

ACSL: ANSI/ISO C Specification Language

Version 1.5 – Carbon-20101202-beta2

Patrick Baudin¹, Pascal Cuoq¹, Jean-Christophe Filliâtre^{4,3}, Claude Marché^{3,4},
Benjamin Monate¹, Yannick Moy^{2,4,3}, Virgile Prevosto¹

¹ CEA LIST, Software Reliability Laboratory, Saclay, F-91191

² France Télécom, Lannion, F-22307

³ INRIA Saclay - Île-de-France, ProVal, Orsay, F-91893

⁴ LRI, Univ Paris-Sud, CNRS, Orsay, F-91405

©2009-2010 CEA LIST and INRIA

This work has been supported by the ‘CAT’ ANR project (ANR-05-RNTL-0030x) and by the ANR CIFRE contract 2005/973.



Contents

1	Introduction	11
1.1	Organization of this document	11
1.2	Generalities about Annotations	12
1.2.1	Kinds of annotations	12
1.2.2	Parsing annotations in practice	13
1.2.3	About preprocessing	13
1.2.4	About keywords	13
1.3	Notations for grammars	13
2	Specification language	15
2.1	Lexical rules	15
2.2	Logic expressions	16
2.2.1	Operators precedence	18
2.2.2	Semantics	19
2.2.3	Typing	20
2.2.4	Integer arithmetic and machine integers	20
2.2.5	Real numbers and floating point numbers	23
2.2.6	C arrays and pointers	25
2.2.7	Structures, Unions and Arrays in logic	26
2.2.8	String literals	28
2.3	Function contracts	28
2.3.1	Built-in constructs <code>\old</code> and <code>\result</code>	29
2.3.2	Simple function contracts	29
2.3.3	Contracts with named behaviors	30
2.3.4	Memory locations and sets of terms	33
2.3.5	Default contracts, multiple contracts	34
2.4	Statement annotations	35
2.4.1	Assertions	35
2.4.2	Loop annotations	35

CONTENTS

2.4.3	Built-in construct <code>\at</code>	40
2.4.4	Statement contracts	41
2.5	Termination	42
2.5.1	Integer measures	42
2.5.2	General measures	42
2.5.3	Recursive function calls	43
2.5.4	Non-terminating functions	44
2.6	Logic specifications	45
2.6.1	Predicate and function definitions	45
2.6.2	Lemmas	46
2.6.3	Inductive predicates	46
2.6.4	Axiomatic definitions	47
2.6.5	Polymorphic logic types	48
2.6.6	Recursive logic definitions	48
2.6.7	Higher-order logic constructions	49
2.6.8	Concrete logic types	50
2.6.9	Hybrid functions and predicates	51
2.6.10	Memory footprint specification: <code>reads clause</code>	52
2.6.11	Specification Modules	53
2.7	Pointers and physical addressing	54
2.7.1	Memory blocks and pointer dereferencing	54
2.7.2	Separation	54
2.7.3	Allocation and deallocation	54
2.8	Sets as first-class values	55
2.9	Abrupt termination	55
2.10	Dependencies information	57
2.11	Data invariants	58
2.11.1	Semantics	59
2.11.2	Model variables and model fields	61
2.12	Ghost variables and statements	62
2.12.1	Volatile variables	65
2.13	Undefined values, dangling pointers	66
2.13.1	Initialization	66
2.13.2	Unspecified values	67

CONTENTS

3 Libraries	69
3.1 Libraries of logic specifications	69
3.1.1 Real numbers	69
3.1.2 Finite lists	69
3.1.3 Sets and Maps	69
3.2 Jessie library: logical addressing of memory blocks	70
3.2.1 Abstract level of pointer validity	70
3.2.2 Strings	70
3.3 Memory leaks	71
4 Conclusion	73
A Appendices	75
A.1 Glossary	75
A.2 Comparison with JML	75
A.2.1 Low-level language vs. inheritance-based one	76
A.2.2 Deductive verification vs. RAC	79
A.2.3 Syntactic differences	79
A.3 Typing rules	80
A.3.1 Rules for terms	80
A.3.2 Typing rules for sets	81
A.4 Specification Templates	81
A.4.1 Accessing a C variable that is masked	81
A.5 Illustrative example	82
A.6 Changes	87
A.6.1 Version 1.5	87
A.6.2 Version 1.4	88
A.6.3 Version 1.3	88
A.6.4 Version 1.2	88
Bibliography	89
List of Figures	91
Index	93



Foreword

This is the version 1.5 of ACSL design. Several features may still evolve in the future. In particular, some features in this document are considered *experimental*, meaning that their syntax and semantics is not already fixed. These features are marked with EXPERIMENTAL. They must also be considered as advanced features, which are not supposed to be useful for a basic use of that specification language.

Acknowledgements

We gratefully thank all the people who contributed to this document: Sylvie Boldo, Jean-Louis Colaço, Pierre Crégut, David Delmas, Stéphane Duprat, Arnaud Gotlieb, Thierry Hubert, Dillon Pariente, Pierre Rousseau, Julien Signoles, Jean Souyris.



Chapter 1

Introduction

This document is a reference manual for the ACSL implementation provided by the Frama-C framework [11]. ACSL is an acronym for “ANSI C Specification Language”. This is a Behavioral Interface Specification Language (BISL) implemented in the FRAMA-C framework. It aims at specifying behavioral properties of C source code. The main inspiration for this language comes from the specification language of the CADUCEUS tool [9, 10] for deductive verification of behavioral properties of C programs. The specification language of Caduceus is itself inspired from the *Java Modeling Language* (JML [19]) which aims at similar goals for Java source code: indeed it aims both at *runtime assertion checking* and *static verification* using the ESC/JAVA2 tool [15], where we aim at *static verification* and *deductive verification* (see Appendix A.2 for a detailed comparison between ACSL and JML).

Going back further in history, JML design was guided by the general *design-by-contract* principle proposed by Bertrand Meyer, who took his own inspiration from the concepts of preconditions and postconditions on a routine, going back at least to Dijkstra, Floyd and Hoare in the late 60’s and early 70’s, and originally implemented in the EIFFEL language.

In this document, we assume that the reader has a good knowledge of the ANSI C programming language [14, 13].

1.1 Organization of this document

In this preliminary chapter we introduce some definitions and vocabulary, and discuss generalities about this specification language. Chapter 2 presents the specification language itself. Chapter 3 presents additional informations about *libraries* of specifications. Finally, Appendix A provides a few additional information. A detailed table of contents is given on page 5. A glossary is given in Appendix A.1.

Not all of the features mentioned in this document are currently implemented in the Frama-C kernel. Those who aren’t yet are signaled as in the following line:

This feature is not currently supported by Frama-C¹

As a summary, the features that are not currently implemented into Frama-C include in particular:

- some built-in predicates and logical functions;

¹Additional remarks on the feature may appear as footnote

- abrupt termination clauses in statement contracts (section 2.4.4);
- definition of logical types (section 2.6);
- specification modules (section 2.6.11);
- model variables and Model fields (section 2.54);
- only basic support for ghost code is provided (section 2.12);
- verification of non interference of ghost code (p. 63);
- specification of volatile variables (section 2.12.1);

1.2 Generalities about Annotations

In this document, we consider that specifications are given as annotations in comments written directly in C source files, so that source files remain compilable. Those comments must start by `/*@` or `//@` and end as usual in C.

In some contexts, it is not possible to modify the source code. It is strongly recommended that a tool which implements ACSL specifications provides technical means to store annotations separately from the source. It is not the purpose of this document to describe such means. Nevertheless, some of the specifications, namely those at a global level, can be given in separate files: logical specifications can be imported (see Section 2.6.11) and a function contract can be attached to a copy of the function profile (see Section 2.3.5).

1.2.1 Kinds of annotations

- Global annotations:
 - *function contract*: such an annotation is inserted just before the declaration or the definition of a function. See section 2.3.
 - *global invariant*: this is allowed at the level of global declarations. See section 2.11.
 - *type invariant*: this allows to declare both structure or union invariants, and invariants on type names introduced by `typedef`. See section 2.11.
 - *logic specifications*: definitions of logic functions or predicates, lemmas, axiomatizations by declaration of new logic types, logic functions, predicates with axioms they satisfy. Such an annotation is placed at the level of global declarations.
- Statement annotations:
 - *assert clause*: these are allowed everywhere a C label is allowed, or just before a block closing brace.
 - *loop annotation* (invariant, variant, assign clauses): is allowed immediately before a loop statement: `for`, `while`, `do ... while`. See Section 2.4.2.
 - *statement contract*: very similar to a function contract, and placed before a statement or a block. Semantical conditions must be checked (no `goto` going inside, no `goto` going outside). See Section 2.4.4.
 - *ghost code*: regular C code, only visible from the specifications, that is only allowed to modify ghost variables. See section 2.12. This includes ghost braces for enclosing blocks.

1.2.2 Parsing annotations in practice

In JML, parsing is done by simply ignoring `/*@`, `/*@` and `*/` at the lexical analysis level. This technique could modify the semantics of the code, for example:

```
1 | return x /*@ +1 */ ;
```

In our language, this is forbidden. Technically, the current implementation of Frama-C isolates the comments in a first step of syntax analysis, and then parses a second time. Nevertheless, the grammar and the corresponding parser must be carefully designed to avoid interaction of annotations with the code. For example, in code such as

```
1 |   if (c) /*@ assert P;
2 |     c=1;
```

the statement `c=1` must be understood as the branch of the `if`. This is ensured by the grammar below, saying that `assert` annotations are not statements themselves, but attached to the statement that follows, like C labels.

1.2.3 About preprocessing

This document considers C source *after* preprocessing. Tools must decide how they handle preprocessing (what to do with annotations, whether macro substitution should be performed, etc.)

1.2.4 About keywords

Additional keywords of the specification language start with a backslash, if they are used in position of a term or a predicate (which are defined in the following). Otherwise they do not start with a backslash (like `ensures`) and they remain valid identifiers.

1.3 Notations for grammars

In this document, grammar rules are given in BNF form. In grammar rules, we use extra notations e^* to denote repetition of zero, one or more occurrences of e , e^+ for repetition of one or more occurrences of e , and $e^?$ for zero or one occurrence of e . For the sake of simplicity, we only describe annotations in the usual `/*@ ... */` style of comments. One-line annotations in `/*@` comments are alike.



Specification language

2.1 Lexical rules

Lexical structure mostly follows that of ANSI C. A few differences must be noted.

- The at sign (@) is a blank character, thus equivalent to a space character.
- Identifiers may start with the backslash character (\).
- Some UTF8 characters may be used in place of some constructs, as shown in the following table:

\geq	\geq	0x2265
\leq	\leq	0x2264
$>$	$>$	0x003E
$<$	$<$	0x003C
\neq	\neq	0x2262
\equiv	\equiv	0x2261
\implies	\implies	0x21D2
\iff	\iff	0x21D4
\wedge	\wedge	0x2227
\vee	\vee	0x2228
\oplus	\oplus	0x22BB
\neg	\neg	0x00AC
- (unary minus)	$-$	0x2212
\forall	\forall	0x2200
\exists	\exists	0x2203
integer	\mathbb{Z}	0x2124
real	\mathbb{R}	0x211D
boolean	\mathbb{B}	0x1D539

- Comments can be put inside ACSL annotations. They use the C++ format, *i.e.* begin with // and extend to the end of current line.

2.2 Logic expressions

This first section presents the language of expressions one can use in annotations. These are called *logic expressions* in the following. They correspond to pure C expressions, with additional constructs that we will introduce progressively.

<i>bin-op</i>	::=	+ - * / %	
		= != <= >= > <	
		&&	boolean operations
		& --> <--> ^	bitwise operations
<i>unary-op</i>	::=	+ -	unary plus and minus
		!	boolean negation
		~	bitwise complementation
		*	pointer dereferencing
		&	address-of operator
<i>term</i>	::=	\true \false	
		<i>integer</i>	integer constants
		<i>real</i>	real constants
		<i>id</i>	variables
		<i>unary-op term</i>	
		<i>term bin-op term</i>	
		<i>term</i> [<i>term</i>]	array access
		{ <i>term</i> \with [<i>term</i>] = <i>term</i> }	array functional modifier
		<i>term</i> . <i>id</i>	structure field access
		{ <i>term</i> \with . <i>id</i> = <i>term</i> }	field functional modifier
		<i>term</i> -> <i>id</i>	
		(<i>type-expr</i>) <i>term</i>	cast
		<i>id</i> (<i>term</i> (, <i>term</i>)*)	function application
		(<i>term</i>)	parentheses
		<i>term</i> ? <i>term</i> : <i>term</i>	ternary condition
		\let <i>id</i> = <i>term</i> ; <i>term</i>	local binding
		sizeof (<i>term</i>)	
		sizeof (<i>C-type-expr</i>)	
		<i>id</i> : <i>term</i>	syntactic naming

Figure 2.1: Grammar of terms

Figures 2.1 and 2.2 present the grammar for the basic constructs of logic expressions. In that grammar, we distinguish between *predicates* and *terms*, following the usual distinction between propositions and terms in classical first-order logic. The grammar for binders and type expressions is given separately in Figure 2.3.

With respect to C pure expressions, the additional constructs are as follows:

Additional connectives C operators && (UTF8: \wedge), || (UTF8: \vee) and ! (UTF8: \neg) are used as logical connectives. There are additional connectives ==> (UTF8: \implies) for implication, <==> (UTF8: \iff) for equivalence and ^^ (UTF8: \oplus) for exclusive or. These logical connectives all have a bitwise counterpart, either C ones like &, |, ~ and ^, or additional ones like bitwise implication --> and bitwise equivalence <-->.

Quantification Universal quantification is denoted by `\forall` $\tau x_1, \dots, x_n; e$ and existential quantification by `\exists` $\tau x_1, \dots, x_n; e$.

Local binding `\let x = e1; e2` introduces the name `x` for expression `e1` which can be used in expression `e2`.

Conditional `c ? e1 : e2`. There is a subtlety here: the condition may be either a boolean term or a predicate. In case of a predicate, the two branches must be also predicates, so that this construct acts as a connective with the following semantics: `c ? e1 : e2` is equivalent to `(c ==> e1) && (! c ==> e2)`.

Syntactic naming `id : e` is a term or a predicate equivalent to `e`. It is different from local naming with `\let`: the name cannot be reused in other terms or predicates. It is only for readability purposes.

