Alarms WP Goals

Full Context Binary Proved Goal

Goal Assertion:

```

Let x_1 = « *x »@L1.
Let x_2 = « *y »@L1.
Let x_3 = « *z »@L1.
Assume {
  Stmt { L1: }
  (* Pre-condition *)
  Have: (y@L1 != x@L1) /\ (z@L1 != x@L1) /\ (z@L1 != y@L1) /\
        (x_2 = (1 + x_1)) /\ (x_3 = (1 + x_2)) /\ (3 <= x_1) /\
        (4 <= x_2) /\ (x_1 <= 5) /\ (x_2 <= 5) /\ (x_3 <= 5).
  (* Assertion *)
  Have: P_rectangle(x_1, x_2, x_3).
}
Prove: P_rectangle(2 * x_1, 2 * x_2, 2 * x_3).

```

Here, we can see that for the proof of the second assertion, WP has collected and added the first assertion to the assumptions. Thus WP considers that SMT solvers can assume this assertion. That means that they can rely on it, but also that it should be proved to be sure that the current verification condition is verified.

Note that WP only collects what we can find on the different paths that allows to reach the assertion. For example, if we modify the code such that the path to the assertion jumps over the first assertion, it does not appear in the verification condition.

```

26 void example_1_p(int* x, int* y, int* z){
27     goto next;
28     /*@ assert rectangle(*x, *y, *z);
29     next: ;
30     /*@ assert rectangle(2* (*x), 2* (*y), 2* (*z));
31 }

```

7. Proof methodologies

```
void example_1_p(int *x, int *y, int *z)
{
    goto next;
    /*@ assert rectangle(*x, *y, *z); */ ;
next: ;
    /*@ assert rectangle(2 * *x, 2 * *y, 2 * *z); */ ;
    return;
}
```

Information Messages (1) Console Properties Values Red Alarms WP Goals

Full Context Binary Proved Goal

Goal Assertion:

```
Let x_1 = « *y »@L1.
Let x_2 = « *x »@L1.
Let x_3 = « *z »@L1.
Assume {
    Stmt { L1: }
    (* Pre-condition *)
    Have: (y@L1 != x@L1) /\ (z@L1 != x@L1) /\ (z@L1 != y@L1) /\
          (x_1 = (1 + x_2)) /\ (x_3 = (1 + x_1)) /\ (3 <= x_2) /\
          (4 <= x_1) /\ (x_2 <= 5) /\ (x_1 <= 5) /\ (x_3 <= 5).
}
Prove: P_rectangle(2 * x_2, 2 * x_1, 2 * x_3).
```

Now let us modify the example a little bit in order to illustrate how assertions can change the way to prove a program. For example, let us first modify the different memory locations (doubling each value) and check that the resulting triangle is a right triangle.

```
34 /*@
35   requires \separated(x, y , z);
36   requires 3 <= *x <= 5 ;
37   requires 4 <= *y <= 5 ;
38   requires *z <= 5 ;
39   requires *x+2 == *y+1 == *z ;
40 */
41 void example_2(int* x, int* y, int* z){
42     *x += 3 ;
43     *y += 4 ;
44     *z += 5 ;
45
46     /*@ assert rectangle(*x, *y, *z);
47 }
```

7. Proof methodologies

```
void example_2(int *x, int *y, int *z)
{
  *x += 3;
  *y += 4;
  *z += 5;
  /*@ assert rectangle(*x, *y, *z); */ ;
  return;
}
```

Information Messages (1) Console Properties Values Red Alarms WP Goals

Full Context Binary Proved Goal

Goal Assertion:

```
Let x_1 = « *y »@L1.
Let x_2 = « *x »@L1.
Let x_3 = « *z »@L1.
Let x_4 = « *z@L1 »@L4.
Let x_5 = 5 + x_4.
Let x_6 = « *x@L1 »@L5.
Let x_7 = « *y@L1 »@L5.
Assume {
  Stmt { L1: }
  (* Pre-condition *)
  Have: (y@L1 != x@L1) /\ (z@L1 != x@L1) /\ (z@L1 != y@L1) /\
        (x_1 = (1 + x_2)) /\ (x_3 = (1 + x_1)) /\ (3 <= x_2) /\
        (4 <= x_1) /\ (x_2 <= 5) /\ (x_1 <= 5) /\ (x_3 <= 5).
  Stmt { *x@L1 = 3 + x_2; }
  Stmt { L3: *y@L1 = 4 + *y@L1; }
  Stmt { L4: *z@L1 = x_5; }
  Stmt { L5: }
}
Prove: P_rectangle(x_6, x_7, x_5).
```

Here, the solver will likely unfold the predicate and directly check that the property is true, indeed from the information we have in this verification condition, we do not really have any other knowledge that would allow us to produce this proof in another way. Now, let us bring new information in annotation:

```
49 /*@
50   requires \separated(x, y, z);
51   requires 3 <= *x <= 5 ;
52   requires 4 <= *y <= 5 ;
53   requires *z <= 5 ;
54   requires *x+2 == *y+1 == *z ;
55 */
56 void example_3(int* x, int* y, int* z){
57   /*@ assert rectangle(2* (*x), 2* (*y), 2* (*z));
58   /*@ ghost L: ;
59
60   *x += 3 ;
61   *y += 4 ;
62   *z += 5 ;
63
64   /*@ assert *x == \at(2* (*x), L) ;
65   /*@ assert *y == \at(2* (*y), L) ;
66   /*@ assert *z == \at(2* (*z), L);
67   /*@ assert rectangle(*x, *y, *z);
68 }
```

We first prove that if we multiply by 2 each of the values, the predicate is true for the new values. The solver in fact basically solves the same problem at first, but this is not what we want to point out now. Then, we re-express the values we have modified in another way: we show that all of them have been multiplied by 2. Now, we can have a new look at the generated verification condition for the last assertion:

7. Proof methodologies

```

void example_3(int *x, int *y, int *z)
{
  /*@ assert rectangle(2 * *x, 2 * *y, 2 * *z); */ ;
  L: /*@ ghost ; */
  *x += 3;
  *y += 4;
  *z += 5;
  /*@ assert *x == \at(2 * *x,L); */ ;
  /*@ assert *y == \at(2 * *y,L); */ ;
  /*@ assert *z == \at(2 * *z,L); */ ;
  /*@ assert rectangle(*x, *y, *z); */ ;
  return;
}

```

Information Messages (1) Console Properties Values Red Alarms WP Goals

Full Context

Proved Goal

Goal Assertion:

```

Let x_1 = « *x »@L1.      Let x_6 = « *y »@L6.
Let x_2 = « *z »@L5.      Let x_7 = « *y »@L1.
Let x_3 = 5 + x_2.        Let x_8 = 2 * x_7.
Let x_4 = « *z »@L1.      Let x_9 = « *x »@L6.
Let x_5 = 2 * x_4.        Let x_10 = 2 * x_1.
Assume {
  Stmt { L1: }
  (* Pre-condition *)
  Have: (y@L1 != x@L1) /\ (z@L1 != x@L1) /\ (z@L1 != y@L1) /\
        (x_7 = (1 + x_1)) /\ (x_4 = (1 + x_7)) /\ (3 <= x_1) /\
        (4 <= x_7) /\ (x_1 <= 5) /\ (x_7 <= 5) /\ (x_4 <= 5).
  (* Assertion *)
  Have: P_rectangle(x_10, x_8, x_5).
  Stmt { *x@L1 = 3 + x_1; }
  Stmt { L4: *y@L1 = 4 + *y@L1; }
  Stmt { L5: *z@L1 = x_3; }
  Stmt { L6: }
  (* Assertion *)
  Have: x_9 = x_10.
  (* Assertion *)
  Have: x_6 = x_8.
  (* Assertion *)
  Have: x_3 = x_5.
}
Prove: P_rectangle(x_9, x_6, x_3).

```

While we have to prove exactly the same property as before (with a bit of renaming), we can see that we have here another way to prove it. Indeed, by just combining this set of properties:

```

1  (* Assertion *)
2  Have: P_rectangle(x_10, x_8, x_5).
3  (* Assertion *)
4  Have: x_9 = x_10.
5  (* Assertion *)
6  Have: x_6 = x_8.
7  (* Assertion *)
8  Have: x_3 = x_5.

```

It is easy to deduce:

```

1  Prove: P_rectangle(x_9, x_6, x_3).

```

By just replacing the values `x_9`, `x_6` and `x_3`. So the solver could use this to avoid unfolding the predicate. However, it will not necessarily do it: SMT solvers are based on heuristic methods, so we can just provide them properties and hope that they will use it.

7. Proof methodologies

Here, the property is simple to prove, so it was not really necessary to add assertions (and make more effort). In other cases, as we will see now, we have to give the right information so that they will find what they need to finish the proof.

7.2.2. Triggering lemmas

We often use assertions to express properties that precisely correspond to one or all of the premises of a lemma or to its conclusion. By doing this, we maximize the chances that the SMT solver “recognize” that what we have written corresponds to a particular lemma and that it should use it.

Let us illustrate this use with the following example. We use axioms and not lemmas as they are used exactly the same way by WP. First consider the following axiomatic definition. We define two predicates `P` and `Q` about a particular memory location `x`. We have two axioms: `ax_1` that states that if `P(x)` is true, then `Q(x)` is true, and a second axiom `ax_2` that state that if the pointed location does not change between two labels (modeled by the predicate `eq`) and `P(x)` holds for one of them, then it holds for the other one.

```
1  /*@
2  predicate eq{L1, L2}(int* x) =
3    \at(*x, L1) == \at(*x, L2) ;
4  */
5
6  /*@
7  axiomatic Ax {
8    predicate P(int* x) reads *x ;
9    predicate Q(int* x) reads *x ;
10
11    axiom ax_1: \forall int* x ; P(x) ==> Q(x);
12    axiom ax_2{L1, L2}:
13      \forall int* x ; eq{L1, L2}(x) ==> P{L1}(x) ==> P{L2}(x);
14  }
15  */
```

And we want to prove the following program:

```
17 /*@ assigns *x ; */
18 void g(int* x);
19
20 /*@
21 requires \separated(x, y);
22 requires P(x) ;
23 ensures Q(x) ;
24 */
25 void example(int* x, int* y){
26   g(y);
27 }
```

However, we can see here that the proof fails on the following verification condition (again we have removed what is not useful for our explanations):

7. Proof methodologies

```

/*@ requires \separated(x, y);
   requires P(x);
   ensures Q(\old(x)); */
void example(int *x, int *y)
{
    g(y);
    return;
}

```

Information Messages (1) Console Properties Values Red Alarms WP Goals

Full Context Binary Non Proved Prop

Goal Post-condition:

Assume {

 Stmt { L1: }

 (* Pre-condition *)

 Have: (y@L1 != x@L1) /\ P_P(μ:Mint@L1, x@L1).

 Stmt { *y@L1 = v_0; }

 Stmt { L3: }

}

Prove: P_Q(μ:Mint@L3, x@L1).

Thus, the SMT solver seems to be unable to use one of the two axioms of our definition: either it cannot show that after the call to `g(y)`, `P(x)` is still true, or it can and thus it cannot show that it implies that `Q(x)` is true. Let us try to add an assertion so that we verify that we can prove `P(x)` after the call:

```

20  /*@
21    requires \separated(x, y);
22    requires P(x) ;
23    ensures Q(x) ;
24  */
25  void example(int* x, int* y){
26      g(y);
27      /*@ assert P(x);
28  }

```

```

/*@ requires \separated(x, y);
   requires P(x);
   ensures Q(\old(x)); */
void example(int *x, int *y)
{
    g(y);
    /*@ assert P(x); */ ;
    return;
}

```

Information Messages (1) Console Properties Values Red Alarms WP Goals

Full Context Binary Non Proved Prop

Goal Assertion:

Assume {

 Stmt { L1: }

 (* Pre-condition *)

 Have: (y@L1 != x@L1) /\ P_P(μ:Mint@L1, x@L1).

 Stmt { *y@L1 = v_0; }

 Stmt { L3: }

}

Prove: P_P(μ:Mint@L3, x@L1).