Functional modifier The composite element modifier is an additional operator in relation to the C structure field and array accessors. The expression `{ s \with .id = v }` denotes the same structure than `s`, except for the field `id` to be modified by `v`. The equivalent expression for an array is `{ t \with [i] = v }` which returns the same array than `t`, except for the `ith` element whose value is updated to `v`. See Section 2.10 for an example of use of these operators.

Logic functions Applications in terms and in propositions are not applications of C functions, but of logic functions or predicates; see Section 2.6 for detail.

Consecutive comparison operators The construct `t1 relop1 t2 relop2 t3 ... tk` with several consecutive comparison operators is a shortcut for `(t1 relop1 t2) && (t2 relop2 t3) && ...`. It is required that the `relopi` operators

<code>rel-op</code>	<code>::=</code>	<code>==</code> <code>!=</code> <code><=</code> <code>>=</code> <code>></code> <code><</code>	
<code>pred</code>	<code>::=</code>	<code>\true</code> <code>\false</code>	
		<code>term (rel-op term)⁺</code>	comparisons (see remark)
		<code>id (term (, term)*)</code>	predicate application
		<code>(pred)</code>	parentheses
		<code>pred && pred</code>	conjunction
		<code>pred pred</code>	disjunction
		<code>pred ==> pred</code>	implication
		<code>pred <==> pred</code>	equivalence
		<code>! pred</code>	negation
		<code>pred ^^ pred</code>	exclusive or
		<code>term ? pred : pred</code>	ternary condition
		<code>pred ? pred : pred</code>	
		<code>\let id = term ; pred</code>	local binding
		<code>\let id = pred ; pred</code>	
		<code>\forall</code> <code>binders ; pred</code>	universal quantification
		<code>\exists</code> <code>binders ; pred</code>	existential quantification
		<code>id : pred</code>	syntactic naming

Figure 2.2: Grammar of predicates

$binders$	$::=$	$binder$	$(, binder)^*$
$binder$	$::=$	$type\text{-}expr$	$variable\text{-}ident$ $(, variable\text{-}ident)^*$
$type\text{-}expr$	$::=$	$logic\text{-}type\text{-}expr$	$ C\text{-}type\text{-}expr$
$logic\text{-}type\text{-}expr$	$::=$	$built\text{-}in\text{-}logic\text{-}type$	
		$ id$	$type\ id$
$built\text{-}in\text{-}logic\text{-}type$	$::=$	$boolean$	$ integer$ $ real$
$variable\text{-}ident$	$::=$	id	$ * variable\text{-}ident$
		$ variable\text{-}ident$	$[]$
		$ (variable\text{-}ident$	$)$

Figure 2.3: Grammar of binders and type expressions

must be in the same “direction”, *i.e.* they must all belong either to $\{<, <=, ==\}$ or to $\{>, >=, ==\}$. Expressions such as $x < y > z$ or $x != y != z$ are not allowed.

To enforce the same interpretation as in C expressions, one may need to add extra parentheses: $a == b < c$ is equivalent to $a == b \ \&\& \ b < c$, whereas $a == (b < c)$ is equivalent to $\backslash let \ x = b < c ; \ a == x$. This situation raises some issues, see example below.

There is a subtlety regarding comparison operators: they are predicates when used in predicate position, and boolean functions when used in term position.

Example 2.1 *Let’s consider the following example:*

```
| int f(int a, int b) { return a < b; }
```

- *the obvious postcondition $\backslash result == a < b$ is not the right one because it is actually a shortcut for $\backslash result == a \ \&\& \ a < b$.*
- *adding parentheses results in a correct post-condition $\backslash result == (a < b)$. Note however that there is an implicit conversion (see Sec. 2.2.3) from the `int` (the type of $\backslash result$) to `boolean` (the type of $(a < b)$)*
- *an equivalent post-condition, which does not rely on implicit conversion, is $(\backslash result != 0) == (a < b)$. Both pairs of parentheses are mandatory.*
- *$\backslash result == (integer)(a < b)$ is also acceptable because it compares two integers. The cast towards `integer` enforces $a < b$ to be understood as a boolean term. Notice that a cast towards `int` would also be acceptable.*
- *$\backslash result != 0 <==> a < b$ is acceptable because it is an equivalence between two predicates.*

2.2.1 Operators precedence

The precedence of C operators is conservatively extended with additional operators, as shown in Figure 2.4. In this table, operators are sorted from highest to lowest priority. Operators of same priority are presented on the same line.

class	associativity	operators
selection	left	[...] -> .
unary	right	! ~ + - * & (cast) sizeof
multiplicative	left	* / %
additive	left	+ -
shift	left	<< >>
comparison	left	< <= > >=
comparison	left	== !=
bitwise and	left	&
bitwise xor	left	^
bitwise or	left	
bitwise implies	left	-->
bitwise equiv	left	<-->
connective and	left	&&
connective xor	left	^^
connective or	left	
connective implies	right	==>
connective equiv	left	<==>
ternary connective	right	...?...:...:
binding	left	\forall \exists \let
naming	right	:

Figure 2.4: Operator precedence

There is a remaining ambiguity between the connective $\dots? \dots : \dots$ and the labelling operator $:$. Consider for instance the expression $x?y:z:t$. The precedence table does not indicate whether this should be understood as $x?(y:z):t$ or $x?y:(z:t)$. Such a case must be considered as a syntax error, and should be fixed by explicitly adding parentheses.

2.2.2 Semantics

The semantics of logic expressions in ACSL is based on mathematical first-order logic [25]. In particular, it is a 2-valued logic with only total functions. Consequently, expressions are never “undefined”. This is an important design choice and the specification writer should be aware of that. (For a discussion about the issues raised by such design choices, in similar specification languages such as JML, see the comprehensive list compiled by Patrice Chalin [4, 5].)

Having only total functions implies that one can write terms such as $1/0$, or $*p$ when p is null (or more generally when it points to a non-properly allocated memory cell). In particular, the predicates $1/0 == 1/0$ and $*p == *p$ are valid, since they are instances of the axiom $\forall x, x = x$ of first-order logic. The reader should not be alarmed, because there is no way to deduce anything useful from such terms. As usual, it is up to the specification designer to write consistent assertions. For example, when introducing the following lemma (see Section 2.6):

```

1 | /*@ lemma div_mul_identity:
2 |   @ \forall real x, real y; y != 0.0 ==> y*(x/y) == x;
3 |   @*/

```

a premise is added to require y to be non zero.

2.2.3 Typing

The language of logic expressions is typed (as in *multi-sorted* first-order logic). Types are either C types or *logic types* defined as follows:

- “mathematical” types: `integer` for unbounded, mathematical integers, `real` for real numbers, `boolean` for booleans (with values written `\true` and `\false`);
- logic types introduced by the specification writer (see Section 2.6).

There are implicit coercions for numeric types:

- C integral types `char`, `short`, `int` and `long`, signed or unsigned, are all subtypes of type `integer`;
- `integer` is itself a subtype of type `real`;
- C types `float` and `double` are subtypes of type `real`.

Notes:

- There is a distinction between booleans and predicates. The expression `x<y` in term position is a boolean, and the same expression is also allowed in predicate position.
- Unlike in C, there is a distinction between booleans and integers. There is an implicit promotion from integers to booleans, thus one may write `x && y` instead of `x != 0 && y != 0`. If the reverse conversion is needed, an explicit cast is required, *e.g.* `(int)(x>0)+1`, where `\false` becomes 0 and `\true` becomes 1.
- Quantification can be made over any type: logic types and C types. Quantification over pointers must be used carefully, since it depends on the memory state where dereferencing is done (see Section 2.2.4 and Section 2.6.9).

Formal typing rules for terms are given in appendix A.3.

2.2.4 Integer arithmetic and machine integers

The following integer arithmetic operations apply to *mathematical integers*: addition, subtraction, multiplication, unary minus. The value of a C variable of an integral type is promoted to a mathematical integer. As a consequence, there is no “arithmetic overflow” in logic expressions.

Division and modulo are also mathematical operations, which coincide with the corresponding C operations on C machine integers, thus following the ANSI C99 conventions. In particular, these are not the usual mathematical Euclidean division and remainder. Generally speaking, division rounds the result towards zero. The results are not specified if divisor is zero; otherwise if q and r are the quotient and the remainder of n divided by d then:

- $|d \times q| \leq |n|$, and $|q|$ is maximal for this property;
- q is zero if $|n| < |d|$;
- q is positive if $|n| \geq |d|$ and n and d have the same sign;

- q is negative if $|n| \geq |d|$ and n and d have opposite signs;
- $q \times d + r = n$;
- $|r| < |d|$;
- r is zero or has the same sign as n .

Example 2.2 *The following examples illustrate the results of division and modulo depending on the sign of their arguments:*

- $5/3$ is 1 and $5\%3$ is 2;
- $(-5)/3$ is -1 and $(-5)\%3$ is -2;
- $5/(-3)$ is -1 and $5\%(-3)$ is 2;
- $(-5)/(-3)$ is 1 and $(-5)\%(-3)$ is -2.

Hexadecimal and octal constants

Hexadecimal and octal constants are always non-negative. Suffixes `u` and `l` for C constants are allowed but meaningless.

Casts and overflows

In logic expressions, casting from mathematical integers to an integral C type t (such as `char`, `short`, `int`, etc.) is allowed and is interpreted as follows: the result is the unique value of the corresponding type that is congruent to the mathematical result modulo the cardinal of this type, that is $2^{8 \times \text{sizeof}(t)}$.

Example 2.3 `(unsigned char)1000` is $1000 \bmod 256$ i.e. 232.

To express in the logic the value of a C expression, one has to add all the necessary casts. For example, the logic expression denoting the value of the C expression `x*y+z` is `(int)((int)(x*y)+z)`. Note that there is no implicit cast from integers to C integral types.

Example 2.4 *The declaration*

```
| //@ logic int f(int x) = x+1 ;
```

is not allowed because `x+1`, which is a mathematical integer, must be casted to `int`. One should write either

```
| //@ logic integer f(int x) = x+1 ;
```

or

```
| //@ logic int f(int x) = (int)(x+1) ;
```

Quantification on C integral types

Quantification over a C integral type corresponds to integer quantification over the corresponding interval.

Example 2.5 *Thus the formula*

```
| \forallall char c; c <= 1000
```

is equivalent to

```
| \forallall integer c; MIN_CHAR <= c <= MAX_CHAR ==> c <= 1000
```

where the the bounds MIN_CHAR and MAX_CHAR are defined in limits.h

Size of C integer types

Note that the size of C types is architecture-dependent. ACSL does not enforce these sizes either, hence the semantics of terms involving such types is also architecture-dependent. The `sizeof` operator may be used in annotations and is assumed to be consistent with the C code. For instance, it should be possible to verify the following code:

```
1 | /*@ ensures \result <= sizeof(int); */
2 | int f() { return sizeof(char); }
```

Constants giving maximum and minimum values of those types may be provided in a library.

Enum types

Enum types are also interpreted as mathematical integers. Casting an integer into an enum in the logic gives the same result as if the cast was performed in the C code.

Bitwise operations

Like arithmetic operations, bitwise operations apply to any mathematical integer: any mathematical integer has an unique infinite 2-complement binary representation with infinitely many zeros (for non-negative numbers) or ones (for negative numbers) on the left. Bitwise operations apply to this representation.

Example 2.6

- $7 \ \& \ 12 == \dots 00111 \ \& \ \dots 001100 == \dots 00100 == 4$
- $-8 \ | \ 5 == \dots 11000 \ | \ \dots 00101 == \dots 11101 == -3$
- $\sim 5 == \sim \dots 00101 == \dots 111010 == -6$
- $-5 \ \ll \ 2 == \dots 11011 \ \ll \ 2 == \dots 11101100 == -20$
- $5 \ \gg \ 2 == \dots 00101 \ \gg \ 2 == \dots 0001 == 1$
- $-5 \ \gg \ 2 == \dots 11011 \ \gg \ 2 == \dots 1110 == -2$

2.2.5 Real numbers and floating point numbers

Floating-point constants and operations are interpreted as mathematical real numbers: a C variable of type float or double is implicitly promoted to a real. Integers are promoted to reals if necessary. Usual binary operations are interpreted as operators on real numbers, hence they never involve any rounding nor overflow

Example 2.7 *In an annotation, $1e+300 * 1e+300$ is equal to $1e+600$, even if that last number exceeds the largest representable number in double precision: there is no "overflow".*

*$2*0.1$ is equal to the real number 0.2, and not to any of its floating-point approximation: there is no "rounding".*

Unlike the promotion of C integer types to mathematical integers, there are special float values which do not naturally map to a real number, namely the IEEE-754 special values for “not-a-number”, $+\infty$ and $-\infty$. See below for a detailed discussions on such special values. However, remember that ACSL’s logic has only total functions. Thus, there are implicit promotion functions `real_of_float` and `real_of_double` whose results on the 3 values above is left unspecified.

In logic, real literals can also be expressed under the hexadecimal form of C99: `0xhh.hhp±dd` where *h* are hexadecimal digits and *dd* is in decimal, denotes number $hh.hh \times 2^{dd}$, e.g. `0x1.Fp-4` is $(1 + 15/16) \times 2^{-4}$.

Usual operators for comparison are interpreted as real operators too. In particular, equality operation \equiv of float (or double) expressions means equality of the real numbers they represent respectively. Or equivalently, $x \equiv y$ for x, y two float variables means `real_of_float(x) \equiv real_of_float(y)` with the mathematical equality of real numbers.

Special predicates are also available to express the comparison operators of float (resp. double) numbers as in C: `\eq_float`, `\gt_float`, `\ge_float`, `\le_float`, `\lt_float`, `\ne_float` (resp. for double).

Casts, infinity and NaNs

Casting from a C integer type or a float type to a float or a double is as in C: the same conversion operations apply.

Conversion of real numbers to float or double values indeed depends on various possible rounding modes defined by the IEEE 754 standard [24, 26]. These modes are defined by a logic type (see section 2.6.8):

```
/*@ type rounding_mode = \Up | \Down | \ToZero | \NearestAway | \NearestEven;
*/
```

Then rounding a real number can be done explicitly using functions

```
logic float \round_float(rounding_mode m, real x);
logic double \round_double(rounding_mode m, real x);
```

Cast operators `(float)` and `(double)` applied to a mathematical integer or real number x are equivalent to apply rounding functions above with the nearest-even rounding mode (which is the default rounding mode in C programs) If the source real number is too large, this may also result into one of the special values $+\infty$ and $-\infty$.