It seems that despite the fact that it is clear that `*x` did not change during the call `g(y)`, and thus that `eq{Pre, Here}(x)` holds after the call, since this property is not directly

7. Proof methodologies

provided in this verification condition, the SMT solver does not use the corresponding axiom `ax_2`. Thus, let us provide this information to the solver:

```
20 /*@
21   requires \separated(x, y);
22   requires P(x) ;
23   ensures  Q(x) ;
24 */
25 void example(int* x, int* y){
26   g(y);
27   //@ assert eq{Pre, Here}(x);
28 }
```

Now, everything is proved as, if we have a look at the verification condition, we can see that this important information is provided so that the SMT solver can use it:

The screenshot shows a verification tool interface. At the top, a C code snippet is displayed with syntax highlighting. Below the code, there are several tabs: Information, Messages (1), Console, Properties, Values, Red Alarms, and WP Goals. The WP Goals tab is selected, showing a goal post-condition. Below the tabs, there are navigation buttons (back, forward, search, etc.) and a dropdown menu set to 'Full Context'. The main area displays the verification conditions for the goal.

```
Goal Post-condition:
Let m_0 = μ:Mint@L3.
Assume {
  Stmt { L1: }
  (* Pre-condition *)
  Have: (y@L1 != x@L1) /\ P_P(μ:Mint@L1, x@L1).
  Stmt { *y@L1 = v_0; }
  Stmt { L3: }
  (* Assertion *)
  Have: P_eq(m_0, μ:Mint@L1, x@L1).
}
Prove: P_Q(m_0, x@L1).
```

7.2.3. A more complex example: sort, again

Let us now consider a more complex example that involves some actual axiomatic definitions. This time, we will prove the insertion sort function:

```
1  #include <stddef.h>
2  #include <limits.h>
3
4  void insert(int* a, size_t beg, size_t last){
5    size_t i = last ;
6    int value = a[i] ;
7
8    while(i > beg && a[i - 1] > value){
9      a[i] = a[i - 1] ;
10     --i ;
11   }
12   a[i] = value ;
13 }
```

7. Proof methodologies

```
14
15 void insertion_sort(int* a, size_t beg, size_t end){
16     for(size_t i = beg+1; i < end; ++i)
17         insert(a, beg, i);
18 }
```

The `insertion_sort` function visits each value, from the beginning of the array, to the end. For each value v , it is inserted (using the function `insert`) at the right place in the range of the already sorted values (at the beginning of the array), by shifting them until it meets a value that is smaller than v or the first cell of the array.

We want to prove the same postcondition as we already proved for the selection sort, that is: we want to create a sorted permutation of the original values. Again, each iteration of the loop must ensure that the new configuration is a permutation of the original values, and that the range from the beginning to the current visited cell is sorted. All these properties are ensured by the `insert` function. If we give a closer look at this function, we can see that it records the value to insert (which is at the end of the range) in the variable `value` and starting from this last position it shifts all value until it meets a value that is smaller than the one we want to insert or the first cell of the array, and finally inserts the value.

First, let us provide a suitable contract and loop invariant for the insertion sort function. The contract is equivalent to the one provided for the selection sort. Note however that the invariant is weaker: we do not need the values that are still not visited to be greater than the visited ones: we insert each value at the right place.

```
63 /*@
64   requires beg < end && \valid(a + (beg .. end-1));
65   assigns a[beg .. end-1];
66   ensures sorted(a, beg, end);
67   ensures permutation{Pre, Post}(a,beg,end);
68 */
69 void insertion_sort(int* a, size_t beg, size_t end){
70     /*@
71     loop invariant beg+1 <= i <= end ;
72     loop invariant sorted(a, beg, i) ;
73     loop invariant permutation{Pre, Here}(a,beg,end);
74     loop assigns a[beg .. end-1], i ;
75     loop variant end-i ;
76     */
77     for(size_t i = beg+1; i < end; ++i) {
78         insert(a, beg, i);
79     }
80 }
```

Now, we can provide a contract for the insert function. The function requires that the first part of the range is already sorted from the beginning to the penultimate, and in exchange, it guarantees that the complete final range is sorted and is a permutation of the original one:

```
43 /*@
44   requires beg < last < UINT_MAX && \valid(a + (beg .. last));
45   requires sorted(a, beg, last) ;
46
47   assigns a[ beg .. last ] ;
48
```

7. Proof methodologies

```
49  ensures permutation{Pre, Post}(a, beg, last+1);
50  ensures sorted(a, beg, last+1) ;
51  */
52  void insert(int* a, size_t beg, size_t last){
53      size_t i = last ;
54      int value = a[i] ;
55
56      while(i > beg && a[i - 1] > value){
57          a[i] = a[i - 1] ;
58          --i ;
59      }
60      a[i] = value ;
61  }
```

Then, we need to provide a suitable invariant to the loop of the `insert` function. And this time, we can see that with our previously defined permutation predicate, we are in trouble. Indeed, our inductive definition of the permutation specifies three cases: a range is permutation of itself, or two (and only two) values have been swapped, or the permutation is transitive. But none of this cases can be applied to our `insert` function, since the shifted range is not obtained by successively exchanging values, and that the other cases obviously do not apply here.

Thus, we need to find a better definition to the notion of permutation. We can notice that what we really need to provide is a way to say “each value that was previously in the array is still in the array and if several values were equivalent, the number of occurrences of these values does not change”. And in fact, this last part of the expression is enough to express our permutation. A permutation of a range is a range such that for all value, the number of occurrences of this value in the array does not change from a program point to another:

```
37  /*@
38  predicate permutation{L1, L2}(int* in, integer from, integer to) =
39      \forall int v ; l_occurrences_of{L1}(v, in, from, to) ==
40                      l_occurrences_of{L2}(v, in, from, to) ;
41  */
```

Starting from this definition, we are able to provide lemmas that allow to reason efficiently about permutation, provided that some properties hold about the array between two program points. For example, it would be possible to define the case `Swap` in our previous inductive definition using a lemma. And it is of course also possible for our shifted range.

Let us determine what are the required lemmas by first considering the function `insert_sort`. The only property that is not proved is the invariant of the loop that expresses the fact that the array is a permutation of the original array. How can we deduce it? (We will consider the proofs of the lemmas later).

We can observe two facts: the first range in the array (from `beg` to `i+1`) is a permutation of the same range at the beginning of the iteration (by the contract of the `insert` function). The second part (from `i+1` to `end`) is unchanged, thus this is also a permutation. Let us use some assertions to see whether we can prove this or not. While the first property is easily deduced, we can see that the second one is not:

7. Proof methodologies

```

1  /*@
2    loop invariant beg+1 <= i <= end ;
3    loop invariant sorted(a, beg, i) ;
4    loop invariant permutation{Pre, Here}(a,beg,end);
5    loop assigns a[beg .. end-1], i ;
6    loop variant end-i ;
7  */
8  for(size_t i = beg+1; i < end; ++i) {
9    //@ ghost L:
10   insert(a, beg, i);
11   //@ assert permutation{L, Here}(a, beg, i+1); // PROVED
12   //@ assert permutation{L, Here}(a, i+1, end); // NOT PROVED
13 }

```

So we need a first lemma for this property. Let us define two predicates `shifted` and `unchanged`, the later being used to define the former (we will see why later) and express that an unchanged range is a permutation.

```

48 /*@
49   predicate shifted{L1, L2}(integer s, int* a, integer beg, integer end) =
50     \forall integer k ; beg <= k < end ==> \at(a[k], L1) == \at(a[s+k], L2) ;
51 */
52 /*@
53   predicate unchanged{L1, L2}(int* a, integer beg, integer end) =
54     shifted{L1, L2}(0, a, beg, end);
55 */

```

```

67 /*@ lemma unchanged_is_permutation{L1, L2}:
68   \forall int* a, integer beg, integer end ;
69   unchanged{L1, L2}(a, beg, end) ==> permutation{L1, L2}(a, beg, end) ;
70 */

```

And now, we can verify that those two sub-arrays are permutations, this is done by adding an assertion that shows that the range `i+1` to `end` is unchanged, in order to trigger our lemma `unchanged_is_permutation`.

```

129 /*@
130   loop invariant beg+1 <= i <= end ;
131   loop invariant sorted(a, beg, i) ;
132   loop invariant permutation{Pre, Here}(a,beg,end);
133   loop assigns a[beg .. end-1], i ;
134   loop variant end-i ;
135 */
136 for(size_t i = beg+1; i < end; ++i) {
137   //@ ghost L:
138   insert(a, beg, i);
139   //@ assert permutation{L, Here}(a, beg, i+1);
140   //@ assert unchanged{L, Here}(a, i+1, end) ;
141   //@ assert permutation{L, Here}(a, i+1, end) ;
142 }

```

Thus, since those two parts are permutations, the global array is a permutation of the values present at the beginning of the iteration. However, this is not directly proved, so we also need a lemma for this:

7. Proof methodologies

```

75 /*@ lemma union_permutation{L1, L2}:
76   \forall int* a, integer beg, split, end, int v ;
77     beg <= split <= end ==>
78     permutation{L1, L2}(a, beg, split) ==>
79     permutation{L1, L2}(a, split, end) ==>
80     permutation{L1, L2}(a, beg, end) ;
81 */

```

And now we can deduce that a loop iteration produces a permutation by adding this conclusion as an assertion:

```

1   /*@ ghost L: ;
2   insert(a, beg, i);
3   /*@ assert permutation{L, Here}(a, beg, i+1);
4   /*@ assert unchanged{L, Here}(a, i+1, end);
5   /*@ assert permutation{L, Here}(a, i+1, end);
6   /*@ assert permutation{L, Here}(a, beg, end); // PROVED

```

Finally, we need to add one more information, the permutation of a permutation is also a permutation. This time, we do not have to add more assertions, the context contains:

- `permutation{Pre, L}(a, beg, end)` (invariant)
- `permutation{L, Here}(a, beg, end)` (assertion)

which is enough to conclude `permutation{Pre, Here}(a, beg, end)` at the end of the loop block using the following lemma:

```

42 /*@ lemma transitive_permutation{L1, L2, L3}:
43   \forall int* a, integer beg, integer end ;
44     permutation{L1, L2}(a, beg, end) ==>
45     permutation{L2, L3}(a, beg, end) ==>
46     permutation{L1, L3}(a, beg, end) ;
47 */

```

Now, we can have a closer look at the insertion function, by first considering how to maintain that the function produces a permutation.

It shifts the different elements to the left until it reaches the beginning of the array or an element that is smaller than the element to insert which is initially at the end of the range, and is inserted at the reached position. The cells from the beginning of the array to the location of insertion are left unchanged, so this is a permutation. We have a lemma for this, but we also have to state that the values are unchanged as an invariant of the loop. The second part of the array is a permutation because we rotate the elements, we need a lemma to express this, and to state at least that the element are shifted by the loop as an invariant. Finally the union of the two permutations is a permutation, and we also have a lemma for this.