Example 2.8 *We have $(float)0.1 \equiv 13421773 \times 2^{-27}$ which is equal to 0.100000001490116119384765625*

Notice also that unlike for integers, suffixes `f` and `l` are meaningful, because they implicitly add a cast operator as above.

This semantics of casts ensures that the float result `r` of a C operation $e_1 \text{ op } e_2$ on floats, if there is no overflow and if the default rounding mode is not changed in the program, has the same real value as the logic expression $(\text{float})(e_1 \text{ op } e_2)$. Notice that this is not true for the equality `\eq_float` of floats: $-0.0 + -0.0$ in C is equal to the float number -0.0 , which is not `\eq_float` to 0.0 , which is the value of the logic expression $(\text{float})(-0.0 + -0.0)$.

Finally, additional predicates are provided on float and double numbers, which check that their argument is NaN or a finite number respectively:

```
1 | predicate \is_NaN(float x); \\
2 | predicate \is_NaN(double x); \\
3 | predicate \is_finite(float x); \\
4 | predicate \is_finite(double x);
```

Quantification

Quantification over a variable of type `real` is of course usual quantification over real numbers.

Quantification over float (resp. double) types is allowed too, and is supposed to range over all real numbers representable as floats (resp doubles). In particular, this does not include NaN, $+\infty$ and $-\infty$ in the considered range.

Mathematical functions

Classical mathematical operations like exponential, sine, cosine, and such are available as built-in:

```
integer \min(integer x, integer y) ;
integer \max(integer x, integer y) ;
real \min(real x, real y) ;
real \max(real x, real y) ;
integer \abs(integer x) ;
real \abs(real x) ;

real \sqrt(real x) ;
real \pow(real x, real y) ;

integer \ceil(real x) ;
integer \floor(real x) ;

real \exp(real x) ;
real \log(real x) ;
real \log10(real x) ;

real \cos(real x) ;
real \sin(real x) ;
real \tan(real x) ;

real \cosh(real x) ;
real \sinh(real x) ;
real \tanh(real x) ;

real \acos(real x) ;
real \asin(real x) ;
real \atan(real x) ;

real \atan2(real y, real x) ;
real \hypot(real x, real y) ;
```


Exact computations

In order to specify properties of rounding errors, it is useful to express something about the so-called *exact* computations [3]: the computations that would be performed in an ideal mode where variables denote true real numbers.

To express such exact computations, two special constructs exist in annotations:

- `\exact(x)` denotes the value of the C variable x (or more generally any C left-value) as if the program was executed with ideal real numbers.
- `\round_error(x)` is a shortcut for $|x - \text{\exact}(x)|$

Example 2.9 *Here is an example of a naive approximation of cosine [2].*

```

/*@ requires \abs(\exact(x)) <= 0x1p-5;
    requires \round_error(x) <= 0x1p-20;
    @ ensures \abs(\exact(\result) - \cos(\exact(x))) <= 0x1p-24;
    @ ensures \round_error(\result) <= \round_error(x) + 0x3p-24;
    */
float cosine(float x) {
    return 1.0f - x * x * 0.5f;
}

```

2.2.6 C arrays and pointers

Address operator, array access, pointer arithmetic and dereferencing

These operators are similar to their corresponding C operators.

address-of operator should be used with caution. Values in logic do not lie in C memory so it does not mean anything to talk about their “address”.

Unlike C, there is no implicit cast from from an array type to a pointer type. Nevertheless, arithmetic and dereferencing over arrays lying in C memory are allowed like in C.

Example 2.10 *Dereferencing a C array is equivalent an access to the first element of the array ; shifting it from i denotes the address of it i^{th} element.*

```

int tab[10] = { 1 } ;
int x ;
int *p = &x;

/*@ requires p == &x
int main(void){
    /*@ assert tab[0]==1 && *p == x;
    /*@ assert *tab == 1;
    int *q = &tab[3];
    /*@ assert q+1 == tab+4;
    ...
}

```

Since pointers can only refers values lying in C memory, $p \rightarrow s$ is always equivalent to $(*p).s$. On the contrary, $t[i]$ is not always equivalent to $*(t+i)$, especially for arrays lying not in C (see Section 2.2.7 for details about arrays as possible logic values).

Function pointers

Pointers to C functions are allowed in logic. The only possible use of them is to check for equality.

Example 2.11

```
int f(int x);
int g(int x);

/*@ requires p == &f || p == &g;
void h(int(*p)(int)) {
...
}
```

2.2.7 Structures, Unions and Arrays in logic

Aggregate C objects (i.e. structures, unions and arrays) are also possible values for terms in logic. They can be passed as parameters (and also returned) to logic functions, tested for equality, etc. like any other values.

Aggregate types can be declared in logic, and their contents may be any logic types themselves. Constructing such values in logic can be performed using a syntax similar to C designated initializers.

Example 2.12 *Array types in logic may be declared either with or without an explicit non-negative length. Access to the length of a logic array can be done with `\length`.*

```
/*@ type point = struct { real x; real y; };
/*@ type triangle = point[3];

/*@ logic point origin = { .x = 0.0 , .y = 0.0 };
/*@ logic triangle t_iso = { [0] = origin,
@                               [1] = { .y = 2.0 , .x = 0.0 }
@                               [2] = { .x = 2.0 , .y = 0.0 } };
@*/

/*@ logic point centroid(triangle t) = {
@     .x = mean3(t[0].x,t[1].x,t[2].x);
@     .y = mean3(t[0].y,t[1].y,t[2].y);
@ };
@*/

/*@ type polygon = point[];
/*@ logic perimeter(polygon p) =
@     \sum(0, \length(p)-1, \lambda integer i;d(p[i],p[(i+1) % \length(p)])) ;
@*/
```

Beware that because of the principle of only total functions in logic, $t[i]$ can be written even if i is outside the array bounds.

Functional updates

Syntax for functional update is similar to initialization of aggregate objects.

Example 2.13 *Functional update of an array is done by*

```
| { t_iso \with [0] = { .x = 3.0 , .y = 3.0 } }
```

Functional update of a structure is done by

```
| { origin \with .x = 3.0 }
```

There is no particular syntax for a functional update of an union, but of course the following equality is not true for an object of an union type:

```
| { { object \with .x = 3.0 }
    \with .y = 2.0 } == { { object \with .y = 2.0 }
                          \with .x = 3.0 }
```

The equality predicate == applies to aggregate values, but it is required that they have the same type. Then equality amounts to recursively check equality of fields. Equality of arrays of different lengths returns false. Beware that equality of unions is also equality of all fields.

C aggregate types

C aggregate types (struct, union or array) naturally map to logic types, by recursively mapping their fields.

Example 2.14 *There is no implicit cast to type of the updated/initialized fields.*

```
| struct S { int x; float y; int t[10]; };
|
| /*@ logic integer f(struct S s) = s.t[3];
| /*@ logic struct S g(integer n, struct S s) = { s \with .x = (int)n };
```

Unlike in C, all fields should be initialized:

```
| /*@ logic struct S h(integer n, int a[10]) = {
|   @ .x = (int)n, .y = (float)0.0, .t = a
|   @ };
| /*@
```

Cast and conversion

Unlike in C, there is no implicit conversion from an array type to a pointer type. On the other hand, there is an implicit conversion from an array of a given size to an array with unspecified size (but not the converse).

Example 2.15

```
| /*@ logic point square[4] = { origin, ... };
|
| /*@ ... perimeter(square);           // well typed
| /*@ ... centroid(square);           // wrongly typed
| /*@ ... centroid((triangle)square); // well-typed (truncation)
```

An explicit cast from an array type to a pointer type is allowed only for arrays that lie in C memory. As in C, the result of the cast is the address of the first element of the array (see Section 2.2.6).

Conversely, an explicit cast from a pointer type to an array type or a structure type is allowed, and acts as collecting the values it points to.

Subtyping and cast recursively applies to fields.

Example 2.16

<i>function-contract</i>	::=	<i>requires-clause</i> *	<i>terminates-clause</i> [?]
		<i>decreases-clause</i> [?]	<i>simple-clause</i> *
		<i>named-behavior</i> *	<i>completeness-clause</i> *
<i>requires-clause</i>	::=	requires	<i>predicate</i> ;
<i>terminates-clause</i>	::=	terminates	<i>pred</i> ;
<i>decreases-clause</i>	::=	decreases	<i>term</i> (for <i>ident</i>) [?] ;
<i>simple-clause</i>	::=	<i>assigns-clause</i>	<i>ensures-clause</i>
		<i>abrupt-clause-fn</i>	
<i>assigns-clause</i>	::=	assigns	<i>locations</i> ;
<i>locations</i>	::=	<i>location</i>	(, <i>location</i>) [*] \nothing
<i>ensures-clause</i>	::=	ensures	<i>predicate</i> ;
<i>named-behavior</i>	::=	behavior	<i>id</i> : <i>behavior-body</i>
<i>behavior-body</i>	::=	<i>assumes-clause</i> *	<i>requires-clause</i> *
		<i>simple-clause</i> *	
<i>assumes-clause</i>	::=	assumes	<i>predicate</i> ;
<i>completeness-clause</i>	::=	complete behaviors	(<i>id</i> (, <i>id</i>) [*]) [?] ;
		disjoint behaviors	(<i>id</i> (, <i>id</i>) [*]) [?] ;

Figure 2.5: Grammar of function contracts

<i>term</i>	::=	\old (<i>term</i>)	old value
		\result	result of a function
<i>pred</i>	::=	\old (<i>pred</i>)	

Figure 2.6: \old and \result in terms

```

struct { float u,v; } p[10];
/*@ assert centroid((point[3])p) == ...
/*@ assert perimeter((point[])p) == ...

```

Precisely, conversion of a pointer p of type τ^* to an logic array of type $\tau[]$ returns a logic array t such that

$$\text{length}(t) = (\text{\block_length}(p) - \text{\offset}(p)) / \text{sizeof}(\tau)$$

2.2.8 String literals

2.3 Function contracts

Figure 2.5 shows a grammar for function contracts. *location* denotes a memory location and is defined in Section 2.3.4. *abrupt-clauses* allow to specify what happens when the function does not return normally but exits abruptly. They are defined in Section 2.9.

This section is organized as follows. First, the grammar for terms is extended with two new constructs. Then Section 2.3.2 introduces *simple contracts*. Finally, Section 2.3.3 defines

more general contracts involving *named behaviors*. The `decreases` and `terminates` clauses are presented later in Section 2.5.

2.3.1 Built-in constructs `\old` and `\result`

Post-conditions usually require to refer to both the function result and values in the pre-state. Thus terms are extended with the following new constructs (shown in Figure 2.6).

- `\old(e)` denotes the value of predicate or term `e` in the pre-state of the function.
- `\result` denotes the returned value of the function.

`\old(e)` and `\result` can be used only in `ensures` clauses, since the other clauses already refer to the pre-state. In addition, `\result` can not be used in the contract of a function which returns `void`.

Note C function parameters are obtained by value from actual parameters that mostly remain unaltered by the function calls. For that reason, formal parameters in function contracts are defined such that they always refer implicitly to their values interpreted in the pre-state. So, `\old` construct is useless for formal parameters (in function contracts only).

2.3.2 Simple function contracts

A simple function contract, having only simple clauses and no named behavior, takes the following form:

```

1 | /*@ requires P1; requires P2; ...
2 |   @ assigns L1; assigns L2; ...
3 |   @ ensures E1; ensures E2; ...
4 | @*/

```

The semantics of such a contract is as follows:

- The caller of the function must guarantee that it is called in a state where the property $P_1 \ \&\& \ P_2 \ \&\& \ \dots$ holds.
- The called function returns a state where the property $E_1 \ \&\& \ E_2 \ \&\& \ \dots$ holds.
- All memory locations of the pre-state that do not belong to the set $L_1 \cup L_2 \cup \dots$ remain allocated and are left unchanged in the post-state. The set $L_1 \cup L_2 \cup \dots$ itself is interpreted in the pre-state.

Notice that the multiplicity of clauses are proposed mainly to improve readability since the contract above is equivalent to the following simplified one:

```

1 | /*@ requires P1 && P2 && ...;
2 |   @ assigns L1, L2, ...;
3 |   @ ensures E1 && E2 && ...;
4 | @*/

```

If no clause `requires` is given, it defaults to `\true`, and similarly for `ensures` clause. Giving no `assigns` clause means that locations assigned by the function are not specified, so the caller has no information at all on this function's side effects. See Section 2.3.5 for more details on default status of clauses.

Example 2.17 *The following function is given a simple contract for computation of the integer square root.*

```

1  /*@ requires x >= 0;
2     @ ensures \result >= 0;
3     @ ensures \result * \result <= x;
4     @ ensures x < (\result + 1) * (\result + 1);
5     @*/
6  int isqrt(int x);

```

The contract means that the function must be called with a nonnegative argument, and returns a value satisfying the conjunction of the three ensures clauses. Inside these ensures clauses, the use of the construct `\old(x)` is not necessary, even if the function modifies the formal parameter `x`, because function calls modify a copy of the effective parameters, and the effective parameters remain unaltered. In fact, `x` denotes the effective parameter of `isqrt` calls which has the same value interpreted in the pre-state than in the post-state.

Example 2.18 *The following function is given a contract to specify that it increments the value pointed to by the pointer given as argument.*

```

1  /*@ requires \valid(p);
2     @ assigns *p;
3     @ ensures *p == \old(*p) + 1;
4     @*/
5  void incrstar(int *p);

```

The contract means that the function must be called with a pointer `p` that points to a safely allocated memory location (see Section 2.7 for details on the `\valid` built-in predicate). It does not modify any memory location but the one pointed to by `p`. Finally, the ensures clause specifies that the value `*p` is incremented by one.

2.3.3 Contracts with named behaviors

The general form of a function contract contains several named behaviors (restricted to two behaviors, in the following, for readability).

```

1  /*@ requires P;
2     @ behavior b1:
3     @   assumes A1;
4     @   requires R1;
5     @   assigns L1;
6     @   ensures E1;
7     @ behavior b2:
8     @   assumes A2;
9     @   requires R2;
10    @   assigns L2;
11    @   ensures E2;
12    @*/

```

The semantics of such a contract is as follows:

- The caller of the function must guarantee that the call is performed in a state where the property $P \ \&\& \ (A_1 \implies R_1) \ \&\& \ (A_2 \implies R_2)$ holds.
- The called function returns a state where the properties $\text{\old}(A_i) \implies E_i$ hold for each i .
- For each i , if the function is called in a pre-state where A_i holds, then each memory location of that pre-state that does not belong to the set L_i remains allocated and is left unchanged in the post-state.