So first, we can give a suitable invariant for the permutation:

- we provide the bounds of `i`
- we state that the first part is left unchanged

7. Proof methodologies

- we state that the last part is shifted to the left

as well as some assertions that we want to be verified:

- first in order to trigger `unchanged_permutation`, we place a first assertion that state that the first part of the array is unchanged, which allows to prove
- the second one that asserts that the first part of the array is a permutation of the original one, and is used in combination with ...
- the third one that asserts that the second part of the array is a permutation of the original one (which allows to trigger `union_permutation` and prove the postcondition).

```

1  /*@
2    loop invariant beg <= i <= last ;
3    loop invariant \forall integer k ; beg <= k <= i ==> a[k] == \at(a[k], Pre) ;
4    loop invariant \forall integer k ; i+1 <= k <= last ==> a[k] == \at(a[k-1], Pre) ;
5
6    loop assigns i, a[beg .. last] ;
7    loop variant i ;
8  */
9  while(i > beg && a[i - 1] > value){
10     a[i] = a[i - 1] ;
11     --i ;
12  }
13
14  a[i] = value ;
15
16  //@ assert unchanged{Pre, Here}(a, beg, i) ; // PROVED
17  //@ assert permutation{Pre, Here}(a, beg, i) ; // PROVED
18
19  //@ assert rotate_left{Pre, Here}(a, i, last+1) ; //PROVED
20  //@ assert permutation{Pre, Here}(a, i, last+1) ; // NOT PROVED

```

Then, for the last assertion, we need a lemma about the rotation of the elements:

```

56 /*@
57 predicate rotate_left{L1, L2}(int* a, integer beg, integer end) =
58     beg < end && \at(a[beg], L2) == \at(a[end-1], L1) &&
59     shifted{L1, L2}(1, a, beg, end - 1) ;
60 */

```

```

71 /*@ lemma rotate_left_is_permutation{L1, L2}:
72     \forall int* a, integer beg, integer end ;
73     rotate_left{L1, L2}(a, beg, end) ==> permutation{L1, L2}(a, beg, end) ;
74 */

```

We also have to help a little bit the provers to show that the range is sorted after the insertion. For this, we provide a new invariant to show that the shifted values are greater than the value to insert, and then we add some assertions to show that the array is sorted before the insertion, that all values before the cell where we insert are lower than the inserted value, and that the range is consequently sorted after the insertion. Leading us to the following annotated `insert` function:

7. Proof methodologies

```

83  /*@
84   requires beg < last < UINT_MAX && \valid(a + (beg .. last));
85   requires sorted(a, beg, last) ;
86
87   assigns a[ beg .. last ] ;
88
89   ensures permutation{Pre, Post}(a, beg, last+1);
90   ensures sorted(a, beg, last+1) ;
91  */
92  void insert(int* a, size_t beg, size_t last){
93      size_t i = last ;
94      int value = a[i] ;
95
96      /*@
97       loop invariant beg <= i <= last ;
98       loop invariant \forall integer k ; i <= k < last ==> a[k] > value ;
99       loop invariant \forall integer k ; beg <= k <= i ==> a[k] == \at(a[k], Pre) ;
100      loop invariant \forall integer k ; i+1 <= k <= last ==> a[k] == \at(a[k-1], Pre) ;
101
102      loop assigns i, a[beg .. last] ;
103      loop variant i ;
104  */
105      while(i > beg && a[i - 1] > value){
106          a[i] = a[i - 1] ;
107          --i ;
108      }
109      //@ assert sorted(a, beg, last+1) ;
110      //@ assert \forall integer k ; beg <= k < i ==> a[k] <= value ;
111      a[i] = value ;
112      //@ assert sorted(a, beg, last+1) ;
113
114      //@ assert unchanged{Pre, Here}(a, beg, i) ;
115      //@ assert permutation{Pre, Here}(a, beg, i) ;
116
117      //@ assert rotate_left{Pre, Here}(a, i, last+1) ;
118      //@ assert permutation{Pre, Here}(a, i, last+1) ;
119  }

```

Overall, we have 6 lemmas to prove:

- `l_occurrences_union`
- `shifted_maintains_occ`
- `unchanged_is_permutation`
- `rotate_left_is_permutation`
- `union_permutation`
- `transitive_permutation`

While the Coq proofs of these lemmas are beyond the scope of this tutorial (and we will see later that in this particular case we can get rid of them), let us give some ideas of how to prove them (nevertheless, the Coq scripts are available on the GitHub repository of this book).

To prove `l_occurrences_union`, we reason by induction on the size of the second part of the array, the basic case is trivial: if the size is 0, we immediately have the equality since `split == to`. So now, we have to prove that if it is true for a range of size i , it is true for a range of size $i + 1$. As we know that it is true until i by our induction hypothesis, we simply analyze the different cases for the last element of the range: either this element is the one we count or not, and it adds the same value on both sides of the equality.

7. Proof methodologies

For `shifted_maintains_occ`, we reason by induction on the complete range, the first case is trivial (empty range), and for the induction case, we just have to show the value added to the range has been shifted and is thus the same.

The property `unchanged_is_permutation` is proved by SMT solvers thanks to the fact that we have expressed it using `shifted`, then the solver can immediately instantiate the previous lemma. If it is not the case (depending on the version of the prover), the proof is done by instantiating `shifted_maintains_occ` with a value of 0 for the shift property.

To prove `rotate_left_is_permutation` we split the range at `L1` into two subparts `beg .. beg+1` and `beg+1 .. end`, and the range at `L2` into two subparts `beg .. end-1` and `end-1 .. end`, using `l_occurrences_union`. We show that the number of occurrences in `beg+1 .. end` at `L1` and `beg .. end-1` at `L2` did not change thanks to `shifted_maintains_occ`, and that the number of occurrences in `beg .. beg+1` at `L1` and `end-1 .. end` at `L2` is the same by case analysis (and using the fact that the values equal).

For `union_permutation`, we instantiate the lemma `l_occurrences_union`. Finally, the lemma `transitive_permutation` is automatically proved by SMT solvers thanks to the transitivity of equality.

7.2.4. How to correctly use assertions?

There is no perfect guideline on when we should use assertions or not. Most of the time we use them to first understand why some proof fails at some point by expressing properties that we expect to be true at some program point. Moreover, most of the time, the verification conditions are too long and complex to be read directly and we somewhat have to keep in mind the different lemmas we have already expressed and to determine whether some deduction requires the use of a lemma that we already expressed, or if it requires to reason by induction on some property or value such that SMT solvers cannot make this deduction and thus we might need to add another lemma.

With a bit of experience, the use of assertions and lemmas become more and more natural, however one has to take care of the fact that it is easy to abuse of these. Basically, the more we add lemmas and assertions, the more the proof context is rich and might contain the required information to produce the proof that we need. However, there is also a risk to add too much information so that the proof context mainly contains garbage that is not useful for the proof but pollute the proof context and makes the job of SMT solvers harder. Thus, we have to find the good trade-off.

7.2.5. Exercises

7.2.5.1. Understanding the proof context

In the following function, the last assertion is automatically proved by SMT solvers, probably by unfolding the predicate and directly proving the corresponding property. Using assertions,

7. Proof methodologies

provide another way to prove the last property. In the proof context, find the generated properties that could allow to prove the assertion and explain how.

```
1  /*@
2    predicate rectangle{L}(integer c1, integer c2, integer h) =
3      c1 * c1 + c2 * c2 == h * h ;
4  */
5
6  /*@
7    requires \separated(x, y, z);
8    requires 3 <= *x <= 5 ;
9    requires 3 <= *y <= 5 ;
10   requires 2 <= *z <= 5 ;
11   requires *x+2 == *y+1 == *z ;
12 */
13 void exercise(int* x, int* y, int* z){
14   *x += 2 * (*x) ;
15   *y += *y ;
16   *y += (*y / 2);
17   *z = 3 * (*z) ;
18   //@ assert rectangle(*x, *y, *z);
19 }
```

7.2.5.2. Trigger lemmas

In the following program, WP fails to prove that the postcondition of the function `g` is verified. Add the right assertion (at the right place), such that the proof succeeds.

```
1  /*@
2    axiomatic Ax {
3      predicate X{L1, L2}(int* p, integer l)
4        reads \at(p[0 .. l-1], L1), \at(p[0 .. l-1], L2) ;
5      predicate Y{L1, L2}(int* p, integer l)
6        reads \at(p[0 .. l-1], L1), \at(p[0 .. l-1], L2) ;
7
8      axiom Ax_axiom_XY {L1,L2}:
9        \forall int* p, integer l, i ; 0 <= i <= l ==> X{L1, L2}(p, i) ==> Y{L1, L2}(p, l) ;
10     axiom transitive{L1,L2,L3}:
11       \forall int* p, integer l ; Y{L1,L2}(p, l) ==> Y{L2,L3}(p, l) ==> Y{L1,L3}(p, l);
12   }
13 */
14
15 /*@
16   assigns p[0 .. l-1] ;
17   ensures X{Pre, Post}(p, l) ;
18 */
19 void f(int* p, unsigned l);
20
21 /*@
22   ensures Y{Pre,Post}(p, l);
23 */
24 void g(int* p, unsigned l){
25   f(p, l) ;
26   f(p, l) ;
27 }
```

7. Proof methodologies

7.2.5.3. Trigger lemmas, under condition

In the following program, WP fails to prove that the postcondition of the function `example` is verified. However, we can notice that the `g` function indirectly ensures that the pointed value is either increased or decreased. Add two assertions that shows which predicate holds depending on the value of `*x`.

```
1  /*@
2  predicate dec{L1, L2}(int* x) =
3    \at(*x, L1) > \at(*x, L2) ;
4  predicate inc{L1, L2}(int* x) =
5    \at(*x, L1) < \at(*x, L2) ;
6  */
7
8  /*@
9  axiomatic Ax {
10   predicate P(int* x) reads *x ;
11   predicate Q(int* x) reads *x ;
12
13   axiom ax_1: \forall int* x ; P(x) ==> Q(x);
14   axiom ax_2{L1, L2}:
15     \forall int* x ; dec{L1, L2}(x) ==> P{L1}(x) ==> P{L2}(x);
16   axiom ax_3{L1, L2}:
17     \forall int* x ; inc{L1, L2}(x) ==> P{L1}(x) ==> P{L2}(x);
18 }
19 */
20
21 /*@
22 assigns *x ;
23 behavior b_1:
24   assumes *x < 0 ;
25   ensures *x >= 0 ;
26 behavior b_2:
27   assumes *x >= 0 ;
28   ensures *x < 0 ;
29 complete behaviors ;
30 disjoint behaviors ;
31 */
32 void g(int* x);
33
34 /*@
35 requires P(x) ;
36 ensures Q(x) ;
37 */
38 void example(int* x){
39   g(x);
40 }
```

The assertions should look like:

```
1  //@ assert *x ... ==> ... ;
2  //@ assert *x ... ==> ... ;
```

Another way to bring some information in the context is to use ghost code. For example, the truth value of a condition appears in the context of a verification condition. Modify the annotation to make the code look like:

7. Proof methodologies

```
1 void example(int* x){
2   g(x);
3   /*@ ghost
4     if ( ... ){
5       /@ assert ... @/
6     } else {
7       /@ assert ... @/
8     }
9   */
10 }
```

Compare the verification condition of each assertion with the previous ones.

Finally, one can in fact notice that “the pointed value is either increased or decreased” can be expressed with a single simple assertion. Write the corresponding annotation.

7.2.5.4. An actual example with sum

The following function increases by 1 the value of a cell in an array, thus it also increases the value of the sum of the content of this array. Write a contract to the function that expresses this fact.

```
1  #include <stddef.h>
2
3  /*@
4   axiomat Sum_array{
5     logic integer sum(int* array, integer begin, integer end) reads array[begin .. (end-1)];
6     axiom empty:
7       \forall int* a, integer b, e; b >= e ==> sum(a,b,e) == 0;
8     axiom range:
9       \forall int* a, integer b, e; b < e ==> sum(a,b,e) == sum(a,b,e-1)+a[e-1];
10   }
11 */
12
13 /*@
14 predicate unchanged{L1, L2}(int* array, integer begin, integer end) =
15   \forall integer i ; begin <= i < end ==> \at(array[i], L1) == \at(array[i], L2) ;
16 */
17
18 /*@
19 lemma sum_separable:
20   \forall int* array, integer begin, split, end ;
21   begin <= split <= end ==> sum(array, begin, end) == ... ?
22 lemma unchanged_sum{L1, L2}:
23   \forall int* array, integer begin, end ;
24   unchanged{L1, L2}(array, begin, end) ==> ... ?
25 */
26
27 void inc_cell(int* array, size_t len, size_t i){
28   array[i]++ ;
29 }
```

In order to prove that this function fulfills its contract, we need to provide some assertions that will guide the proof. More precisely, we have to show that since all values before the modified cell has not been modified, the sum has not been modified in this part of the array, and same for the cells that follow the modified cell.

Thus we need two lemmas:

7. Proof methodologies

- `sum_separable` should express that we can split an array into two subparts, count in each of them and sum the results to get the sum of the entire array,
- `unchanged_sum` should express that if a range of an array has not changed between two labels, the sum of the content is the same.

Complete the code of the lemmas and use assertions to ensure that they will be used to complete the proof. We do not ask the proof of the lemmas, the Coq proofs are available on the GitHub repository of this book.

7.3. More on ghost code: lemma functions and lemma macros

Assertions provide a way to give clues to the verification condition generator so the SMT solvers will get enough information to make the proofs we need. However, it is sometimes hard to write the assertion that will create exactly the property needed by the SMT solver to trigger the right lemma (for example, since the generator makes some optimization on the verification condition that might slightly modify it and the context of the proof). Furthermore, we rely on lemmas that often need to be proved with the Coq proof assistant, and that means that we need to learn Coq.