2.3. FUNCTION CONTRACTS

Notice that the `requires` clauses in the behaviors are proposed mainly to improve readability (to avoid some duplication of formulas), since the contract above is equivalent to the following simplified one:

```
1 /*@ requires P && (A1 ==> R1) && (A2 ==> R2);
2   @ assigns L1,L2;
3   @ ensures \old(A1) ==> E1 && \old(A2) ==> E2;
4   @*/
```

Note that a simple contract such as

```
1 /*@ requires P; assigns L; ensures E; */
```

is actually equivalent to a single named behavior as follows:

```
1 /*@ requires P;
2   @ behavior <any name>:
3   @   assumes \true;
4   @   assigns L;
5   @   ensures E;
6   @*/
```

Similarly, global `assigns` and `ensures` clauses are equivalent to a single named behavior. More precisely, the following contract

```
1 /*@ requires P;
2   @ assigns L;
3   @ ensures E;
4   @ behavior b1: ...
5   @ behavior b2: ...
6   @ ...
7   @*/
```

is equivalent to

```
1 /*@ requires P;
2   @ behavior <any name>:
3   @   assumes \true;
4   @   assigns L;
5   @   ensures E;
6   @ behavior b1: ...
7   @ behavior b2: ...
8   @ ...
9   @*/
```

Example 2.19 *In the following, `bsearch(t,n,v)` searches for element `v` in array `t` between indices `0` and `n-1`.*

```
1 /*@ requires n >= 0 && \valid(t+(0..n-1));
2   @ assigns \nothing;
3   @ ensures -1 <= \result <= n-1;
4   @ behavior success:
5   @   ensures \result >= 0 ==> t[\result] == v;
6   @ behavior failure:
7   @   assumes t_is_sorted : \forall integer k1, integer k2;
8   @       0 <= k1 <= k2 <= n-1 ==> t[k1] <= t[k2];
9   @   ensures \result == -1 ==>
10  @       \forall integer k; 0 <= k < n ==> t[k] != v;
11  @*/
12 int bsearch(double t[], int n, double v);
```

The precondition requires array `t` to be allocated at least from indices `0` to `n-1`. The two named behaviors correspond respectively to the successful behavior and the failing behavior.

Since the function is performing a binary search, it requires the array `t` to be sorted in increasing order: this is the purpose of the predicate named `t_is_sorted` in the `assumes` clause of the behavior named `failure`.

See 2.4.2 for a continuation of this example.

Example 2.20 *The following function illustrates the importance of different assigns clauses for each behavior.*

```

1  /*@ behavior p_changed:
2     @   assumes n > 0;
3     @   requires \valid(p);
4     @   assigns *p;
5     @   ensures *p == n;
6  /*@ behavior q_changed:
7     @   assumes n <= 0;
8     @   requires \valid(q);
9     @   assigns *q;
10    @   ensures *q == n;
11    @*/
12 void f(int n, int *p, int *q) {
13     if (n > 0) *p = n; else *q = n;
14 }

```

Its contract means that it assigns values pointed to by p or by q , conditionally on the sign of n .

Completeness of behaviors

Notice that in a contract with named behaviors, it is not required that the disjunction of the A_i is true, *i.e.* it is not mandatory to provide a “complete” set of behaviors. If such a condition is sought, it is possible to add the following clause to a contract:

```

/*@ ...
 @ complete behaviors  $b_1, \dots, b_n$ ;
 @*/

```

It specifies that the set of behaviors b_1, \dots, b_n is complete *i.e.* that

```
| R ==> (A1 || A2 || ... || An)
```

holds, where R is the precondition of the contract. The simplified version of that clause

```

/*@ ...
 @ complete behaviors;
 @*/

```

means that all behaviors given in the contract should be taken into account.

Similarly, it is not required that two distinct behaviors are disjoint. If desired, this can be specified with the following clause:

```

/*@ ...
 @ disjoint behaviors  $b_1, \dots, b_n$ ;
 @*/

```

It means that the given behaviors are pairwise disjoint *i.e.* that, for all distinct i and j ,

```
| R ==> ! (Ai && Aj)
```

holds. The simplified version of that clause

```

/*@ ...
 @ disjoint behaviors;
 @*/

```

means that all behaviors given in the contract should be taken into account.

$tset$	$::=$	<code>\emptyset</code>	empty set
		<code>tset -> id</code>	
		<code>tset . id</code>	
		<code>* tset</code>	
		<code>& tset</code>	
		<code>tset [tset]</code>	
		<code>term[?] .. term[?]</code>	range
		<code>\union (tset (, tset)*)</code>	union of locations
		<code>\inter (tset (, tset)*)</code>	intersection
		<code>tset + tset</code>	
		<code>(tset)</code>	
		<code>{ tset binders (; pred)[?] }</code>	set comprehension
		<code>{ term }</code>	explicit singleton
		<code>term</code>	implicit singleton
$pred$	$::=$	<code>\subset (tset , tset)</code>	set inclusion

Figure 2.7: Grammar for sets of terms

2.3.4 Memory locations and sets of terms

There are several places where one needs to describe a set of memory locations: in `assigns` clauses of function contracts, or in `loop assigns` clauses. A *memory location* is then any set of terms denoting a set of l-values. Moreover, a location given as argument to an `assigns` clause must be a set of modifiable l-values, as described in Section A.1. More generally, we introduce syntactic constructs to denote *sets of terms* which are also useful for `\separated` predicate (see Section 2.7.2)

The grammar for sets of terms is given in Figure 2.7. The semantics is given below, where s denotes any $tset$.

- `\emptyset` denotes the empty set.
- a simple term denotes a singleton set.
- `s->id` denotes the set of `x->id` for each $x \in s$.
- `s.id` denotes the set of `x.id` for each $x \in s$.
- `*s` denotes the set of `*x` for each $x \in s$.
- `&s` denotes the set of `&x` for each $x \in s$.
- `s1[s2]` denotes the set of `x1[x2]` for each $x_1 \in s_1$ and $x_2 \in s_2$.
- `t1 .. t2` denotes the set of integers between t_1 and t_2 , included. If $t_1 > t_2$, this the same as `\emptyset`
- `\union(s1, ..., sn)` denotes the union of s_1, s_2, \dots and s_n ;
- `\inter(s1, ..., sn)` denotes the intersection of s_1, s_2, \dots and s_n ;
- `s1+s2` denotes the set of `x1+x2` for each $x_1 \in s_1$ and $x_2 \in s_2$;

- (s) denotes the same set as s ;
- $\{ s \mid b ; P \}$ denotes set comprehension, that is the union of the sets denoted by s for each value b of binders satisfying predicate P (binders b are bound in both s and P).

Note that `assigns \nothing` is equivalent to `assigns \empty`; it is left for convenience.

Example 2.21 *The following function sets to 0 each cell of an array.*

```

1  /*@ requires \valid(t+(0..n-1));
2     @ assigns t[0..n-1];
3     @ assigns *(t+(0..n-1));
4     @ assigns *(t+{ i | integer i ; 0 <= i < n });
5     @*/
6  void reset_array(int t[],int n) {
7     int i;
8     for (i=0; i < n; i++) t[i] = 0;
9  }
```

It is annotated with three equivalent `assigns` clauses, each one specifying that only the set of cells $\{t[0], \dots, t[n-1]\}$ is modified.

Example 2.22 *The following function increments each value stored in a linked list.*

```

1  struct list {
2     int hd;
3     struct list *next;
4  };
5
6  // reachability in linked lists
7  /*@ inductive reachable{L}(struct list *root, struct list *to) {
8     @ case empty{L}: \forall struct list *l; reachable(l,l) ;
9     @ case non_empty{L}: \forall struct list *l1,*l2;
10    @     \valid(l1) && reachable(l1->next,l2) ==> reachable(l1,l2) ;
11    @ }
12 */
13
14 // The requires clause forbids to give a circular list
15 /*@ requires reachable(p,\null);
16    @ assigns { q->hd | struct list *q ; reachable(p,q) } ;
17    @*/
18 void incr_list(struct list *p) {
19     while (p) { p->hd++ ; p = p->next; }
20 }
```

The `assigns` clause specifies that the set of modified memory locations is the set of fields `q->hd` for each pointer `q` reachable from `p` following `next` fields. See Section 2.6.3 for details about the declaration of the predicate `reachable`.

2.3.5 Default contracts, multiple contracts

A C function can be defined only once but declared several times. It is allowed to annotate each of these declarations with contracts. Those contracts are seen as a single contract with the union of the `requires` clauses and behaviors.

On the other hand a function may have no contract at all, or a contract with missing clauses. Missing `requires` and `ensures` clauses default to `\true`. If no `assigns` clause is given, it remains unspecified. If the function under consideration has only a declaration but no body, then it means that it potentially modifies “everything”, hence in practice it will be impossible to verify anything about programs calling that function; in other words giving it a contract is in practice mandatory. On the other hand, if that function has a body, giving no `assigns` clause means in practice that it is left to tools to compute an over-approximation of the sets of assigned locations.

<i>compound-statement</i>	::=	{ <i>declaration</i> * <i>statement</i> * <i>assertion</i> ⁺ }
<i>statement</i>	::=	<i>assertion</i> <i>statement</i>
<i>assertion</i>	::=	/*@ assert <i>pred</i> ; */
		/*@ for <i>id</i> (, <i>id</i>)* : assert <i>pred</i> ; */

Figure 2.8: Grammar for assertions

2.4 Statement annotations

Annotations on C statements are of three kinds:

- Assertions: allowed before any C statement or at end of blocks.
- Loop annotations: invariant, assigns clause, variant; allowed before any loop statement: while, for, and do ... while.
- Statement contracts: allowed before any C statement, specifying their behavior in a similar manner to C function contracts.

2.4.1 Assertions

The syntax of assertions is given in Figure 2.8, as an extension of the grammar of C statements.

- `assert p` means that `p` must hold in the current state (the sequence point where the assertion occurs).
- The variant `for id1, ..., idk: assert p` associates the assertion to the named behaviors `idi`, each of them being a behavior identifier for the current function (or a behavior of an enclosing block as defined later in Section 2.4.4). It means that this assertion must hold only for the considered behaviors.

2.4.2 Loop annotations

The syntax of loop annotations is given in Figure 2.9, as an extension of the grammar of C statements.

Loop invariants

The semantics of loop invariants is defined as follows: a simple loop annotation of the form

```

1 | /*@ loop invariant I;
2 |   @ loop assigns L;
3 |   @*/
4 | ...

```

specifies that the following conditions hold.

- The predicate `I` holds before entering the loop (in the case of a for loop, this means right after the initialization expression).

<i>statement</i>	::=	<i>/*@ loop-annot */</i> <i>while (expr) statement</i> <i>/*@ loop-annot */</i> <i>for</i> <i>(expr ; expr ; expr)</i> <i>statement</i> <i>/*@ loop-annot */</i> <i>do statement</i> <i>while (expr) ;</i>
<i>loop-annot</i>	::=	<i>loop-clause*</i> <i>loop-behavior*</i> <i>loop-variant?</i>
<i>loop-clause</i>	::=	<i>loop-invariant</i> <i>loop-assigns</i>
<i>loop-invariant</i>	::=	<i>loop invariant pred ;</i>
<i>loop-assigns</i>	::=	<i>loop assigns locations ;</i>
<i>loop-behavior</i>	::=	<i>for id (, id)* :</i> <i>loop-clause*</i> annotation for behavior <i>id</i>
<i>loop-variant</i>	::=	<i>loop variant term ;</i> <i>loop variant term for id ;</i> variant for relation <i>id</i>

Figure 2.9: Grammar for loop annotations

- The predicate *I* is an inductive invariant, that is if *I* is assumed true in some state where the condition *c* is also true, and if execution of the loop body in that state ends normally at the end of the body or with a `continue` statement, *I* is true in the resulting state. If the loop condition has side effects, these are included in the loop body in a suitable way:
 - for a `while (c) s` loop, *I* must be preserved by the side-effects of *c* followed by *s*;
 - for a `for (init;c;step) s` loop, *I* must be preserved by the side-effects of *c* followed by *s* followed by `step`;
 - for a `do s while (c);` loop, *I* must be preserved by *s* followed by the side-effects of *c*.

Note that if *c* has side-effects, the invariant might not be true at the exit of the loop: the last “step” starts from a state where *I* holds, performs the side-effects of *c*, which in the end evaluates to false and exits the loop. Likewise, if a loop is exited through a `break` statement, *I* does not necessarily hold, as side effects may occur between the last state in which *I* was supposed to hold and the `break` statement.

- At any loop iteration, any location that was allocated before entering the loop, and is not member of *L* (interpreted in the current state), must remain allocated and has the same value as before entering the loop. In fact, the `loop assigns` clause specifies an inductive invariant for the locations that are not member of *L*.

A loop annotation preceded by `for id_1,...,id_k`: is similar as above, but applies only for behaviors `id_1,...,id_k`, hence in particular holds only under the assumption of their `assumes` clauses.

Remarks

- The `\old` construct is not allowed¹ in loop annotations. The `\at` form should be used to refer to another state (see Section 2.4.3).
- When a loop exits with `break` or `return` or `goto`, it is not required that the loop invariant holds. In such cases, locations that are not member of L can be deallocated or assigned between the end of the previous iteration and the exit statement.

Example 2.23 Here is a continuation of example 2.19. Note the use of a loop invariant associated to a function behavior.

```

1  /*@ requires n >= 0 && \valid(t+(0..n-1));
2    @ assigns \nothing;
3    @ ensures -1 <= \result <= n-1;
4    @ behavior success:
5    @   ensures \result >= 0 ==> t[\result] == v;
6    @ behavior failure:
7    @   assumes t_is_sorted : \forall integer k1, int k2;
8    @       0 <= k1 <= k2 <= n-1 ==> t[k1] <= t[k2];
9    @   ensures \result == -1 ==>
10   @       \forall integer k; 0 <= k < n ==> t[k] != v;
11   @*/
12 int bsearch(double t[], int n, double v) {
13   int l = 0, u = n-1;
14   /*@ loop invariant 0 <= l && u <= n-1;
15     @ for failure: loop invariant
16     @   \forall integer k; 0 <= k < n && t[k] == v ==> l <= k <= u;
17     @*/
18   while (l <= u ) {
19     int m = l + (u-1)/2; // better than (l+u)/2
20     if (t[m] < v) l = m + 1;
21     else if (t[m] > v) u = m - 1;
22     else return m;
23   }
24   return -1;
25 }
```

Loop variants

Optionally, a loop annotation may include a loop variant of the form

```

/*@ loop variant m;
   @*/
...