In this section, we will see some techniques that we can use to make all of this more predictable and does not require from us to use the Coq proof assistant. While these techniques cannot be used in any case (and we will explain what are the cases when it is not applicable), they are quite efficient to get almost full automatic proof. This relies on ghost code.

7.3.1. Proof by induction

Previously, we mentioned that SMT solvers are bad at reasoning by induction (most of them), and this is the reason why we often need to express lemmas that we then prove using the Coq proof assistant that allows us to write our proof by induction. However, in the section 4.2 about loops, we find a subsection 4.2.1 named “Induction and invariant” where we explain how to prove that a loop does the right job ... by induction. What is this sorcery?!

In fact, it is quite simple. When we prove a loop invariant by induction using SMT solvers, they do not have to perform the reasoning by induction themselves. The job of splitting the proof into two subproofs, one for the establishment of the invariant (the base case of the proof), and one for the preservation (the induction case) is performed by the verification condition generator. So when the verification conditions are transmitted to the SMT solvers, this work is not needed anymore.

How can we exploit this idea? We explained before that ghost code can be used to provide more information than what is explicitly provided by the source code. For this, we add some code (and possibly annotations about this code) that allows to deduce more properties. Let us illustrate this with a simple example. In a previous exercise (5.4.4.2), we wanted to prove the following function call (we have excluded the postconditions to shorten the example):

7. Proof methodologies

```

1  #include <stddef.h>
2  #include <limits.h>
3
4  /*@ predicate sorted(int* arr, integer end) =
5      \forall integer i, j ; 0 <= i <= j < end ==> arr[i] <= arr[j] ;
6      predicate element_level_sorted(int* array, integer end) =
7          \forall integer i ; 0 <= i < end-1 ==> array[i] <= array[i+1] ;
8  */
9
10 /*@ requires \valid_read(arr + (0 .. len-1));
11     requires sorted(arr, len) ;
12 */
13 size_t bsearch(int* arr, size_t len, int value);
14
15 /*@ requires \valid_read(arr + (0 .. len-1));
16     requires element_level_sorted(arr, len) ;
17 */
18 unsigned bsearch_callee(int* arr, size_t len, int value){
19     return bsearch(arr, len, value);
20 }

```

For this, the solution that was asked in the exercise was to provide a lemma that states that if a range is “locally sorted”, meaning that each element is greater or equals to the one that precedes it, then we can say that it is “globally sorted”, that is to say for each pair of indices i and j , if $i \leq j$ then the j^{th} of the array is greater or equals to the i^{th} element. Then, the precondition could be proved by SMT solvers, but not the lemma itself that requires a Coq proof. Can we do something using a loop?

The answer is yes. Before calling the function, we can build a proof that shows that because the array is locally sorted, we can deduce that it is globally sorted (which is basically a proof of the lemma we would need). We want to prove that the range is globally sorted. To write this proof by hand, we would proceed by induction on the size of the range. We have two cases. First, the range is empty and the property trivially true. Second, let us suppose that some range of size i with $i < length$ ($length$ being the size of the complete range), is globally sorted and show that if it is the case, then the range of size $i + 1$ is sorted. This is easy because, by our precondition, we know that the i^{th} element is greater than the $(i - 1)^{th}$ element, that is itself greater than all the preceding elements.

Now, how can we translate this into ghost code? We write a loop that goes from 0 (our base case), to the end of the range `len` and provide as an invariant that the array is globally sorted from 0 to the current step of the loop. We also add some assertions to help the provers (namely the fact that the current element is greater than the one that precedes it):

```

15 /*@ requires \valid_read(arr + (0 .. len-1));
16     requires element_level_sorted(arr, len) ;
17 */
18 unsigned bsearch_callee(int* arr, size_t len, int value){
19     /*@ ghost
20         /@
21         loop invariant 0 <= i <= len ;
22         loop invariant sorted(arr, i) ;
23         loop assigns i ;
24         loop variant len-i ;
25         @/
26         for(size_t i = 0 ; i < len ; ++i){
27             /*@ assert 0 < i ==> arr[i-1] <= arr[i] ; @/
28         }

```

7. Proof methodologies

```

29  */
30  return bsearch(arr, len, value);
31  }

```

And we can see that all verification conditions are easily verified by SMT solvers, without requiring to write a Coq proof or a lemma. The verification conditions that are created respectively for the establishment and the preservation of the invariant correspond to the two cases we needed to prove in our proof by induction.

Base case : size 0

```

size t i = (unsigned int)0;
@/ loop invariant 0 ≤ i ≤ len;
loop invariant sorted(arr, i);
loop assigns i;
loop variant len - i;
@/
while (i < len) {

```

Information Messages (2) Console Properties Values Red Alarms WP Goals

Global All Results

Raw Obligation Binary Proved Goal

Goal Invariant (established):

Assume {

(* Pre-condition *)

Have: P element level_sorted(Mint_0, arr_0, len_0) /\

valid_rd(Malloc_0, shift_sint32(arr_0, 0), len_0).

}

Prove: P_sorted(Mint_0, arr_0, 0).

```

size t i = (unsigned int)0;
@/ loop invariant 0 ≤ i ≤ len;
loop invariant sorted(arr, i);
loop assigns i;
loop variant len - i;
@/
while (i < len) {

```

Information Messages (2) Console Properties Values Red Alarms WP Goals

Global All Results

Raw Obligation Binary Proved Goal

Goal Invariant (preserved):

Let x = Mint_0[shift_sint32(arr_0, i - 1)].

Let x_1 = Mint_0[shift_sint32(arr_0, i)].

Induction Hypothesis: Property true at rank i

Assume {

(* Pre-condition *)

Have: P element level_sorted(Mint_0, arr_0, len_0) /\

valid_rd(Malloc_0, shift_sint32(arr_0, 0), len_0).

(* Invariant *)

Have: P_sorted(Mint_0, arr_0, i).

(* Invariant *)

Have: (0 ≤ i) /\ (i ≤ len_0).

(* Then *)

Have: i < len_0.

(* Assertion *)

Have: ((0 < i) -> (x ≤ x_1)).

Prove: P_sorted(Mint_0, arr_0, to_uint32(1 + i)).

Prove that it is true at rank i+1

This kind of code is called a proof carrying code: we have written some code and annotations that carries a proof that some property we want to verify holds.

Note that here, as we can write quite a lot of ghost code, we augment our risk to modify the behavior of the program by mistake. Most of the time we do not need to modify any

7. Proof methodologies

memory location, thus taking care to check (few) assigned memory locations (for example, loop iterators), that we do not create a runtime error (through RTE plugin) and that the ghost code terminates is enough to be confident about the verification.

In this example, we had to write the ghost code directly in annotation of the program, and that mean that if we have another call somewhere else in the code with some similar precondition, we would have to do it again. Let us make this easier by using lemma functions.

7.3.2. Lemma function

The principle of lemma functions is basically the same as lemmas: from some premises, we want to prove some conclusion. And once it has been done, we want to use it at some other place to directly deduce the conclusion from the premises without having to do the proof again, by instantiating it with actual values.

The way to do this is to use a function, using the `requires` to express the premises of the lemma, and the `ensures` to express the conclusion of the lemma. The universally quantified variables can either still be quantified, or correspond to a parameter of the function. Namely, if a variable is only bounded to premises, or only to conclusions, it can be translated into a universally quantified variable, provided that it is not necessary to bind its value to a C variable in our proof carrying code (a quantified variable is not visible from the C code). If it is bounded to both premises and conclusion, it must be a parameter of the function (as we cannot quantify a variable for all a function contract in ACSL).

Let us first consider an example where we do not use (directly) universally quantified variables in the contract, with our previous example about sorted values. From the property `element_level_sorted(arr, len)`, we want to deduce `sorted(arr, len)`. The corresponding lemma could be:

```
1  /*@
2   lemma element_level_sorted_is_sorted:
3     \forall int* arr, integer len ;
4       element_level_sorted(arr, len) ==> sorted(arr, len) ;
5  */
```

So let us write a function that takes two parameters: `arr` and `len`, requires that the array is locally sorted and ensures that it is globally sorted:

```
16 /*@
17   requires element_level_sorted(arr, len) ;
18   assigns \nothing ;
19   ensures sorted(arr, len);
20 */
21 void element_level_sorted_implies_sorted(int* arr, size_t len);
```

Note that this function must assign `\nothing`, indeed we use it to deduce some properties about the program, with ghost code, and thus it should not modify the content of the array, else the ghost code would modify the behavior of the program. Now let us provide a body to this function, the proof carrying code that guarantees that the conclusion is verified, provided

7. Proof methodologies

that the precondition holds. It corresponds to the code we previously wrote to prove the precondition of the call to `bsearch`:

```
16  /*@
17   requires element_level_sorted(arr, len) ;
18   assigns  \nothing ;
19   ensures  sorted(arr, len);
20  */
21  void element_level_sorted_implies_sorted(int* arr, size_t len){
22    /*@
23     loop invariant 0 <= i <= len ;
24     loop invariant sorted(arr, i) ;
25     loop assigns i ;
26     loop variant len-i ;
27    */
28    for(size_t i = 0 ; i < len ; ++i){
29      /*@ assert 0 < i ==> arr[i-1] <= arr[i] ;
30     */
31    }
```

With this specified loop, we get an inductive proof that the lemma holds, now we can use this lemma function by simply calling it when we need to perform this deduction:

```
34  /*@ requires \valid_read(arr + (0 .. len-1));
35   requires element_level_sorted(arr, len) ;
36  */
37  unsigned bsearch_callee(int* arr, size_t len, int value){
38    /*@ ghost element_level_sorted_implies_sorted(arr, len) ;
39   return bsearch(arr, len, value);
40  }
```

Which asks us to establish the premises thanks to the precondition of the lemma function (and which we trivially get from the precondition of the `bsearch_callee` function), and provides us the conclusion for free thanks to the postcondition of the lemma function (and we can use it as a precondition for the call to `bsearch`).

As we explained, when universally quantified variables are bounded to both conclusion and premises, they must be parameters, and it is the case here for the variables `arr` and `len`. Whereas the quantified variable that are used in the predicates:

```
4  /*@ predicate sorted(int* arr, integer end) =
5     \forall integer i, j ; 0 <= i <= j < end ==> arr[i] <= arr[j] ;
6     predicate element_level_sorted(int* array, integer end) =
7         \forall integer i ; 0 <= i < end-1 ==> array[i] <= array[i+1] ;
8  */
```

as they are only bound to respectively the premise and the conclusion remain universally quantified (even if it is hidden by the predicate). We could for example have written the contract like this:

7. Proof methodologies

```
16  /*@
17   requires \forall integer i ; 0 <= i < len-1 ==> arr[i] <= arr[i+1] ;
18   assigns  \nothing ;
19   ensures  \forall integer i, j ; 0 <= i <= j < len ==> arr[i] <= arr[j] ;
20  */
21  void element_level_sorted_implies_sorted(int* arr, size_t len){
22    /*@
23     loop invariant 0 <= i <= len ;
24     loop invariant sorted(arr, i) ;
25     loop assigns i ;
26     loop variant len-i ;
27    */
28    for(size_t i = 0 ; i < len ; ++i){
29      //@ assert 0 < i ==> arr[i-1] <= arr[i] ;
30    }
31  }
```

where we perfectly see that variables are still universally quantified. However, we are not forced to maintain them universally quantified, and we could perfectly translate them into parameters (provided that the conclusion we want to get from the premises still makes sense). Let us for example translate the `i` and `j` of the conclusion into parameters:

```
34  /*@
35   requires \forall integer i ; 0 <= i < len-1 ==> arr[i] <= arr[i+1] ;
36   assigns  \nothing ;
37   ensures  0 <= i <= j < len ==> arr[i] <= arr[j] ;
38  */
39  void element_level_sorted_implies_greater(int* arr, size_t len, size_t i, size_t j){
40    //@ ghost element_level_sorted_implies_sorted(arr, len);
41  }
```

Which is also perfectly fine and we could for example use this function to deduce some properties about the content of the array. Note that here, we use a call to the previous lemma function to make the proof easier. We even can go further by transferring the “premise of our conclusion” as another premise of a new lemma:

```
44  /*@
45   requires \forall integer i ; 0 <= i < len-1 ==> arr[i] <= arr[i+1] ;
46   requires 0 <= i <= j < len ;
47   assigns  \nothing ;
48   ensures  arr[i] <= arr[j] ;
49  */
50  void element_level_sorted_implies_greater_2(int* arr, size_t len, size_t i, size_t j){
51    //@ ghost element_level_sorted_implies_sorted(arr, len);
52  }
```

All these lemmas state the same global relation, the difference is related to the amount of information that is required to instantiate them (and thus the precision of the property that we get in return).

Finally, let us present a last usage of lemma functions. On all previous examples, we have considered only universally quantified variable. In fact, what we have said before is applicable to existentially quantified variables: if they are bound to both premises and conclusions, they must be parameters, else they can either be parameters or remain quantified. However, about

7. Proof methodologies

existentially quantified variables, we sometimes can go further by building a function that directly provide a witness for an existentially quantified formula.

For example, let us consider the axiomatic definition for occurrence counting, and imagine that at some point in a program, we want to prove the following assertion from the precondition:

```
24  /*@
25   requires \valid(in + (0 .. len-1)) ;
26   requires l_occurrences_of(v, in, 0, len) > 0 ;
27  */
28  void foo(int v, int* in, size_t len){
29      //@ assert \exists integer n ; 0 <= n < len && in[n] == v ;
30
31      // ... code
32  }
```

Of course, there exists some index `n` such that `in[n]` is `v`, else the number of occurrences of this value would be 0. But, instead of just proving that such an index exists, let us directly find some index that respects the constraints on `n` by using a lemma function that returns it:

```
24  /*@
25   requires \valid(in + (0 .. len-1)) ;
26   requires l_occurrences_of(v, in, 0, len) > 0 ;
27   assigns \nothing ;
28   ensures 0 <= \result < len && in[\result] == v ;
29  */
30  size_t occ_not_zero_some_is_v(int v, int* in, size_t len){
31      /*@
32       loop invariant 0 <= i < len ;
33       loop invariant l_occurrences_of(v, in, 0, i) == 0 ;
34       loop assigns i ;
35       loop variant len-i ;
36      */
37      for(size_t i = 0 ; i < len ; ++i){
38          if(in[i] == v) return i ;
39      }
40      //@ assert \false ;
41      return -1 ;
42  }
```

If we only look at the body of the function, it has two behaviors: either some cell of the array contains `v` and the function returns its index, or there is not, and then the function returns -1. The first behavior is easy to show, the return statement is performed in a branch where we know that the considered index corresponds to a cell that is in the range of the array and has a value `v`.

We prove that the second behavior respects the postcondition by showing that it leads to a contradiction. If there is no cell of value `v`, then the number of occurrences of `v` is 0, this is expressed by the second invariant that shows that as we have not met any `v` from the beginning of the loop, the number of occurrences is 0. However, the precondition of the function states that the number of occurrences is more than 0 which leads to a contradiction that we model model by an assertion of false (note that this is not necessary, we explicitly write it for our explanation) which means here that this path is infeasible.

7. Proof methodologies

Finally, we can call this function to show that there exists some index that allows our assertion to be validated:

```
44  /*@
45   requires \valid(in + (0 .. len-1)) ;
46   requires l_occurrences_of(v, in, 0, len) > 0 ;
47  */
48  void foo(int v, int* in, size_t len){
49   /*@ ghost size_t witness = occ_not_zero_some_is_v(v, in, len);
50   /*@ assert \exists integer n ; 0 <= n < len && in[n] == v ;
51
52   // ... code
53 }
```

The use of lemma functions makes reasoning by induction feasible for lemmas without the need of interactive proof. Furthermore, the triggering of lemmas becomes more predictable as we instantiate them by hand. However, while lemmas can consider multiple labels:

```
1  /*@
2   lemma my_lemma{L1, L2}: P{L1} ==> P{L2} ;
3  */
```

Lemma functions do not provide an equivalent mechanism as they are basically normal C functions that cannot take labels in input. Let us show what we can do when we need such a construct.

7.3.3. Lemma macro

When we have to deal with multiple labels, the idea is to directly “inject” the proof carrying code at the place where it is needed exactly as we did at the beginning of the section. However, we do not want to write this code by hand everytime we need such a proof, so let us use macros to do it.

For now, let us translate our previous code into a macro instead of a function. As we use this macro in ghost code (thus, in annotation) we have to take care to use the ghost annotation syntax to write the invariant of the loop and the assertions:

```
16  #define element_level_sorted_implies_sorted(_arr, _len) \
17   /*@ assert element_level_sorted(_arr, _len) ; @/ \
18   /*@ loop invariant 0 <= _i <= _len ; \
19   loop invariant sorted(_arr, _i) ; \
20   loop assigns _i ; \
21   loop variant _len-_i ; @/ \
22   for(size_t _i = 0 ; _i < _len ; ++_i){ \
23   /*@ assert 0 < _i ==> _arr[_i-1] <= _arr[_i] ; @/ \
24   } \
25   /*@ assert sorted(_arr, _len) ; @/
26
27  /*@ requires \valid_read(arr + (0 .. len-1));
28  requires element_level_sorted(arr, len) ;
29  */
30  unsigned bsearch_callee(int* arr, size_t len, int value){
31   /*@ ghost element_level_sorted_implies_sorted(arr, len) ;
```

7. Proof methodologies

```

32  return bsearch(arr, len, value);
33  }

```

Instead of providing a pre and a postcondition, we state these properties using assertions before and after the proof carrying code. The proof carrying code itself is basically the same as before, and it is used exactly as it was used in the case of functions. However, we can see that it makes an important difference once it has been preprocessed by Frama-C as the block of code and annotations is directly injected in the function `bsearch_callee`.

```

/*@ requires \valid_read(arr + (0 .. len - 1));
   requires element_level_sorted(arr, len);
   */
   unsigned int bsearch_callee(int *arr, size_t len, int value)
   {
       unsigned int tmp;
       /*@ assert element_level_sorted(arr, len); */
       /*@ ghost {
           size_t _i = (unsigned int)0;
           /*@ loop invariant 0 ≤ _i ≤ len;
              loop invariant sorted(arr, _i);
              loop assigns _i;
              loop variant len - _i;
           */
           while (_i < len) {
               /*@ assert 0 < _i ⇒ *(arr + (_i - 1)) ≤ *(arr + _i); */
               ;
               _i += (size_t)1;
           }
       }
       /*@
       /*@ assert sorted(arr, len); */
       tmp = bsearch(arr, len, value);
       return tmp;
   }

```

So in fact, we use the macro to generate the code we previously wrote. In this case, it is not really interesting as a function call allows us to make things more modular. So let us study a case where we do not have any other choice than using a macro.

We illustrate using the following lemma:

```

4  /*@
5  predicate shifted{L1, L2}(int* arr, integer fst, integer last, integer shift) =
6      \forall integer i ; fst <= i < last ==> \at(arr[i], L1) == \at(arr[i+shift], L2) ;
7
8  lemma shift_ptr{L1, L2}:
9      \forall int* arr, integer fst, integer last, integer s1, s2 ;
10     shifted{L1, L2}(arr, fst+s1, last+s1, s2) ==> shifted{L1, L2}(arr+s1, fst, last, s2) ;
11 */

```

In order to prove the following program:

```

13 /*@
14 requires \valid(array+(beg .. end+shift-1)) ;
15 requires shift + end <= UINT_MAX ;
16 assigns array[beg+shift .. end+shift-1];
17 ensures shifted{Pre, Post}(array, beg, end, shift) ;

```

7. Proof methodologies

```

18  */
19  void shift_array(int* array, size_t beg, size_t end, size_t shift);
20
21  /*@
22   requires \valid(array+(0 .. len+s1+s2-1)) ;
23   requires s1+s2 + len <= UINT_MAX ;
24   assigns array[s1 .. s1+s2+len-1];
25   ensures shifted{Pre, Post}(array+s1, 0, len, s2) ;
26  */
27  void callee(int* array, size_t len, size_t s1, size_t s2){
28   shift_array(array, s1, s1+len, s2) ;
29  }

```

Where the lemma `shift_ptr` is necessary to prove the postcondition of `callee` from the postcondition of `shift_array`. Our goal is of course to get rid of the lemma, replacing it by a lemma macro.

There is no precise guideline for designing a macro used from the injection of proof carrying code. However, most lemmas stated about multiple labels are quite similar in the way they relate labels. So let us illustrate with this example, most of the time designing a macro in such a situation will more or less follow the same scheme.

In order to build the macro, we need a context where we can work on it. We build the context using a function, let us name this function `context_to_prove_shift_ptr`. The idea is to use the function to build the macro in isolation of the rest of the program to make the verification of the property easier. However, while lemma functions are then called to deduce some properties in some other function, this function will never be called, its only role is to provide us a “place” where we can build our proof. In particular, as we need multiple memory labels, our function **needs** to modify the content of the memory (else, there is a single memory state for all the function).

Let us illustrate with our current problem to make all of this clearer. First, we create a macro `shift_array` that will contain our proof carrying code, for now let us just indicate that it is an empty statement. In the parameters of this lemma, we take the labels that are considered. Note that the rules we previously mentioned about quantified variables still apply to macros.

```

1  #define shift_ptr(_L1, _L2, _arr, _fst, _last, _s1, _s2) ;

```

Then we create our context function:

```

22  /*@
23   assigns arr[fst+s1+s2 .. last+s1+s2] ;
24   ensures shifted{Pre, Post}(arr, fst+s1, last+s1, s2) ;
25  */
26  void assign_array(int* arr, size_t fst, size_t last, size_t s1, size_t s2);
27
28  /*@
29   requires fst <= last ;
30   requires s1+s2+last <= UINT_MAX ;
31  */
32  void context_to_prove_shift_ptr(int* arr, size_t fst, size_t last, size_t s1, size_t s2){
33   L1: ;
34   assign_array(arr, fst, last, s1, s2);
35   L2: ;

```

7. Proof methodologies

```

36  //@ assert shifted{L1, L2}(arr, fst+s1, last+s1, s2) ;
37
38  //@ ghost shift_ptr(L1, L2, arr, fst, last, s1, s2) ;
39
40  //@ assert shifted{L1, L2}(arr+s1, fst, last, s2) ;
41  }

```

Let us decompose this code, starting from the context function. In input, we receive all the variables of the lemma. We also state some properties about the bounds of the integer values we consider, basically these should be requirements that are not related to memory states, or related to the first one. Then, we introduce the label `L1` and we call the function `assign_array` that leads us to the label `L2`. The role of this call is to ensure that WP will create a new memory label (thus, it will not consider that the memory is the same), and to establish our premises. Indeed, if we have a look at the contract of `assign_array`, we see that it assigns the array (which guarantees the creation of a new memory label) and in postcondition, it ensures that the content of the array, between the pre and the postcondition (thus, when we call it: `L1` and `L2`) respects the premise of our lemma (which we repeat on line 36, by adding an assertion). Then we use our `shift_ptr` macro (that will later contain the proof carrying code), and we then expect to be able to prove the postcondition of our lemma (line 40).

By doing this, we ensure that we built a context that only contains the information we need to build the proof carrying code that allows us to deduce the conclusion (line 40) from the premise (line 36). Now let us write the macro.

```

9  #define shift_ptr(L1, L2, _arr, _fst, _last, _s1, _s2)\
10  /@ assert shifted{L1, L2}(_arr, _fst+_s1, _last+_s1, _s2) ; @/ \
11  /@ loop invariant _fst <= _i <= _last ; \
12  loop invariant shifted{L1, L2}(_arr+_s1, _fst, _i, _s2) ; \
13  loop assigns _i ; \
14  loop variant _last-_i ; @/ \
15  for(size_t _i = _fst ; _i < _last ; ++_i){ \
16  /@ assert \let _h_i = \at(_i, Here) ; \
17  \at(_arr[_h_i+_s1], L1) == \at(_arr[_h_i+_s1+_s2], L2) ; @/ \
18  } \
19  /@ assert shifted{L1, L2}(_arr+_s1, _fst, _last, _s2) ; @/

```

We will not detail this code which is quite similar to what we have written in the beginning of this section. the only small subtlety is the assert that helps the SMT solvers to relate the memory locations between `L1` and `L2` together on lines 16–17. With this macro, we can see that the assertion at the end of the function `context_to_prove_shift_ptr` is correctly, proved. Thus, we expect the macro to help the provers to get a similar conclusion in a similar context (that is to say, a context where we now that `shifted` holds for some array between two memory labels).