```

where m is a term of type `integer`.

The semantics is as follows: for each loop iteration that terminates normally or with `continue`, the value of m at end of the iteration must be smaller than its value at the beginning of the iteration. Moreover, its value at the beginning must be nonnegative. Note that the value of m at loop exit might be negative. It does not compromise termination of the loop. Here is an example:

Example 2.24

¹This is not checked

```

assertion ::= /*@ invariant pred ; */
           | /*@ for id (, id)* : invariant pred ; */

```

Figure 2.10: Grammar for general inductive invariants

```

1 void f(int x) {
2     /*@ loop variant x;
3     while (x >= 0) {
4         x -= 2;
5     }
6 }

```

It is also possible to specify termination orderings other than the usual order on integers, using the additional `for` modifier. This is explained in Section 2.5.

General inductive invariants

It is actually allowed to pose an inductive invariant anywhere inside a loop body. For example, it makes sense for a `do s while (c);` loop to contain an invariant right after statement `s`. Such an invariant is a kind of assertions, as shown in Figure 2.10.

Example 2.25 *In the following example, the natural invariant holds at this point ($\backslash\max$ and $\backslash\lambda$ are introduced later in Section 2.6.7). It would be less convenient to set an invariant at the beginning of the loop.*

```

1 /*@ requires n > 0 && \valid(t+(0..n-1));
2   @ ensures \result == \max(0,n-1,(\lambda integer k ; t[k]));
3   @*/
4 double max(double t[], int n) {
5     int i = 0; double m,v;
6     do {
7         v = t[i++];
8         m = v > m ? v : m;
9         /*@ invariant m == \max(0,i-1,(\lambda integer k ; t[k])); */
10    } while (i < n);
11    return m;
12 }

```

More generally, loops can be introduced by `gotos`. As a consequence, such invariants may occur anywhere inside a function's body. The meaning is that the invariant holds at that point, much like an `assert`. Moreover, the invariant must be inductive, *i.e.* it must be preserved across a loop iteration. Several invariants are allowed at different places in a loop body. These extensions are useful when dealing with complex control flows.

Example 2.26 *Here is a program annotated with an invariant inside the loop body:*

```

1 /*@ requires n > 0;
2   @ ensures \result == \max(0,n-1,\lambda integer k ; t[k]);
3   @*/
4 double max_array(double t[], int n) {
5     double m; int i=0;
6     goto L;
7     do {
8         if (t[i] > m) { L: m = t[i]; }
9         /*@ invariant
10        @   0 <= i < n && m == \max(0,i,\lambda integer k ; t[k]);

```


2.4. STATEMENT ANNOTATIONS

```

11     @*/
12     i++;
13     }
14     while (i < n);
15     return m;
16 }

```

The control-flow graph of the code is as follows



The invariant is inductively preserved by the two paths that go from node “inv” to itself.

Example 2.27 The program

```

1     int x = 0;
2     int y = 10;
3
4     /*@ loop invariant 0 <= x < 11;
5     @*/
6     while (y > 0) {
7         x++;
8         y--;
9     }

```

is not correctly annotated, even if it is true that x remains smaller than 11 during the execution. This is because it is not true that the property $x < 11$ is preserved by the execution of $x++ ; y--$. A correct loop invariant could be $0 \leq x < 11 \ \&\& \ x+y == 10$. It holds at loop entrance and is preserved (under the assumption of the loop condition $y > 0$).

Similarly, the following general invariants are not inductive:

```

1     int x = 0;
2     int y = 10;
3
4     while (y > 0) {
5         x++;
6         /*@ invariant 0 < x < 11;
7         y--;
8         /*@ invariant 0 <= y < 10;
9     }

```

since $0 \leq y < 10$ is not a consequence of hypothesis $0 < x < 11$ after executing $y--$; and $0 < x < 11$ cannot be deduced from $0 \leq y < 10$ after looping back through the condition $y > 0$ and executing $x++$. Correct invariants could be:

```

1     while (y > 0) {
2         x++;

```

```

3 |     /*@ invariant 0 < x < 11 && x+y == 11;
4 |     y--;
5 |     /*@ invariant 0 <= y < 10 && x+y == 10;
6 | }

```

2.4.3 Built-in construct `\at`

Statement annotations usually need another additional construct `\at(e, id)` referring to the value of the expression `e` in the state at label `id`. The label `id` can be either a regular C label, or a label added within a ghost statement as described in Section 2.12. This label must be declared in the same function as the occurrence of `\at(e, id)`, but unlike `gotos`, more restrictive scoping rules must be respected:

- the label `id` must occur before the occurrence of `\at(e, id)` in the source;
- the label `id` must not be inside an inner block.

These rules are exactly the same rules as for the visibility of local variables within C statements (see [14], Section A11.1)

Default logic labels

There are four predefined logic labels: `Pre`, `Here`, `Old` and `Post`. `\old(e)` is in fact syntactic sugar for `\at(e, Old)`.

- The label `Here` is visible in all statement annotations, where it refers to the state where the annotation appears; and in all contracts, where it refers to the pre-state for the `requires`, `assumes`, `assigns`, `variant`, `terminates`, clauses and the post-state for other clauses. It is also visible in data invariants, presented in Section 2.11.
- The label `Old` is visible in `assigns` and `ensures` clauses of all contracts (both for functions and for statement contracts described below in Section 2.4.4), and refers to the pre-state of this contract.
- The label `Pre` is visible in all statement annotations, and refers to the pre-state of the function it occurs in.
- The label `Post` is visible in `assigns` and `ensures` clauses of all contracts, and it refers to the post-state.

There is one particular case for `assigns` and `ensures` clauses of function contracts where formal parameters of functions cannot refer to the label `Post`. In such clauses formal parameters always refer implicitly to the label `Pre`, and any `\at` construct can modify the interpretation of formal parameters.

Note that no logic label is visible in global logic declarations such as lemmas, axioms, definition of predicate or logic functions. When such an annotation needs to refer to a given memory state, it has to be given a label binder: this is described in Section 2.6.9.

Example 2.28 *The code below implements the famous extended Euclid's algorithm for computing the greatest common divisor of two integers x and y , while computing at the same time two Bézout coefficients p and q such that $p \times x + q \times y = \text{gcd}(x, y)$. The loop invariant for the Bézout property needs to refer to the value of x and y in the pre-state of the function.*

2.4. STATEMENT ANNOTATIONS

```
1  /*@ requires x >= 0 && y >= 0;
2  @ behavior bezoutProperty:
3  @ ensures (*p)*x+(*q)*y == \result;
4  @*/
5  int extended_Euclid(int x, int y, int *p, int *q) {
6  int a = 1, b = 0, c = 0, d = 1;
7  /*@ loop invariant x >= 0 && y >= 0 ;
8  @ for bezoutProperty: loop invariant
9  @ a*\at(x,Pre)+b*\at(y,Pre) == x &&
10 @ c*\at(x,Pre)+d*\at(y,Pre) == y ;
11 @ loop variant y;
12 @*/
13 while (y > 0) {
14 int r = x % y;
15 int q = x / y;
16 int ta = a, tb = b;
17 x = y; y = r;
18 a = c; b = d;
19 c = ta - c * q; d = tb - d * q;
20 }
21 *p = a; *q = b;
22 return x;
23 }
```

Example 2.29 *Here is a toy example illustrating tricky issues with `\at` and labels:*

```
1  int i;
2  int t[10];
3
4  /*@ ensures 0 <= \result <= 9;
5  int any();
6
7  /*@ assigns i,t[\at(i,Post)];
8  @ ensures
9  @ t[i] == \old(t[\at(i,Here)]) + 1;
10 @ ensures
11 @ \let j = i; t[j] == \old(t[j]) + 1;
12 @*/
13 void f() {
14 i = any();
15 t[i]++;
16 }
```

The two ensures clauses are equivalent. The simpler clause `t[i] == \old(t[i]) + 1` would be wrong because in `\old(t[i])`, `i` denotes the value of `i` in the pre-state.

Also, the assigns clause `i,t[i]` would be wrong too because again in `t[i]`, the value of `i` in the pre-state is considered.

2.4.4 Statement contracts

The grammar for statement contracts is given in Figure 2.11. It is similar to function contracts, but without decreases clause. Additionally, a statement contract may refer to enclosing named behaviors, with the form `for id:...`. Such contracts are only valid for the corresponding behaviors, in particular only under the corresponding assumes clause.

The ensures clause does not constraint the post-state when the annotated statement terminates abruptly with a `break`, `continue` or `return` statement, or a call to the `exit` function. To specify such behaviors, *abrupt clauses* (described in Section 2.9) need to be used.

2.5 Termination

<code>statement</code>	<code>::=</code>	<code>/*@ statement-contract */ statement</code>
<code>statement-contract</code>	<code>::=</code>	<code>(for id (, id)* :)? requires-clause*</code> <code>simple-clause-stmt* named-behavior-stmt*</code>
<code>simple-clause-stmt</code>	<code>::=</code>	<code>(simple-clause abrupt-clause-stmt)</code>
<code>named-behavior-stmt</code>	<code>::=</code>	<code>behavior id : behavior-body-stmt</code>
<code>behavior-body-stmt</code>	<code>::=</code>	<code>assumes-clause*</code> <code>requires-clause* simple-clause-stmt*</code>

Figure 2.11: Grammar for statement contracts

The property of termination concerns both loops and recursive function calls. Termination is guaranteed by attaching a measure function to each loop and each recursive function. By default, a measure is an integer expression, and measures are compared using the usual ordering over integers (Section 2.5.1). It is also possible to define measures into other domains and/or using a different ordering relation (Section 2.5.2).

2.5.1 Integer measures

Functions are annotated with integer measures with the syntax

```
| /*@ decreases e;
```

and loops are annotated similarly with the syntax

```
| /*@ loop variant e;
```

where the logic expression `e` has type `integer`. For recursive calls, or for loops, this expression must decrease for the relation `R` defined by

```
| R(x,y) <==> x > y && x >= 0
```

In other words, the measure must be a decreasing sequence of integers which remain nonnegative, except possibly for the last value of the sequence (See example 2.24).

Example 2.30 *In example 2.23, a loop variant `u-1` decreases at each iteration, and remains nonnegative, except at the last iteration where it may become negative.*

2.5.2 General measures

More general measures on other types can be provided, using the keyword `for`. For functions it becomes

```
| /*@ decreases e for R;
```

and for loops

```
| /*@ loop variant e for R;
```

In those cases, the logic expression `e` has some type `τ` and `R` must be relation on `τ`, that is a binary predicate declared as

```
| /*@ predicate R(τ x, τ y) ...
```

(see Section 2.6 for details). Of course, to guarantee termination, it must be proved that R is a well-founded relation.

Example 2.31 *The following example illustrates a variant annotation using a pair of integers, ordered lexicographically.*

```

1 | /*@ ensures \result >= 0;
2 | int dummy();
3 |
4 | /*@ type intpair = (integer, integer);
5 |
6 | /*@ predicate lexico(intpair p1, intpair p2) =
7 |   @ \let (x1,y1) = p1 ;
8 |   @ \let (x2,y2) = p2 ;
9 |   @ x1 < x2 && 0 <= x2 ||
10 |  @ x1 == x2 && 0 <= y2 && y1 < y2;
11 |   @*/
12 |
13 | /*@ requires x >= 0 && y >= 0;
14 | void f(int x, int y) {
15 |   /*@ loop invariant x >= 0 && y >= 0;
16 |   @ loop variant (x,y) for lexico;
17 |   @*/
18 |   while (x > 0 && y > 0) {
19 |
20 |     if (dummy()) {
21 |       x--; y = dummy();
22 |     }
23 |     else y--;
24 |   }
25 | }
```

2.5.3 Recursive function calls

The precise semantics of measures on recursive calls, especially in the general case of mutually recursive functions, is given as follows. We call *cluster* a set of mutually recursive functions which is a strongly connected component of the call graph. Within each cluster, each function must be annotated with a `decreases` clause with the same relation R (syntactically). Then, in the body of any function f of that cluster, any recursive call to a function g must occur in a state where the measure attached to g is smaller (w.r.t R) than the measure of f in the pre-state of f . This also applies when g is f itself.

Example 2.32 *Here are the classical factorial and Fibonacci functions:*

```

1 | /*@ requires n <= 12;
2 |   @ decreases n;
3 |   @*/
4 | int fact(int n) {
5 |   if (n <= 1) return 1;
6 |   return n * fact(n-1);
7 | }
8 |
9 |
10 | /*@ decreases n;
11 | int fib(int n) {
12 |   if (n <= 1) return 1;
13 |   return fib(n-1) + fib(n-2);
14 | }
```

Example 2.33 *This example illustrates mutual recursion:*

```

1
2 /*@
3   requires n>=0;
4   decreases n;
5 */
6 int even(int n) {
7   if (n == 0) return 1;
8   return odd(n-1);
9 }
10
11 /*@
12   requires x>=0;
13   decreases x;
14 */
15 int odd(int x) {
16   if (x == 0) return 0;
17   return even(x-1);
18 }

```

2.5.4 Non-terminating functions

EXPERIMENTAL

There are cases where a function is not supposed to terminate. For instance, the `main` function of a reactive program might be a `while(1)` which indefinitely waits for an event to process. More generally, a function can be expected to terminate only if some preconditions are met. In those cases, a `terminates` clause can be added to the contract of the function, under the following form:

```

| /*@ terminates p;

```

The semantics of such a clause is as follows: if `p` holds, then the function is guaranteed to terminate (more precisely, its termination must be proved). If such a clause is not present (and in particular if there is no function contract at all), it defaults to `terminates \true`; that is the function is supposed to always terminate, which is the expected behavior of most functions.