Finally, we can complete the proof of our function `callee` by using our lemma macro:

```

52  /*@
53  requires \valid(array+(0 .. len+s1+s2-1)) ;
54  requires s1+s2 + len <= UINT_MAX ;
55  assigns array[s1 .. s1+s2+len-1];
56  ensures shifted{Pre, Post}(array+s1, 0, len, s2) ;
57  */

```

7. Proof methodologies

```
58 void callee(int* array, size_t len, size_t s1, size_t s2){
59     shift_array(array, s1, s1+len, s2) ;
60     //@ ghost shift_ptr(Pre, Here, array, 0, len, s1, s2) ;
61 }
```

As one could notice, while this technique allows to inject the proof carrying code with a single line of code, it can inject quite a lot of code and annotations each time we use it. Furthermore, once we inject the code in the location where we expect it to be actually useful, the corresponding context can sometimes be already complex. Thus, we could need to slightly modify the code of the macro in order to add more information that is unnecessary in a clean context.

All of this can make the proof context bigger, and harder to use for SMT solvers. There are other limitations to this technique and the careful reader might have notice them. Let us now talk about it.

7.3.4. Limitations

The main limitation of lemma functions and lemma macros is the fact that we are limited to C types. For example, if we compare our lemma `element_level_sorted_is_sorted` with its corresponding lemma function, the original type for the variable `len` is a mathematical integer while in the lemma function its type is `size_t`. It means that while the lemma is true for any integer, and so it could be used no matter if in the program the type of the variable that represents the size is an `int`, or an `unsigned` (or another integer type), on the opposite, our lemma function can be used only if this type can be safely converted to `size_t`. However, this limitation is often not a problem: we just have to express our specification for the biggest type we have to consider in our program and most of the time, it will be enough. And if it is not, we can for example duplicate the lemma for the types we are interested in. Most of the time this limitation is not a big deal since during a verification, we just tend to work with the same types as the program uses.

In some cases, however, it can constrain our modeling of some properties, and it is mainly related to the logic types we can use to model some actual data structures. For example, in order to model a linked list, one could use the ACSL logic type `\list<Type>`, and express an inductive or axiomatic definition in order to define how an actual linked list can be modeled by a logic list, thus we could have some lemmas about logic lists. For example:

```
1  /*@
2    lemma in_list_in_sublist:
3      \foralll \list<int> l, l1, l2, int element ;
4        l == (l1 ^ l2) ==>      // Here, ^ denotes lists concatenation
5          (in_list(element, l) <==> (in_list(element, l1) || in_list(element, l2))) ;
6  */
```

We cannot write lemma functions with proof carrying code for this property as we have no way to use this type in C code, and thus, no way to write a loop and an invariant that would allow us to prove this property.

The other limitation is related to lemma macros and what we already mentioned in the previous part about assertions. By adding too many assertions, the proof context can become too big

7. Proof methodologies

and complex, thus hard to manipulate for SMT solvers. Using lemma macros that can generate quite a lot code and annotations can lead to bigger proof contexts, thus it should be used with care.

Finally, depending on the property to prove, it can be hard to find a proof carrying code. Indeed, proof assistants like Coq have been designed to be able to express proofs even for complex properties, mainly relying on an high level view of our problems, while C has been designed to write programs, and with really detailed low level view of our problems. Thus, it can be sometimes difficult to write a C program to handle some properties and even more to find a suitable invariant for the loops it would involve.

7.3.5. Back to the insertion sort

Now let us go back to our proof of the insertion sort algorithm and see how we can get rid of all our interactive proofs for this function. Note however that in this proof, we often need to deal with macros since the program has not been particularly written with the idea to formally verify it later (for this, the reader can refer to the version proposed in the book ACSL by Example which can be adapted with a similar technique and is easier to prove). Thus, in this example, we push the solvers to their limit because of big proof contexts. With this example, depending on how powerful the machine is, we might need to increase the proof timeout to 120 seconds (which is already quite long for a SMT solver). In this example, we will illustrate three actual usage of ghost code that we have seen so far:

- directly writing code to build a proof,
- writing (and using) lemma functions,
- writing (and using) lemma macros.

We also make use of assertions to make the proof context richer so SMT solvers succeed in proving the properties we are interested in. Some parts of the annotations are equivalent to what we have done previously. First, we use some assertions that were also useful in our previous proof. We will recall their purpose for each function. Second, we re-use the same axiomatic definition for occurrences counting. Furthermore, we keep the following predicate definitions:

```
25 /*@
26 predicate sorted(int* a, integer b, integer e) =
27   \forall integer i, j; b <= i <= j < e ==> a[i] <= a[j];
28
29 predicate shifted{L1, L2}(integer s, int* a, integer beg, integer end) =
30   \forall integer k ; beg <= k < end ==> \at(a[k], L1) == \at(a[s+k], L2) ;
31
32 predicate unchanged{L1, L2}(int* a, integer beg, integer end) =
33   shifted{L1, L2}(0, a, beg, end);
34
35 predicate rotate_left{L1, L2}(int* a, integer beg, integer end) =
36   beg < end && \at(a[beg], L2) == \at(a[end-1], L1) &&
37   shifted{L1, L2}(1, a, beg, end - 1) ;
38
39 predicate permutation{L1, L2}(int* in, integer from, integer to) =
40   \forall int v ; l_occurrences_of{L1}(v, in, from, to) ==
41     l_occurrences_of{L2}(v, in, from, to) ;
42 */
```

7. Proof methodologies

As we will need them, as well as the lemma about the transitivity of occurrences counting as it is automatically proved by SMT solvers (thus we can keep it since it does not require an interactive proof from us):

```
43 /*@ lemma transitive_permutation{L1, L2, L3}:
44   \forall int* a, integer beg, integer end ;
45     permutation{L1, L2}(a, beg, end) ==>
46     permutation{L2, L3}(a, beg, end) ==>
47     permutation{L1, L3}(a, beg, end) ;
48 */
```

Let us start with the `insertion_sort` function itself. In this function, we made use of three assertions:

```
87 /*@
88   requires beg < end && \valid(a + (beg .. end-1));
89   assigns a[beg .. end-1];
90   ensures sorted(a, beg, end);
91   ensures permutation{Pre, Post}(a,beg,end);
92 */
93 void insertion_sort(int* a, size_t beg, size_t end){
94   /*@
95     loop invariant beg+1 <= i <= end ;
96     loop invariant sorted(a, beg, i) ;
97     loop invariant permutation{Pre, Here}(a,beg,end);
98     loop assigns a[beg .. end-1], i ;
99     loop variant end-i ;
100   */
101   for(size_t i = beg+1; i < end; ++i) {
102     //@ ghost L:
103     insert(a, beg, i);
104     //@ assert permutation{L, Here}(a, beg, i+1);
105     //@ assert unchanged{L, Here}(a, i+1, end) ;
106     //@ assert permutation{L, Here}(a, beg, end) ;
107   }
108 }
```

The first one makes sure that the part of the array where we just inserted a value is a permutation of the same range of values before the call to `insert`, as this is the postcondition of the function, it is not necessary but let us keep it for illustration. The last assertion is the property we want to prove in order to get enough knowledge to use the lemma that states that permutation is transitive (and show that after the block of the loop since our array is a permutation of the array at the beginning, which is itself a permutation of the original one, then after the body of the loop we have maintained that the array is a permutation of the original one).

The second assertion says that the second part of the array is unchanged, and we want to use this knowledge to show that the number of occurrences of the values is unchanged. Here we could use a combination of lemma functions and macros to prove that the complete range is a permutation (as we will do for the other function) however, directly writing the code is here a bit simpler (requires less proofs, as we will see later) so let us directly write the code that will create the proof of our property.

In order to show that the complete range is a permutation, we have to show that the number of occurrences did not change. We know that the first part of the array is a permutation of the same range at the beginning of the body of the loop. Thus, we already know that the number

7. Proof methodologies

of occurrences of any v did not change for a part of our array. Let us continue the occurrences counting for the rest of our array, knowing that the second part is unchanged (when `i+1` is lower than `end` as else, we do not have anything to count):

```

267 for(size_t i = beg+1; i < end; ++i) {
268     //@ ghost L:
269     insert(a, beg, i);
270     //@ assert permutation{L, Here}(a, beg, i+1);
271     //@ assert unchanged{L, Here}(a, i+1, end) ;
272
273     /*@ ghost
274     if(i+1 < end){
275         //@ loop invariant i+1 <= j <= end ;
276         loop invariant \forall int v ;
277             l_occurrences_of{L}(v, a, beg, \at(j, Here)) ==
278             l_occurrences_of(v, a, beg, j) ;
279         loop assigns j ;
280         loop variant end - j ;
281     }
282     for(size_t j = i+1 ; j < end ; ++j);
283 }
284 */
285 //@ assert permutation{L, Here}(a, beg, end) ;
286 }

```

which is enough to ensure that the `insertion_sort` function respects its specification as long as we finish the proof of the function `insert`. This second function makes a more complex action. Let us depart from this annotated version:

```

50 /*@
51 requires beg < last < UINT_MAX && \valid(a + (beg .. last));
52 requires sorted(a, beg, last) ;
53
54 assigns a[ beg .. last ] ;
55
56 ensures permutation{Pre, Post}(a, beg, last+1);
57 ensures sorted(a, beg, last+1) ;
58 */
59 void insert(int* a, size_t beg, size_t last){
60     size_t i = last ;
61     int value = a[i] ;
62
63     /*@
64     loop invariant beg <= i <= last ;
65     loop invariant \forall integer k ; i <= k < last ==> a[k] > value ;
66     loop invariant \forall integer k ; beg <= k <= i ==> a[k] == \at(a[k], Pre) ;
67     loop invariant \forall integer k ; i+1 <= k <= last ==> a[k] == \at(a[k-1], Pre) ;
68
69     loop assigns i, a[beg .. last] ;
70     loop variant i ;
71 */
72 while(i > beg && a[i - 1] > value){
73     a[i] = a[i - 1] ;
74     --i ;
75 }
76 a[i] = value ;
77 //@ assert sorted(a, beg, last+1) ;
78
79 //@ assert rotate_left{Pre, Here}(a, i, last+1) ;
80 //@ assert permutation{Pre, Here}(a, i, last+1) ;
81
82 //@ assert unchanged{Pre, Here}(a, beg, i) ;
83 //@ assert permutation{Pre, Here}(a, beg, i) ;

```

7. Proof methodologies

```
84 }
```

Again, the proof that this function maintains a permutation of the array is the hardest part of the job. The fact that the function guarantees that the value are sorted is already easily established. Using the same technique as for `insertion_sort` is not so easy here. Indeed, the second part of the array has been rotated which makes the properties slightly more complex. So, let us show that we can split the array at the position where we insert into two parts, in which we respectively show that:

- for the first part, since it is unchanged, for any v , the number of occurrences did not change either,
- for the second part, since it is rotated, for any v , the number of occurrences did not change.

First let us define a lemma function that allows to explicitly split a range of values into two subparts in which we can count separately.

```
57 /*@
58   requires beg <= split <= end ;
59
60   assigns \nothing ;
61
62   ensures \forall int v ;
63     l_occurrences_of(v, a, beg, end) ==
64     l_occurrences_of(v, a, beg, split) + l_occurrences_of(v, a, split, end) ;
65 */
66 void l_occurrences_of_explicit_split(int* a, size_t beg, size_t split, size_t end){
67   /*@
68     loop invariant split <= i <= end ;
69     loop invariant \forall int v ; l_occurrences_of(v, a, beg, i) ==
70       l_occurrences_of(v, a, beg, split) + l_occurrences_of(v, a, split, i) ;
71     loop assigns i ;
72     loop variant end - i ;
73   */
74   for(size_t i = split ; i < end ; ++i);
75 }
```

We can note that this property is proved in a way that is quite similar to what we wrote for the body of our loop in `insertion_sort`, we start from the point where we want to count and show that the property remains true until the end of the array.

We can use our function to split the array at the right place after the loop. However, we can only do it for the new content of the array, indeed, in order to establish it for the original array, we have to call the function on the original array, when we still do not know the value of i . Thus, let us write another version of the “split” property that shows that we can split the array at any index, thus make the `split` variable a universally quantified variable, and use the previous function to prove that it is true:

```
77 /*@
78   requires beg <= end ;
79
80   assigns \nothing ;
81
```

7. Proof methodologies

```

82  ensures \forall int v, size_t split ; beg <= split <= end ==>
83      l_occurrences_of(v, a, beg, end) ==
84      l_occurrences_of(v, a, beg, split) + l_occurrences_of(v, a, split, end) ;
85  */
86  void l_occurrences_of_split(int* a, size_t beg, size_t end){
87      /*@
88      loop invariant beg <= i <= end ;
89      loop invariant \forall int v, size_t split ; beg <= split < i ==>
90          l_occurrences_of(v, a, beg, end) ==
91          l_occurrences_of(v, a, beg, split) + l_occurrences_of(v, a, split, end) ;
92      loop assigns i ;
93      loop variant end - i ;
94      */
95      for(size_t i = beg ; i < end ; ++i){
96          //@ ghost l_occurrences_of_explicit_split(a, beg, i, end);
97      }
98  }

```

And we can split our original array and the new one:

```

1  void insert(int* a, size_t beg, size_t last){
2      size_t i = last ;
3      int value = a[i] ;
4
5      // split before modifying
6      //@ ghost l_occurrences_of_split(a, beg, last+1);
7
8      /*@ LOOP ANNOT */
9      while(i > beg && a[i - 1] > value){
10         a[i] = a[i - 1] ;
11         --i ;
12     }
13     a[i] = value ;
14     // Assertions ...
15
16     // split after modifying, now we know "i"
17     //@ ghost l_occurrences_of_explicit_split(a, beg, i, last+1);
18 }

```

Now, the only remaining parts of the proof are first to show that an unchanged array is a permutation and second that the rotate operation also maintains a permutation. Here, we need macros. Let us start with the easiest: the unchanged property that we already almost exactly proved in the `insertion_sort` function. We start by building the context for our proof:

```

138 /*@
139   assigns arr[fst .. last-1] ;
140   ensures unchanged{Pre, Post}(arr, fst, last);
141 */
142 void unchanged_permutation_premise(int* arr, size_t fst, size_t last);
143
144 /*@
145   requires fst <= last ;
146 */
147 void context_to_prove_unchanged_permutation(int* arr, size_t fst, size_t last){
148     L1: ;
149     unchanged_permutation_premise(arr, fst, last);
150     L2: ;
151     //@ ghost unchanged_permutation(L1, L2, arr, fst, last) ;
152
153     //@ assert permutation{L1, L2}(arr, fst, last) ;

```

7. Proof methodologies

```
154 }
```

The function `unchanged_permutation_premise` ensures that we have modified the array (thus created a new memory state) and that the array is unchanged from the precondition to the postcondition. We can build our lemma macro:

```
125 #define unchanged_permutation(_L1, _L2, _arr, _fst, _last) \
126   /* assert unchanged{_L1, _L2}(_arr, _fst, _last) ; */ \
127   /* loop invariant _fst <= _i <= _last ; \
128     loop invariant \forall int _v ; \
129       l_occurrences_of{_L1}(_v, _arr, _fst, \at(_i, Here)) == \
130       l_occurrences_of{_L2}(_v, _arr, _fst, \at(_i, Here)) ; \
131     loop assigns _i ; \
132     loop variant _last - _i ; \
133   */ \
134   for(size_t _i = _fst ; _i < _last ; ++_i) ; \
135   /* assert permutation{_L1, _L2}(_arr, _fst, _last); */
```

Which almost corresponds to what we have written previously in the `insert_sort` function and we can use the macro where it is needed in the `insert` function.

```
245 /*@ assert unchanged{Pre, Here}(a, beg, i) ;
246 /*@ ghost unchanged_permutation(Pre, Here, a, beg, i) ;
247 /*@ assert permutation{Pre, Here}(a, beg, i) ;
```

The only remaining property to prove is the hardest one and is about the `rotate_left` predicate. Let us first write our context to prepare the macro.

```
183 /*@
184   assigns arr[fst .. last-1] ;
185   ensures rotate_left{Pre, Post}(arr, fst, last);
186 */
187 void rotate_left_permutation_premise(int* arr, size_t fst, size_t last);
188
189 /*@
190   requires fst < last ;
191 */
192 void context_to_prove_rotate_left_permutation(int* arr, size_t fst, size_t last){
193   L1: ;
194   /*@ ghost l_occurrences_of_explicit_split(arr, fst, last-1, last) ;
195     rotate_left_permutation_premise(arr, fst, last);
196   L2: ;
197   /*@ ghost rotate_left_permutation(L1, L2, arr, fst, last) ;
198
199   /*@ assert permutation{L1, L2}(arr, fst, last) ;
200 }
```

How can we prove this property? Basically, one has to notice that since all the elements from the beginning to the penultimate have been shifted of one index to the right the number of occurrences in the shifted part did not change. Then one has to show that the number of occurrences of any v in respectively the last cell in the original array, and the first cell in the new array is the same (since the corresponding element is the same). Again, we rely on the split function to count separately the elements that are shifted and the one which is moved

7. Proof methodologies

from the end to the beginning. However, the call corresponding to the original array has again to be put before the code that modifies the memory (see line 194) in the previous code, and we will have to take that in account when we will insert our use of the macro in the `insert` function.

Let us now present the macro that we use to prove that the lemma holds:

```

156 #define rotate_left_permutation(_L1, _L2, _arr, _fst, _last)      \
157   /* assert rotate_left{_L1, _L2}(_arr, _fst, _last) ; */        \
158   /* loop invariant _fst+1 <= _i <= _last ;                      \
159     loop invariant \forall int _v ;                                \
160       l_occurrences_of{_L1}(_v, _arr, _fst, \at(_i-1, Here)) ==  \
161       l_occurrences_of{_L2}(_v, _arr, _fst+1, \at(_i, Here)) ;    \
162     loop assigns _i ;                                             \
163     loop variant _last - _i ;                                     \
164   */                                                              \
165   for(size_t _i = _fst+1 ; _i < _last ; ++_i) {                 \
166     /* assert \at(_arr[\at(_i-1, Here)], _L1) ==                 \
167        \at(_arr[\at(_i, Here)], _L2) ;                          \
168     */                                                            \
169   }                                                                \
170   l_occurrences_of_explicit_split(_arr, _fst, _fst+1, _last) ;   \
171   /* assert \forall int _v ;                                       \
172     l_occurrences_of{_L1}(_v, _arr, _fst, _last) ==             \
173     l_occurrences_of{_L1}(_v, _arr, _fst, _last-1) +            \
174     l_occurrences_of{_L1}(_v, _arr, _last-1, _last) ;          \
175   */                                                              \
176   /* assert \at(_arr[_fst], _L2) == \at(_arr[_last-1], _L1) ==> \
177     (\forall int _v ;                                             \
178       l_occurrences_of{_L2}(_v, _arr, _fst, _fst+1) ==         \
179       l_occurrences_of{_L1}(_v, _arr, _last-1, _last)) ;        \
180   */                                                              \
181   /* assert permutation{_L1, _L2}(_arr, _fst, _last); */ @/

```

The loop invariant is pretty similar to what we have written so far, the only difference is that it takes in account the shift of the elements. Furthermore, to prove the invariant we had to add an assertion to help the solvers notice that the last element of both ranges is the same (note however that depending on the versions of the solvers or how powerful is the machine, this might be unneeded sometimes). A more important difference compared to our previous examples is that here, we need to provide more information to SMT solvers by adding other ghost functions calls (line 170, in order to split the first element of the array), as well as assertions to guide the last steps of the proof:

- 171–175: we recall that in the original array we can split the last element,
- 176–180: we show that as the first element of the array is the last element of the original array (176), the number of occurrences for any value in these ranges is the same (177–179).

We can use the macro in our program:

```

1  /*@ assert rotate_left{Pre, Here}(a, i, last+1) ;
2  /*@ ghost rotate_left_permutation(Pre, Here, a, i, last+1) ;
3  /*@ assert permutation{Pre, Here}(a, i, last+1) ;

```

However, we have to show that the considered range at label `Pre` can be split at `last`. For this, we use another variant of the split function, that shows that any subrange can be split before the last element (if it is non empty):

7. Proof methodologies

```

100  /*@
101    requires beg < end ;
102
103    assigns \nothing ;
104
105    ensures \forall int v, size_t any ; beg <= any < end ==>
106      l_occurrences_of(v, a, any, end) ==
107      l_occurrences_of(v, a, any, end-1) + l_occurrences_of(v, a, end-1, end) ;
108  */
109  void l_occurrences_of_from_any_split_last(int* a, size_t beg, size_t end){
110    /*@
111      loop invariant beg <= i <= end-1 ;
112      loop invariant \forall int v, size_t j ;
113        beg <= j < i ==>
114          l_occurrences_of(v, a, j, end) ==
115          l_occurrences_of(v, a, j, end-1) + l_occurrences_of(v, a, end-1, end) ;
116      loop assigns i ;
117      loop variant (end - 1) - i ;
118    */
119    for(size_t i = beg ; i < end-1 ; ++i){
120      //@ ghost l_occurrences_of_explicit_split(a, i, end-1, end);
121    }
122  }

```

That we have to call before the loop in the `insert` function:

```

211  void insert(int* a, size_t beg, size_t last){
212    size_t i = last ;
213    int value = a[i] ;
214
215    //@ ghost l_occurrences_of_split(a, beg, last+1);
216    //@ ghost l_occurrences_of_from_any_split_last(a, beg, last+1);

```

Note that depending on the version of the solvers, the assertion on lines 176 to 180 of the macro, about the element at the beginning/the end of the array might fail due to the complexity of the proof context. Let us help the solvers a last time by adding a last lemma (automatically proved by SMT solvers) that states this relation for any array and position in the array:

```

49  /*@ lemma one_same_element_same_count{L1, L2}:
50    \forall int* a, int* b, int v, integer pos_a, pos_b ;
51    \at(a[pos_a], L1) == \at(b[pos_b], L2) ==>
52      l_occurrences_of{L1}(v, a, pos_a, pos_a+1) ==
53      l_occurrences_of{L2}(v, b, pos_b, pos_b+1) ;
54  */

```

which guarantees that our resulting annotated insertion function is entirely proved:

```

202  /*@
203    requires beg < last < UINT_MAX && \valid(a + (beg .. last));
204    requires sorted(a, beg, last) ;
205
206    assigns a[ beg .. last ] ;
207
208    ensures permutation{Pre, Post}(a, beg, last+1);
209    ensures sorted(a, beg, last+1) ;
210  */

```

7. Proof methodologies

```
211 void insert(int* a, size_t beg, size_t last){
212     size_t i = last ;
213     int value = a[i] ;
214
215     //@ ghost l_occurrences_of_split(a, beg, last+1);
216     //@ ghost l_occurrences_of_from_any_split_last(a, beg, last+1);
217
218     /*@
219     loop invariant beg <= i <= last ;
220     loop invariant \forall integer k ; i <= k < last ==> a[k] > value ;
221     loop invariant \forall integer k ; beg <= k <= i ==> a[k] == \at(a[k], Pre) ;
222     loop invariant \forall integer k ; i+1 <= k <= last ==> a[k] == \at(a[k-1], Pre) ;
223
224     loop assigns i, a[beg .. last] ;
225     loop variant i ;
226     */
227     while(i > beg && a[i - 1] > value){
228         a[i] = a[i - 1] ;
229         --i ;
230     }
231     a[i] = value ;
232     //@ assert sorted(a, beg, last+1) ;
233
234     /*@ assert
235     \forall int v ;
236     l_occurrences_of{Pre}(v, a, \at(i, Here), last+1) ==
237     l_occurrences_of{Pre}(v, a, \at(i, Here), last) +
238     l_occurrences_of{Pre}(v, a, last, last +1);
239     */
240
241     //@ assert rotate_left{Pre, Here}(a, i, last+1) ;
242     //@ ghost rotate_left_permutation(Pre, Here, a, i, last+1) ;
243     //@ assert permutation{Pre, Here}(a, i, last+1) ;
244
245     //@ assert unchanged{Pre, Here}(a, beg, i) ;
246     //@ ghost unchanged_permutation(Pre, Here, a, beg, i) ;
247     //@ assert permutation{Pre, Here}(a, beg, i) ;
248
249     //@ ghost l_occurrences_of_explicit_split(a, beg, i, last+1);
250 }
```

We finally highlight how the proof context can make the proof harder for SMT solvers. Basically, if we swap the proofs for each part of the array, that is, starting with the “unchanged” part and after the “rotate” part, the proof has more chance to fail, since it would make the proof context bigger for the hardest proof. For the same reason, it is here almost needed (again depending on how powerful is the machine and the solvers) to separate the proof of the absence of runtime errors, else it pollutes the proof context and it may fail.