Note that nothing is specified for the case where `p` does not hold: the function may terminate or not. In particular, `terminates \false`; does not imply that the function loops forever. A possible specification for a function that never terminates is the following:

```

1 /*@ ensures \false;
2   terminates \false;
3 */
4 void f() { while(1); }

```

Example 2.34 *A concrete example of a function that may not always terminate is the `incr_list` function of example 2.22. In fact, The following contract is also acceptable for this function:*

```

1
2 // this time, the specification accepts circular lists, but does not ensure
3 // that the function terminates on them (as a matter of fact, it does not).
4 /*@ terminates reachable(p, \null);
5   @ assigns { q->hd | struct list *q ; reachable(p,q) } ;
6   @*/
7 void incr_list(struct list *p) {
8   while (p) { p->hd++ ; p = p->next; }
9 }

```

$C\text{-global-decl}$	$::=$	<code>/*@ logic-def⁺ */</code>	
$logic\text{-def}$	$::=$	<code>logic-const-def</code> <code> </code> <code>logic-function-def</code> <code> </code> <code>predicate-def</code> <code> </code> <code>lemma-def</code>	
$type\text{-var}$	$::=$	<code>id</code>	
$type\text{-expr}$	$::=$	<code>type-var</code> <code> </code> <code>id</code> <code>< type-expr</code> <code>(, type-expr)* ></code>	type variable polymorphic type
$type\text{-var-binders}$	$::=$	<code>< type-var</code> <code>(, type-var)* ></code>	
$poly\text{-id}$	$::=$	<code>id</code> <code> </code> <code>id type-var-binders</code>	normal identifier polymorphic object identifier
$logic\text{-const-def}$	$::=$	<code>logic type-expr</code> <code>poly-id = term ;</code>	
$logic\text{-function-def}$	$::=$	<code>logic type-expr</code> <code>poly-id parameters =</code> <code>term ;</code>	
$predicate\text{-def}$	$::=$	<code>predicate</code> <code>poly-id parameters[?] =</code> <code>pred ;</code>	
$parameters$	$::=$	<code>(parameter</code> <code>(, parameter)*)</code>	
$parameter$	$::=$	<code>type-expr id</code>	
$lemma\text{-decl}$	$::=$	<code>lemma poly-id :</code> <code>pred ;</code>	

Figure 2.12: Grammar for global logic definitions

2.6 Logic specifications

The language of logic expressions used in annotations can be extended by declarations of new logic types, and new constants, logic functions and predicates. These declarations follows the classical setting of *algebraic specifications*. The grammar for these declarations is given in Figure 2.12.

2.6.1 Predicate and function definitions

New functions and predicates can be *defined* by explicit expressions, given after an equal sign.

Example 2.35 *The following definitions*

```
1| /*@ predicate is_positive(integer x) = x > 0;
```

<code>logic-def</code>	<code>::=</code>	<code>inductive-def</code>
<code>inductive-def</code>	<code>::=</code>	<code>inductive</code> <code>poly-id parameters?</code> <code>{ indcase* }</code>
<code>indcase</code>	<code>::=</code>	<code>case poly-id : pred ;</code>

Figure 2.13: Grammar for inductive definitions

```

2 | /*@ logic integer get_sign(real x) =
3 |   @   x > 0.0 ? 1 : ( x < 0.0 ? -1 : 0);
4 |   @*/

```

illustrates the definition of a new predicate `is_positive` with an integer parameter, and a new logic function `sign` with a real parameter returning an integer.

2.6.2 Lemmas

Lemmas are user-given propositions, a facility that might help theorem provers to establish validity of ACSL specifications.

Example 2.36 *The following lemma*

```

1 | /*@ lemma mean_property: \forall integer x,y; x <= y ==> x <= (x+y)/2 <= y;

```

is a useful hint for program like binary search.

Of course, a complete verification of an ACSL specification has to provide a proof for each lemma.

2.6.3 Inductive predicates

A predicate may also be defined by an inductive definition. The grammar for those style of definitions is given on Figure 2.13.

In general, an inductive definition of a predicate P has the form

```

1 | /*@ inductive P( $x_1, \dots, x_n$ ) {
2 |   @   case  $c_1$  :  $p_1$ ;
3 |   ...
4 |   @   case  $c_k$  :  $p_k$ ;
5 |   @ }
6 |   @*/

```

where each c_i is an identifier and each p_i is a proposition.

The semantics of such a definition is that P is the least fixpoint of the cases, i.e. is the smallest predicate (in the sense that it is false the most often) satisfying the propositions p_1, \dots, p_k . With this general form, the existence of a least fixpoint is not guaranteed, so tools might enforce syntactic conditions on the form of inductive definitions. A standard syntactic restriction could be to allow only propositions p_i of the form

```

| \forall y_1, \dots, y_m, h_1 ==> \dots ==> h_l ==> P( $t_1, \dots, t_n$ )

```

where P occurs only positively in hypotheses h_1, \dots, h_l (definite Horn clauses, http://en.wikipedia.org/wiki/Horn_clause).

<code>logic-def</code>	<code>::=</code>	<code>axiomatic-decl</code>	
<code>axiomatic-decl</code>	<code>::=</code>	<code>axiomatic id { logic-decl* }</code>	
<code>logic-decl</code>	<code>::=</code>	<code>logic-def</code>	
		<code>logic-type-decl</code>	
		<code>logic-const-decl</code>	
		<code>logic-predicate-decl</code>	
		<code>logic-function-decl</code>	
		<code>axiom-decl</code>	
<code>logic-type-decl</code>	<code>::=</code>	<code>type logic-type ;</code>	
<code>logic-type</code>	<code>::=</code>	<code>id</code>	
		<code>id type-var-binders</code>	polymorphic type
<code>logic-const-decl</code>	<code>::=</code>	<code>logic type-expr poly-id ;</code>	
<code>logic-function-decl</code>	<code>::=</code>	<code>logic type-expr</code>	
		<code>poly-id parameters ;</code>	
<code>logic-predicate-decl</code>	<code>::=</code>	<code>predicate</code>	
		<code>poly-id parameters?</code>	
<code>axiom-decl</code>	<code>::=</code>	<code>axiom poly-id : pred ;</code>	

Figure 2.14: Grammar for axiomatic declarations

Example 2.37 *The following introduce a predicate `isgcd(x,y,d)` meaning that `d` is the greatest common divisor of `x` and `y`.*

```

1  /*@ inductive is_gcd(integer a, integer b, integer d) {
2    @ case gcd_zero:
3    @   \forall integer n; is_gcd(n,0,n);
4    @ case gcd_succ:
5    @   \forall integer a,b,d; is_gcd(b, a % b, d) ==> is_gcd(a,b,d);
6    @ }
7  @*/

```

This definition uses definite Horn clauses, hence is consistent.

Example 2.22 already introduced an inductive definition of reachability in linked-lists, and was also based on definite Horn clauses thus consistent.

2.6.4 Axiomatic definitions

Instead of an explicit definition, one may introduce an *axiomatic* definitions for a set of types, predicates and logic functions, which amounts to declare the expected profiles and a set of axioms. The grammar for those constructions is given on Figure 2.14.

Example 2.38 *The following axiomatization introduce a theory of finite lists of integers a la LISP.*

```

1  /*@ axiomatic IntList {
2    @ type int_list;
3    @ logic int_list nil;
4    @ logic int_list cons(integer n,int_list l);
5    @ logic int_list append(int_list l1,int_list l2);

```

```

6 | @ axiom append_nil:
7 | @   \forall int_list l; append(nil,l) == l;
8 | @ axiom append_cons:
9 | @   \forall integer n, int_list l1,l2;
10 | @     append(cons(n,l1),l2) == cons(n,append(l1,l2));
11 | @ }
12 | @*/

```

Like inductive definitions, there is no syntactic conditions which would guarantee axiomatic definitions to be consistent. It is usually up to the user to ensure that the introduction of axioms does not lead to a logical inconsistency.

Example 2.39 *The following axiomatization*

```

1 | /*@ axiomatic sign {
2 | @   logic integer get_sign(real x);
3 | @   axiom sign_pos: \forall real x; x >= 0. ==> get_sign(x) == 1;
4 | @   axiom sign_neg: \forall real x; x <= 0. ==> get_sign(x) == -1;
5 | @ }
6 | @*/

```

is inconsistent since it implies $\text{sign}(0.0) == 1$ and $\text{sign}(0.0) == -1$, hence $-1 == 1$

2.6.5 Polymorphic logic types

EXPERIMENTAL

We consider here an algebraic specification setting based on multi-sorted logic, where types can be *polymorphic* that is parametrized by other types. For example, one may declare the type of polymorphic lists as

```

1 | /*@ type list<A>;

```

One can then consider for instance list of integers (`list <integer>`), list of pointers (e.g. `list <char*>`), list of list of reals (`list<list <real> >`²), etc.

The grammar of Figure 2.12 contains rules for declaring polymorphic types and using polymorphic type expressions.

2.6.6 Recursive logic definitions

Explicit definitions of logic functions and predicates can be recursive. Declarations in the same bunch of logic declarations are implicitly mutually recursive, so that mutually recursive functions are possible too.

Example 2.40 *The following logic declaration*

```

1 | /*@ logic integer max_index{L}(int t[], integer n) =
2 | @   (n==0) ? 0 :
3 | @   (t[n-1]==0) ? n : max_index(t, n-1);
4 | @*/

```

defines a logic function which returns the maximal index i between 0 and $n-1$ such that $t[i]=0$.

²In this latter case, note that the two '>' must be separated by a space, to avoid confusion with the shift operator.

$term$	$::=$	$\backslash\lambda\ binder\ ;\ term$	abstraction
		$ ext\text{-}quantifier\ (term\ ,\ term\ ,\ term\)$	
$ext\text{-}quantifier$	$::=$	$\backslash\max\ \ \backslash\min\ \ \backslash\sum$	
		$ \ \backslash\product\ \ \backslash\numof$	

Figure 2.15: Grammar for higher-order constructs

Notice that there is no syntactic condition on such recursive definitions, such as limitation to primitive recursion. In essence, a recursive definition of the form $f(\mathit{args}) = e$; where f occurs in expression e is just a shortcut for axiomatic declaration of f with an axiom $\backslash\forall\mathit{all}\ \mathit{args}; f(\mathit{args}) = e$. In other words, recursive definitions are not guaranteed to be consistent, in the same way that axiomatics may introduce inconsistency. Of course, tools might provide a way to check consistency.

2.6.7 Higher-order logic constructions

EXPERIMENTAL

Figure 2.15 introduces new term constructs for higher-order logic.

Abstraction The term $\backslash\lambda\ \tau_1\ x_1, \dots, \tau_n\ x_n; \mathit{t}$ denotes the n -ary logic function which maps x_1, \dots, x_n to t . It has the same precedence as $\backslash\forall\mathit{all}$ and $\backslash\exists\mathit{ists}$

Extended quantifiers Terms $\backslash\mathit{quant}(t_1, t_2, t_3)$ where quant is \max \min \sum \product or νmof are extended quantifications. t_1 and t_2 must have type $\mathit{integer}$, and t_3 must be a unary function with an integer argument, and a numeric value (integer or real) except for $\backslash\numof$ for which it should have a boolean value. Their meanings are given as follows:

$$\begin{aligned} \backslash\max(i, j, f) &= \max\{f(i), f(i+1), \dots, f(j)\} \\ \backslash\min(i, j, f) &= \min\{f(i), f(i+1), \dots, f(j)\} \\ \backslash\sum(i, j, f) &= f(i) + f(i+1) + \dots + f(j) \\ \backslash\product(i, j, f) &= f(i) \times f(i+1) \times \dots \times f(j) \\ \backslash\numof(i, j, f) &= \#\{k \mid i \leq k \leq j \wedge f(k)\} \\ &= \backslash\sum(i, j, \backslash\lambda\ \mathit{integer}\ k; f(k) ? 1 : 0) \end{aligned}$$

If $i > j$ then $\backslash\sum$ and $\backslash\numof$ above are 0, $\backslash\product$ is 1, and $\backslash\max$ and $\backslash\min$ are unspecified (see Section 2.2.2).

Example 2.41 *Function that sums the element of an array of doubles.*

```

1  /*@ requires n >= 0 && \valid(t+(0..n-1)) ;
2  @ ensures \result == \sum(0,n-1,\lambda integer k; t[k]);
3  @*/
4  double array_sum(double t[],int n) {
5      int i;
6      double s = 0.0;
7      /*@ loop invariant 0 <= i <= n;
8       @ loop invariant s == \sum(0,i-1,\lambda integer k; t[k]);
9       @ loop variant n-i;
10     */
11     for(i=0; i < n; i++) s += t[i];
12     return s;
13 }
```

<code>logic-def</code>	<code>::= type logic-type = logic-type-def ;</code>	
<code>logic-type-def</code>	<code>::= record-type sum-type type-expr</code>	type abbreviation
<code>record-type</code>	<code>::= { type-expr id (; type-expr id)* ;? }</code>	
<code>sum-type</code>	<code>::= ? constructor (constructor)*</code>	
<code>constructor</code>	<code>::= id id (type-expr (, type-expr)*)</code>	constant constructor non-constant constructor
<code>type-expr</code>	<code>::= (type-expr (, type-expr)+)</code>	product type
<code>term</code>	<code>::= term . id \match term { match-cases } (term (, term)+) { (. id = term ;)+ } \let (id (, id)+) = term ; term</code>	record field access pattern-matching tuples records
<code>match-cases</code>	<code>::= match-case+</code>	
<code>match-case</code>	<code>::= case pat : term</code>	
<code>pat</code>	<code>::= id id (pat (, pat)*) pat pat - cst { (. id = pat)* } (pat (, pat)*) pat as id</code>	constant constructor non-constant constructor or pattern any pattern numeric constant record pattern tuple pattern pattern binding

Figure 2.16: Grammar for concrete logic types and pattern-matching

2.6.8 Concrete logic types

EXPERIMENTAL

Logic types may not only be declared but also be given a definition. Defined logic types can be either under record types, or sum types. These definitions may be recursive. For record types, the field access notation $t.id$ can be used, and for sum types, a pattern-matching construction is available. The grammar rules for these additional constructions are given in Figure 2.16

Example 2.42 *The declaration*

```
1 | @@ type list<A> = Nil | Cons(A, list<A>);
```

<code>poly-id</code>	<code>::=</code>	<code>id</code>	normal identifier
		<code>id type-var-binders</code>	identifier for polymorphic object
		<code>id label-binders</code>	normal identifier with labels
		<code>id label-binders</code> <code>type-var-binders</code>	polymorphic identifier with labels
<code>label-binders</code>	<code>::=</code>	<code>{ id (, id)* }</code>	

Figure 2.17: Grammar for logic declarations with labels

introduces a concrete definition of finite lists. The logic definition

```

1 |
2 | /*@ logic integer list_length<A>(list<A> l) =
3 |   @   \match l {
4 |     @   case Nil : 0
5 |     @   case Cons(h,t) : 1+list_length(t)
6 |   @   };
7 | /*/

```

defines the length of a list by recursion and pattern-matching.