7.3.6. Exercises

7.3.6.1. Sum of N integers

Using lemma functions, we can prove the lemma about the sum of the N first integers that we previously expressed. You might have done this proof when you were in high school, now it is time to do it in C and ACSL. Write a contract for the following function that expresses in postcondition that the sum of the N first integers is $N(N+1)/2$. Complete the body of the function with a loop in order to prove this property. We advise to slightly transform the invariant in order to be sure that the property does not contain any division (division on integers have some properties that can make them hard to deal with for SMT solvers depending on the constraint that exists on the used values).

7. Proof methodologies

```
1  /*@
2   logic integer sum_of_n_integers(integer n) =
3   (n <= 0) ? 0 : sum_of_n_integers(n-1) + n ;
4  */
5
6  /*@
7   assigns \nothing ;
8   ensures ... ;
9  */
10 void lemma_value_of_sum_of_n_integers_2(unsigned n){
11     // ...
12 }
```

Now, let us generalize to any bounds with the sum of all the integers between `fst` and `lst`. We provide the logic function and the contract. Write a body for the function such that the postcondition is proved. Note that again, we advise to express the invariant without division.

```
14 /*@
15 logic integer sum_of_range_of_integers(integer fst, integer lst) =
16 ((lst <= fst) ? lst : lst + sum_of_range_of_integers(fst, lst-1)) ;
17 */
18
19 /*@
20 requires fst <= lst ;
21 assigns \nothing ;
22 ensures ((lst-fst+1)*(fst+lst))/2 == sum_of_range_of_integers(fst, lst) ;
23 */
24 void lemma_value_of_sum_of_range_of_integers(int fst, int lst){
25     // ...
26 }
```

7.3.6.2. Properties about occurrences counting

In this exercise, we want to prove some interesting properties about the axiomatically defined logic function `l_occurrences_of` :

```
1  #include <stddef.h>
2
3  void occ_bounds(int v, int* arr, size_t len){
4      // ...
5  }
6
7  void not_in_occ_0(int v, int* arr, size_t len){
8      // ...
9  }
10
11 void occ_monotonic(int v, int* arr, size_t pos, size_t more){
12     // ...
13 }
14
15 void occ_0_not_in(int v, int* arr, size_t len){
16     // ...
17     // needs occ_monotonic
18 }
19
20 size_t occ_pos_find(int v, int* arr, size_t len){
21     // ...
22 }
```

7. Proof methodologies

```
22 // needs occ_monotonic
23 }
24
25 void occ_pos_exists(int v, int* arr, size_t len){
26 // ...
27 // should use occ_pos_find
28 }
```

The function `occ_bounds` should state that the number of occurrences of `v` in the array is comprised between 0 and `len`.

The function `not_in_occ_0` should state that if `v` is not in the array then the number of occurrences of `v` in the array is 0.

The function `occ_monotonic` should state that the number of occurrences of `v` in the array from 0 to `pos` is lower or equals to the number of occurrences of `v` in the array from 0 to `more`, if `more` is greater or equals to `pos`.

The function `occ_0_not_in` should state that if the number of occurrences of `v` in the array is 0 then `v` is not in the array. Note that you will probably need to use `occ_monotonic`.

The function `occ_pos_find` should find an index `i` such that the value `arr[i]` is `v`, provided that the number of occurrences of `v` is positive. Note that you will probably need to use `occ_monotonic`.

Finally the function `occ_pos_exists` should translate the contract of the previous function using an existentially quantified variable, and use the previous function to obtain the proof for free.

For all these functions, WP should be parameterized with the control of the absence of runtime errors as well as the option `-warn-unsigned-overflow`.

7.3.6.3. An actual example with sum

Take back the proof performed in the previous chapter for the exercise 7.2.5.4. Modify the annotations in order to ensure that no more classic lemmas are necessary. The skeleton of the file follows:

```
1 #include <limits.h>
2 #include <stddef.h>
3
4 /*@
5   axiomatic Sum_array{
6     logic integer sum(int* array, integer begin, integer end) reads array[begin .. (end-1)];
7
8     axiom empty:
9       \forall int* a, integer b, e; b >= e ==> sum(a,b,e) == 0;
10    axiom range:
11      \forall int* a, integer b, e; b < e ==> sum(a,b,e) == sum(a,b,e-1)+a[e-1];
12  }
13 */
14
15 /*@
16   predicate unchanged{L1, L2}(int* array, integer begin, integer end) =
17     \forall integer i ; begin <= i < end ==> \at(array[i], L1) == \at(array[i], L2) ;
```

7. Proof methodologies

```
18 */
19
20 void sum_separable(int* array, size_t begin, size_t split, size_t end){
21     // ...
22 }
23
24 #define unchanged_sum(_L1, _L2, _arr, _beg, _end) ;
25
26
27 /*@
28   requires i < len ;
29   requires array[i] < INT_MAX ;
30   requires \valid(array + (0 .. len-1));
31   assigns array[i];
32   ensures sum(array, 0, len) == sum{Pre}(array, 0, len)+1;
33 */
34 void inc_cell(int* array, size_t len, size_t i){
35     // ...
36     array[i]++;
37     // ...
38 }
```

7. Proof methodologies

As we try to prove more complex properties, particularly when programs involve loops, there is a part of “trial and error” in order to understand what the provers miss to establish the proof.

It can miss hypotheses. In this case, we can try to add assertions to guide the prover, or write ghost code with the right invariant that allows to make a part of the reasoning by ourselves when it is too hard for SMT solvers.

With some experience, we can read the content of the proof obligations or try to perform the proof with the Coq interactive prover to see whether the proof seems to be possible. Sometimes, the prover just needs more time, in such a case, we can (sometimes a lot) augment the timeout value. Of course, the property can be too hard for the prover and ghost code might be sometimes unsuitable, and in this case, we have to write the proof ourselves with an interactive prover.

Finally, the implementation can be indeed incorrect, and in this case we have to fix it. Here, we use test and not proof, because a test allows us to prove the presence of a bug and to analyze this bug.

8. Conclusion

Voilà, c'est fini ...

Jean-Louis Aubert, Bleu Blanc Vert, 1989

... for this introduction to the proof of C programs using Frama-C and WP.

Along this tutorial, we have seen how we can use these tools to specify what we expect of our programs and verify that the source code we have produced indeed corresponds to the specification we have provided. This specification is provided using annotations of our functions that includes the contract they must respect. These contracts are properties required about the input to ensure that the function will correctly work, which is specified by properties about the output of the function and enforced by the tool that allow us to check specific problems related to the use of C (namely, the absence of runtime errors).

Starting from specified programs, WP is able to produce the weakest precondition of our functions, provided what we want in postcondition, and to ask some provers if the specified precondition is compatible with the computed one. The reasoning is completely modular, which allows to prove functions in isolation from each other and to compose the results.

WP cannot currently work with dynamic allocation. A function that would use it could not be proved. However, even without dynamic allocation, a lot of function can be proved since they work with data-structures that are already allocated. And these functions can then be called with the certainty that they perform a correct job. If we cannot or do not want to prove the client code of a function, we can still write something like this:

```
1  /*@
2   requires some_properties_on(a);
3   requires some_other_on(b);
4
5   assigns ...
6   ensures ...
7  */
8  void my_function(int* a, int b){
9      //this is indeed the "assert" defined in "assert.h"
10     assert(/*properties on a*/ && "must respect properties on a");
11     assert(/*properties on b*/ && "must respect properties on b");
12 }
```

Which allows us to benefit from the robustness of our function having the possibility to debug an incorrect call in a source code that we cannot or do not want to prove.

Writing specifications is sometimes long or tedious. Higher-level features of ACSL (predicates, logic functions, axiomatizations) allow us to lighten this work, as well as our programming languages allow us to define types containing other types, functions calling functions, bringing

8. Conclusion

us to the final program. But, despite this, write specification in a formal language, no matter which one, is generally a hard task.

However, this **formalization** of our need is **crucial**. Concretely, such a formalization is a work every developer should do. And not only in order to prove a program. Even the definition of tests for a function requires to correctly understand its goal if we want to test what is necessary and only what is necessary. And writing specification in a formal language is incredibly useful (even if it can be sometimes frustrating) to get a clear specification.

Formal languages, that are close to mathematics, are precise. Mathematics have this: they are precise. What is more terrible than reading a specification written in a natural language, with complex sentences, using conditional forms, imprecise terms, ambiguities, compiled in administrative documents composed of hundreds of pages, and where we need to determine, “so, what this function is supposed to do? And what do I have to validate about it?”.

Formal methods are probably not used enough currently. Sometimes because of mistrust, sometimes because of ignorance, sometimes because of prejudice based on ideas born at the beginning of the tools, 30 years ago. Our tools evolve, the ways we develop change, probably faster than in any other technical domain. Saying that these tools could never be used for real life programs would certainly be a too big shortcut. After all, we see everyday how much development is different from what it were 10 years, 20 years, 40 years ago and can barely imagine how much it will be different in 10 years, 20 years, 40 years.

During the past few years, safety and security questions have become more and more visible and crucial. Formal methods also progressed a lot, and the improvement they bring for these questions are greatly appreciated. For example, this [this link ↗](#) brings to the report of a conference about security that brought together people from academic and industrial world, in which we can read:

Adoption of formal methods in various areas (including verification of hardware and embedded systems, and analysis and testing of software) has dramatically improved the quality of computer systems. We anticipate that formal methods can provide similar improvement in the security of computer systems.

...

Without broad use of formal methods, security will always remain fragile.

Formal Methods for Security, 2016

8.1. Going further

8.1.1. With Frama-C

Frama-C provides different ways to analyse programs. Among these tools, the most commonly used and interesting to know from a static and dynamic verification point of view are certainly those ones:

- abstract interpretation analysis using [EVA ↗](#) ,
- the transformation of annotation into runtime verification using [E-ACSL ↗](#) .

8. Conclusion

The goal of the first one is to compute the domain of the different variables at each program point. When we precisely know these domains, we can determine if these variables can produce errors when they are used. However this analysis is executed on the whole program and not modularly. It is also strongly dependent of the type of domain we use (we will not enter into details here) and it is not so good at keeping the relations between variables. On the other side, it is really completely automatic, we do not even need to give loop invariant! The most manual part of the work is to determine whether or not an alarm is a true error or a false positive.

The second analysis allows to generate from an original program, a new program where the assertions are transformed into runtime verifications. If these assertions fail, the program fails. If they are valid, the program has the same behavior it would have without the assertions. An example of use is to generate the verification of absence of runtime errors as assertions using `-rte` and then to use E-ACSL to generate the program containing the runtime verification that these assertions do not fail.

There exist a lot of different plugins for very different tasks in Frama-C.

Finally, a last possibility that will motivate the use of Frama-C is the ability to develop their own analysis plugins using the API provided by the Frama-C kernel. A lot of tasks can be realized by the analysis of the source code and Frama-C allows to build them as easily as possible.

8.1.2. With deductive proof

Along this tutorial we used WP to generate proof obligation starting from programs with their specification. Next we have used automatic solvers to assure that these properties were indeed verified.

When we use other solvers than Alt-Ergo and Coq, the communication with this solver is provided by a translation to the Why3 language that will next be used to bridge the gap to automatic solvers. But this is not the only way to use Why3. It can also be used itself to write programs and prove them. It especially provides a set of theories for some common data structures.

There are some proofs that cannot be discharged by automatic solvers. In such a case, we have to provide these proofs interactively. WP, like Why3, can extract its verification conditions to Coq, and it is very interesting to study this language. In the context of Frama-C, we can produce reusable lemmas libraries proved with Coq. But Coq can also be used for many different tasks, including programming. Note that Why3 can also extract its proof obligations to Isabelle or PVS that are also proof assistants.

Finally, there exists other program logics, for example separation logic or concurrent program logics. Again these notions are interesting to know in the context of Frama-C: if we cannot directly use them, they can inspire the way we specify our program in Frama-C for the proof with WP. They could also be implemented into new plugins to Frama-C.

A whole new world of methods to explore.