2.6.9 Hybrid functions and predicates

Logic functions and predicates may take both (pure) C types and logic types arguments. Such an hybrid predicate (or function) can either be defined with the same syntax as before (or axiomatized).

Be it defined either directly by an expression or through a set of axioms, an hybrid function (or predicate) usually depends on one or more program points, because it depends upon memory states, *via* expressions such as:

- pointer dereferencing: `*p, p->f;`
- array access: `t[i];`
- address-of operator: `&x;`
- built-in predicate depending on memory: `\valid`

To make such a definition safe, it is mandatory to add after the declared identifier a set of labels, between curly braces, as shown on Figure 2.17. Expressions as above must then be enclosed into the `\at` construct to refer to a given label. However, to ease reading of such logic expressions, it is allowed to omit a label whenever there is only one label in the context.

Example 2.43 *The following annotations declare a function which returns the number of occurrences of a given double in an array of doubles between the given indexes, together with the related axioms. It should be noted that without labels, this axiomatization would be inconsistent, since the function would not depend on the values stored in `t`, hence the two last axioms would say both that `a==b+1` and `a==b` for some `a` and `b`.*

```

1 | /*@ axiomatic NbOcc {
2 |   @   // nb_occ(t,i,j,e) gives the number of occurrences of e in t[i..j]
3 |   @   // (in a given memory state labelled L)
4 |   @   logic integer nb_occ{L}(double t[], integer i, integer j,
5 |   @   double e);

```

```

logic-function-decl ::= logic type-expr poly-id
                    parameters reads-clause ;

logic-predicate-decl ::= predicate poly-id
                     parameters? reads-clause ;

reads-clause ::= reads locations

logic-function-def ::= logic type-expr poly-id
                    parameters reads-clause = term ;

logic-predicate-def ::= predicate poly-id
                     parameters? reads-clause = pred ;

```

Figure 2.18: Grammar for logic declarations with reads clauses

```

6  @ axiom nb_occ_empty{L}:
7  @   \forall double t[], e, integer i, j;
8  @   i > j ==> nb_occ(t,i,j,e) == 0;
9  @ axiom nb_occ_true{L}:
10 @   \forall double t[], e, integer i, j;
11 @   i <= j && t[j] == e ==>
12 @     nb_occ(t,i,j,e) == nb_occ(t,i,j-1,e) + 1;
13 @ axiom nb_occ_false{L}:
14 @   \forall double t[], e, integer i, j;
15 @   i <= j && t[j] != e ==>
16 @     nb_occ(t,i,j,e) == nb_occ(t,i,j-1,e);
17 @ }
18 @*/

```

Example 2.44 This second example defines a predicate which indicates whether two arrays of the same size are a permutation of each other. It illustrates the use of more than a single label. Thus, the `\at` operator is mandatory here. Indeed the two arrays may come from two distinct memory states. Typically, one of the post condition of a sorting function would be `permut{Pre,Post}(t,t)`.

```

1  /*@ axiomatic Permut {
2  @   // permut{L1,L2}(t1,t2,n) is true whenever t1[0..n-1] in state L1
3  @   // is a permutation of t2[0..n-1] in state L2
4  @   predicate permut{L1,L2}(double *t1, double *t2, integer n);
5  @   axiom permut_refl{L}:
6  @   \forall double *t, integer n; permut{L,L}(t,t,n);
7  @   axiom permut_sym{L1,L2} :
8  @   \forall double *t1, *t2, integer n;
9  @     permut{L1,L2}(t1,t2,n) ==> permut{L2,L1}(t2,t1,n) ;
10 @   axiom permut_trans{L1,L2,L3} :
11 @   \forall double *t1, *t2, *t3, integer n;
12 @     permut{L1,L2}(t1,t2,n) && permut{L2,L3}(t2,t3,n)
13 @     ==> permut{L1,L3}(t1,t3,n) ;
14 @   axiom permut_exchange{L1,L2} :
15 @   \forall double *t1, *t2, integer i, j, n;
16 @     \at(t1[i],L1) == \at(t2[j],L2) &&
17 @     \at(t1[j],L1) == \at(t2[i],L2) &&
18 @     (\forall integer k; 0 <= k < n && k != i && k != j ==>
19 @       \at(t1[k],L1) == \at(t2[k],L2))
20 @     ==> permut{L1,L2}(t1,t2,n);
21 @ }
22 @*/

```

2.6.10 Memory footprint specification: reads clause

EXPERIMENTAL

Logic declaration can be augmented with a `reads` clause, with the syntax given in Figure 2.18, which extends the one of Figure 2.12. This feature allows to specify the *footprint* of a hybrid predicate or function, that is the set of memory locations which it depends on. From such an information, one might deduce properties of the form $f\{L_1\}(args) = f\{L_2\}(args)$ if it is known that between states L_1 and L_2 , the memory changes are disjoint from the declared footprint.

Example 2.45 *The following is the same as example 2.43 augmented with a reads clause.*

```

1  /*@ axiomatic Nb_occ {
2    @ logic integer nb_occ{L}(double t[], integer i, integer j,
3      @                               double e)
4    @   reads t[i..j];
5    @
6    @ axiom nb_occ_empty{L}: // ...
7
8    @
9    @ // ...
10   @ }
11   @*/

```

If for example a piece of code between labels `L_1` and `L_2` modifies `t[k]` for some index `k` outside `i..j`, then one can deduce that `nb_occ{L_1}(t,i,j,e)==nb_occ{L_2}(t,i,j,e)`.

2.6.11 Specification Modules

Specification modules can be provided to encapsulate several logic definitions, for example

```

1  /*@ module List {
2    @
3    @   type list<A> = Nil | Cons(A , list<A>);
4    @
5    @   logic integer length<A>(list<A> l) =
6    @     \match l {
7    @       case Nil : 0
8    @       case Cons(h,t) : 1+length(t) } ;
9    @
10   @   logic A fold_right<A,B>((A -> B -> B) f, list<A> l, B acc) =
11   @     \match l {
12   @       case Nil : acc
13   @       case Cons(h,t) : f(h,fold_right(f,t,acc)) } ;
14   @
15   @   logic list<A> filter<A>((A -> boolean) f, list<A> l) =
16   @     fold_right((\lambda A x, list<A> acc;
17   @       f(x) ? Cons(x,acc) : acc), Nil) ;
18   @
19   @
20   @ }
21   @*/

```

Module components are then accessible using a qualified notation like `List::length`.

Predefined algebraic specifications can be provided as libraries (see section 3), and imported using a construct like

```
1 | /*@ open List;
```

where the file `list.acsl` contains logic definitions, like the `List` module above.

2.7 Pointers and physical addressing

2.7.1 Memory blocks and pointer dereferencing

The following built-in functions and predicate allows to deal with the memory state:

- `\base_addr` returns the base address of an allocated pointer
`| \base_addr : α * \rightarrow char*`
- `\block_length` returns the length of the allocated block of a pointer
`| \block_length : α * \rightarrow size_t`
- `\valid` applies to a set of terms (see Section 2.3.4) of some pointer type. `\valid(s)` holds if and only if dereferencing any $p \in s$ is safe. In particular, `\valid(\emptyset)` holds.

Some shortcuts are provided:

- `\null` is an extra notation for the null pointer (*i.e.* a shortcut for `(void*)0`). Note that as in C itself (see [13], par. 6.3.2.3), the constant 0 can have any pointer type.
- `\offset(p)` returns the offset between `p` and its base address

`\offset` : α * \rightarrow size_t
`\offset(p) = (char*)p - \base_addr(p)`
 the following property holds: for any set of pointers `s`, `\valid(s)` if and only if for all $p \in s$:

`| \offset(p) >= 0 && \offset(p) + sizeof(*p) <= \block_length(p)`

2.7.2 Separation

EXPERIMENTAL

`\separated(pt_1, \dots, pt_n)` means that for each $i \neq j$, the intersection of locations `* pt_i` and `* pt_j` is empty. Each pt_i is a set of terms of some pointer type as defined in Section 2.3.4.

2.7.3 Allocation and deallocation

EXPERIMENTAL

The following built-in predicates allow to deal with allocation and deallocation of memory blocks. They can be used in a postcondition

- `\fresh(p)` indicates that `p` was not allocated in the pre-state.
- `\freed(p)`, indicates that `p` was allocated in the pre-state but that it is not the case in the post-state.

2.8 Sets as first-class values

Sets of terms, as defined in Section 2.3.4, can be used as first-class values in annotations. All the elements of such a set must share the same type (modulo the usual implicit conversions). Sets have the built-in type `set<A>` where `A` is the type of terms contained in the set.

In addition, it is possible to consider sets of pointers to values of different types. In this case, the set is of type `set<char*>` and each of its elements `e` is converted to `(char*)e + (0..sizeof(*e)-1)`.

Example 2.46 *Here is an example where we defined the footprint of a structure, that is the set of locations that can be accessed from an object of this type.*

```

1 | struct S {
2 |     char *x;
3 |     int *y;
4 | };
5 |
6 | /*@ logic set<char*> footprint(struct S s) = \union(s.x,s.y) ;
7 |
8 | /*@ logic set<char*> footprint2(struct S s) =
9 |     @ \union(s.x,(char*)s.y+(0..sizeof(s.y)-1)) ;
10 |    */
11 |
12 | /*@ axiomatic Conv {
13 |     axiom conversion: \forallall struct S s;
14 |         footprint(s) == \union(s.x,(char*) s.y + (0 .. sizeof(int) - 1));
15 |     }
16 | */

```

Notice that in the first definition, since `union` is made with a `set<char*>` and a `set<int*>`, the result is a `set<char*>` (accordingly to typing of `union`). In other words, the two definitions above are equivalent.

This logic function can be used as argument of `\separated` or in `assigns` clause.

Thus, the `\separated` predicate satisfies the following property (with l_1 of type `set< τ_1 *>` and l_2 of type `set< τ_2 *>`)

```

1 | \separated(l1,l2) <==>
2 |   (\forallall \tau1* p; \forallall \tau2* q;
3 |     \subset(p,l1) && \subset(q,l2) ==>
4 |     (\forallall \integer i,j;
5 |       0 <= i < \sizeof(\tau1) && 0 <= j < \sizeof(\tau2) ==>
6 |       (char*)p + i != (char*)q + j))

```

and a clause `assigns l1, ..., ln` is equivalent to the postcondition

```

| \forallall char* p; \separated(\union(&l1, ..., &ln), p) ==> *p == \old(*p)

```

2.9 Abrupt termination

EXPERIMENTAL

The `ensures` clause of function and statement contracts does not constraint the post-state when the annotated function and statement terminates respectively abruptly. In such cases, *abrupt clauses* can be used inside *simple clause* or *behavior body*. The allowed constructs are shown in Figure 2.19.

<code>abrupt-clauses-fn</code>	<code>::=</code>	<code>exits-clause</code>
<code>exits-clause</code>	<code>::=</code>	<code>exits predicate ;</code>
<code>abrupt-clauses-stmt</code>	<code>::=</code>	<code>exits-clause</code> <code> breaks-clause continues-clause returns-clause</code>
<code>breaks-clause</code>	<code>::=</code>	<code>breaks predicate ;</code>
<code>continues-clause</code>	<code>::=</code>	<code>continues predicate ;</code>
<code>returns-clause</code>	<code>::=</code>	<code>returns predicate ;</code>
<code>term</code>	<code>::=</code>	<code>\exit_status</code>

Figure 2.19: Grammar of contracts about abrupt terminations

The clauses `breaks`, `continues` and `returns` can only be found in a statement contract and state properties on the program state which hold when the annotated statement terminates abruptly with the corresponding statement (`break`, `continue` or `return`).

Inside these clauses, the construct `\old(e)` is allowed and denotes, like for statement contracts `assigns` and `ensures`, the value of `e` in the pre-state of the statement. More generally, the visibility in *abrupt clauses* of predefined logics labels (presented in Section 2.4.3) is the same as in `ensures` clauses.

For the `returns` case, the `\result` construct is allowed and bounds to the returned value (if not a void function).

Example 2.47 *Here is an example which illustrates each of these special clauses for statement contracts.*

```

1 | int f(int x) {
2 |
3 |   while (x > 0) {
4 |
5 |     /*@ breaks x % 11 == 0 && x == \old(x);
6 |        @ continues (x+1) % 11 != 0 && x % 7 == 0 && x == \old(x)-1;
7 |        @ returns (\result+2) % 11 != 0 && (\result+1) % 7 != 0
8 |        @          && \result % 5 == 0 && \result == \old(x)-2;
9 |        @ ensures (x+3) % 11 != 0 && (x+2) % 7 != 0 && (x+1) % 5 != 0
10 |        @          && x == \old(x)-3;
11 |        @*/
12 |     {
13 |       if (x % 11 == 0) break;
14 |       x--;
15 |       if (x % 7 == 0) continue;
16 |       x--;
17 |       if (x % 5 == 0) return x;
18 |       x--;
19 |     }
20 |   }
21 |   return x;
22 | }
```

The `exits` clause can be used either in function or statement contracts to give behavioral properties to the `main` function or to any function that may exit the program, *e.g.* by calling the `exit` function.

In such clauses, `\old(e)` is allowed and denotes the value of `e` in the pre-state of the function or statement, and `\exit_status` is bound to the return code, *e.g.* the value returned by `main`


```

assigns-clause ::= assigns locations (\from locations)? ;
                  | assigns term \from locations = term ;

```

Figure 2.20: Grammar for dependencies information

or the argument passed to `exit`. The construct `\exit_status` can be used only in `exits` and `assigns` clauses. On the contrary, `\result` cannot be used in `exits` clauses.

Example 2.48 *Here is a complete specification of the `exit` function which performs an unconditional exit of the main function:*

```

1  /*@ assigns \nothing;
2    @ ensures \false;
3    @ exits   \exit_status == status;
4    @*/
5  void exit(int status);
6
7  int status;
8
9  /*@ assigns status;
10   @ exits   !cond && \exit_status == 1 && status == val;
11   @*/
12 void may_exit(int cond, int val) {
13   if (! cond) {
14     status = val;
15     exit(1);
16   }
17 }

```

Note that the specification of the `may_exit` function is incomplete since it allows modifications of the variable `status` when no exit is performed. Using behaviors, it is possible to distinguish between the exit case and the normal case, as in the following specification:

```

8
9  /*@ behavior no_exit :
10   @   assumes cond;
11   @   assigns \nothing;
12   @   exits   \false;
13   @ behavior no_return :
14   @   assumes !cond;
15   @   assigns status;
16   @   exits   \exit_status == 1 && status == val;
17   @   ensures \false;
18   @*/
19 void may_exit(int cond, int val) ;

```

2.10 Dependencies information

EXPERIMENTAL

An extended syntax of `assigns` clauses, described in Figure 2.20 allows to specify data dependencies and *functional expressions*.

Such a clause indicates that the assigned values can only depend upon the locations mentioned in the `\from` part of the clause. Again, this is an over-approximation: all of the locations involved in the computation of the modified values must be present, but some of locations might not be used in practice. If the `\from` clause is absent, all of the locations reachable at the given point of the program are supposed to be used. Moreover, for a single location, it is

possible to give the precise relation between its final value and the value of its dependencies. This expression is evaluated in the pre-state of the corresponding contract.

Example 2.49 *The following example is a variation over the `array_sum` function in example 2.41, in which the values of the array are added to a global variable `total`.*

```

1 | double total = 0.0;
2 |
3 | /*@ requires n >= 0 && \valid(t+(0..n-1)) ;
4 |   @ assigns total
5 |     \from t[0..n-1] = total + \sum(0,n-1,\lambda int k; t[k]);
6 |   @*/
7 | void array_sum(double t[],int n) {
8 |   int i;
9 |   for(i=0; i < n; i++) total += t[i];
10 |   return;
11 | }
```

Example 2.50 *The composite element modifier operators are useful additional constructs for such functional expressions.*

```

1 | struct buffer { int pos ; char buf[80]; } line;
2 |
3 | /*@ requires 80 > line.pos >= 0 ;
4 |   @ assigns line
5 |   @   \from line =
6 |     { line \with .buf =
7 |       { line.buf \with [line.pos] = (char)'\0' } };
8 |   @*/
9 | void add_eol() {
10 |   line.buf[line.pos] = '\0' ;
11 | }
```

2.11 Data invariants

Data invariants are properties on data that are supposed to hold permanently during the lifetime of these data. In ACSL, we distinguish between:

- *global* invariants and *type* invariants: the former only apply to specified global variables, whereas the latter are associated to a static type, and apply to any variables of the corresponding type;
- *strong* invariants and *weak* invariants: strong invariants must be valid at any time during program execution (more precisely at any *sequence point* as defined in the C standard), whereas weak invariants must be valid at *function boundaries* (function entrance and exit) but can be violated in between.

The syntax for declaring data invariants is given in Figure 2.21. The strength modifier defaults to *weak*.

Example 2.51 *In the following example, we declare*

1. a weak global invariant `a_is_positive` which specifies that global variable `a` should remain positive (weakly, so this property might be violated temporarily between functions calls);
2. a strong type invariant for variables of type `temperature`;

<code>declaration</code>	<code>::=</code>	<code>/*@ data-inv-decl */</code>
<code>data-inv-decl</code>	<code>::=</code>	<code>data-invariant type-invariant</code>
<code>data-invariant</code>	<code>::=</code>	<code>inv-strength[?] global invariant</code> <code>id : pred ;</code>
<code>type-invariant</code>	<code>::=</code>	<code>inv-strength[?] type invariant</code> <code>id (C-type-expr id) = pred ;</code>
<code>inv-strength</code>	<code>::=</code>	<code>weak strong</code>

Figure 2.21: Grammar for declarations of data invariants

3. a weak type invariant for variables of type `struct S`.

```

1 | int a;
2 | /*@ global invariant a_is_positive: a >= 0 ;
3 |
4 | typedef double temperature;
5 | /*@ strong type invariant temp_in_celsius(temperature t) =
6 |   @ t >= -273.15 ;
7 |   @*/
8 |
9 | struct S {
10 |   int f;
11 | };
12 | /*@ type invariant S_f_is_positive(struct S s) = s.f >= 0 ;

```

2.11.1 Semantics

The distinction between strong and weak invariants has to do with the sequence points where the property is supposed to hold. The distinction between global and type invariants has to do with the set of values on which they are supposed to hold.

- Weak global invariants are properties which apply to global data and hold at any function entrance and function exit.
- Strong global invariants are properties which apply to global data and hold at any step during execution (starting after initialization of these data).
- A weak type invariant on type τ must hold at any function entrance and exit, and applies to any global variable or formal parameter with static type τ . If the result of the function is of type τ , the result must also satisfy its weak invariant at function exit. However, it says nothing of fields, array elements, memory locations, etc. of type τ .
- A strong type invariant on type τ must hold at any step during execution, and applies to any global variable, local variable, or formal parameter with static type τ . If the result of the function has type τ , the result must also satisfy its strong invariant at function exit. Again, it says nothing of fields, array elements, memory locations, etc. of type τ .

Example 2.52 *The following example illustrates the use of a data invariant on a local static variable.*

```

1 void out_char(char c) {
2     static int col = 0;
3     //@ global invariant I : 0 <= col <= 79;
4     col++;
5     if (col >= 80) col = 0;
6 }

```

Example 2.53 Here is a longer example, the famous Dijkstra's Dutch flag algorithm.

```

1 typedef enum { BLUE, WHITE, RED } color;
2 /*@ type invariant isColor(color c) =
3     @ c == BLUE || c == WHITE || c == RED ;
4     @*/
5
6 /*@ predicate permut{L1,L2}(color *t1, color *t2, integer n) =
7     @ \at(\valid_range(t1,0,n),L1) && \at(\valid_range(t2,0,n),L2) &&
8     @ \numof(0,n,\lambda integer i; \at(t1[i],L1) == BLUE) ==
9     @ \numof(0,n,\lambda integer i; \at(t2[i],L2) == BLUE)
10    @ &&
11    @ \numof(0,n,\lambda integer i; \at(t1[i],L1) == WHITE) ==
12    @ \numof(0,n,\lambda integer i; \at(t2[i],L2) == WHITE)
13    @ &&
14    @ \numof(0,n,\lambda integer i; \at(t1[i],L1) == RED) ==
15    @ \numof(0,n,\lambda integer i; \at(t2[i],L2) == RED);
16    @*/
17
18 /*@ requires \valid(t+i) && \valid(t+j);
19    @ assigns t[i],t[j];
20    @ ensures t[i] == \old(t[j]) && t[j] == \old(t[i]);
21    @*/
22 void swap(color t[], int i, int j) {
23     int tmp = t[i];
24     t[i] = t[j];
25     t[j] = tmp;
26 }
27 typedef struct flag {
28     int n;
29     color *colors;
30 } flag;
31 /*@ type invariant is_colored(flag f) =
32     @ f.n >= 0 && \valid(f.colors+(0..f.n-1)) &&
33     @ \forall integer k; 0 <= k < f.n ==> isColor(f.colors[k]) ;
34     @*/
35
36 /*@ predicate isMonochrome{L}(color *t, integer i, integer j,
37     @ color c) =
38     @ \forall integer k; i <= k <= j ==> t[k] == c ;
39     @*/
40
41 /*@ assigns f.colors[0..f.n-1];
42     @ ensures
43     @ \exists integer b, integer r;
44     @ isMonochrome(f.colors,0,b-1,BLUE) &&
45     @ isMonochrome(f.colors,b,r-1,WHITE) &&
46     @ isMonochrome(f.colors,r,f.n-1,RED) &&
47     @ permut{Old,Here}(f.colors,f.colors,f.n-1);
48     @*/
49 void dutch_flag(flag f) {
50     color *t = f.colors;
51     int b = 0;
52     int i = 0;
53     int r = f.n;
54     /*@ loop invariant
55     @ (\forall integer k; 0 <= k < f.n ==> isColor(t[k])) &&
56     @ 0 <= b <= i <= r <= f.n &&
57     @ isMonochrome(t,0,b-1,BLUE) &&
58     @ isMonochrome(t,b,i-1,WHITE) &&
59     @ isMonochrome(t,r,f.n-1,RED) &&

```

```

60     @ permut{Pre,Here}(t,t,f.n-1);
61     @ loop assigns b,i,r,t[0 .. f.n-1];
62     @ loop variant r - i;
63     @*/
64     while (i < r) {
65         switch (t[i]) {
66             case BLUE:
67                 swap(t, b++, i++);
68                 break;
69             case WHITE:
70                 i++;
71                 break;
72             case RED:
73                 swap(t, --r, i);
74                 break;
75         }
76     }
77 }

```

2.11.2 Model variables and model fields

A *model variable* is a variable introduced in the specification with the keyword `model`. Its type must be a logic type. Analogously, *structures may have model fields*. These are used to provide abstract specifications to functions whose concrete implementation must remain private.

The precise syntax for declaring model variables and fields is given in Figure 2.22. It is presented as additions to the regular C grammar for variable declarations and structure field declarations.

Informal semantics of model variables is as follows.

- Model variables can only appear in specifications. They are not lvalues, thus they cannot be assigned directly (unlike ghost variables, see below).
- Nevertheless, a function contract might state that a model variable is assigned.
- When a function contract mentions model variables:
 - the precondition is implicitly existentially quantified over those variables;
 - the postconditions are universally quantified over the old values of model variables, and existentially quantified over the new values.

Thus, in practice, the only way to prove that a function body satisfies a contract with model variables is to provide an invariant relating model variables and concrete variables, as in the example below.

<i>declaration</i>	::=	<i>C-declaration</i>	
		<i>/*@ model C-declaration */</i>	model variable
<i>struct-declaration</i>	::=	<i>C-struct-declaration</i>	
		<i>/*@ model C-struct-declaration */</i>	model field

Figure 2.22: Grammar for declarations of model variables and fields

Example 2.54 Here is an example of a specification for a function which generates fresh integers. The contract is given in term of a model variable which is intended to represent the set of “forbidden” values, e.g. the values that have already been generated.

```

1  /* public interface */
2
3  //@ model set<integer> forbidden = \empty;
4
5  /*@ assigns forbidden;
6     @ ensures ! \subset(\result ,\old(forbidden))
7     @   && \subset(\result ,forbidden) && \subset(\old(forbidden),forbidden);
8     @*/
9  int gen();

```

The contract is expressed abstractly, telling that

- the forbidden set of values is modified;
- the value returned is not in the set of forbidden values, thus it is “fresh”;
- the new set of forbidden values contains both the value returned and the previous forbidden values.

An implementation of this function might be as follows, where a decision has been made to generate values in increasing order, so that it is sufficient to record the last value generated. This decision is made explicit by an invariant.

```

1  /* implementation */
2
3  int gen() {
4     static int x = 0;
5     /*@ global invariant I: \forall integer k;
6        @   Set::mem(k,forbidden) ==> x > k;
7        @*/
8     return x++;
9  }

```

Remarks Although the syntax of model variables is close to JML model variables, they differ in the sense that the type of a model variable is a logic type, not a C type. Also, the semantics above is closer to the one of B machines [1]. It has to be noticed that program verification with model variables does not have a well-established theoretical background [20, 18], so we deliberately do not provide a precise semantics in this document .

2.12 Ghost variables and statements

Ghost variables and statements are like C variables and statements, but visible only in the specifications. They are introduced by the ghost keyword at the beginning of the annotation (i.e. `/*@ ghost ... */` or `//@ ghost ...` for a one-line ghost code, as mentioned in section 1.2). The grammar is given in Figure 2.23, in which only the first form of annotation is used. In this figure, the *C*-* non-terminals refer to the corresponding grammar rules of the ISO standard, without any ACSL extension. Any non terminal of the form *ghost-non-term* for which no definition is given in the figure represents the corresponding *C-non-term* entry, in which any *entry* is substituted by *ghost-entry*.

The variations with respect to the C grammar are the following:

- Comments must be introduced by `//` and extend until the end of the line (the ghost code itself is placed inside a C comment. `/* ... */` would thus lead to incorrect C code).
- It is however possible to write multi-line annotations for ghost code. These annotations are enclosed between `/@` and `@/`. As in normal annotations, `@`s at the beginning of a line and at the end of the comment (before the final `@/`) are considered as blank.
- Logical types, such as `integer` or `real` are authorized in ghost code.
- A non-ghost function can take ghost parameters. If such a ghost clause is present in the declarator, then the list of ghost parameters must be non-empty and fixed (no `vararg` ghost). The call to the function must then provide the appropriate number of ghost parameters.
- Any non-ghost *if-statement* which does not have a non-ghost `else` clause can be augmented with a ghost one. Similarly, a non-ghost `switch` can have a ghost `default : clause` if it does not have a non-ghost one (there are however semantical restrictions for valid ghost labelled statements in a `switch`, see next paragraph for details).

Semantics of Ghost Code The question of semantics is essential for ghost code. Informally, the semantics requires that ghost statements do not change the regular program execution³ This implies several conditions, including e.g:

- Ghost code cannot modify a non-ghost C variable.
- Ghost code cannot modify a non-ghost structure field.
- If `p` is a ghost pointer pointing to a non-ghost memory location, then it is forbidden to assign `*p`.
- Body of a ghost function is ghost code, hence do not modify non-ghost variables or fields.
- If a non-ghost C function is called in ghost code, it must not modify non-ghost variables or fields.
- If a structure has ghost fields, the `sizeof` of the structure is the same as the structure without ghost fields. Also, alignment of fields remains unchanged.
- The control-flow graph of a function must not be altered by ghost statements. In particular, no ghost `return` can appear in the body of a non-ghost function. Similarly, ghost `goto`, `break`, and `continue` cannot jump outside of the innermost non-ghost enclosing block.

Semantics is specified as follows. First, the execution of a program with ghost code involves a *ghost memory heap* and a *ghost stack*, disjoint from the regular heap and stack. Ghost variables lie in the ghost heap, so as the ghost field of structures. Thus, every memory side-effect can be classified as ghost or non-ghost. Then, the semantics is that memory side-effects of ghost code must always be in the ghost heap or the ghost stack.

Notice that this semantics is not statically decidable. It is left to tools to provide approximations, correct in the sense that any code statically detected as ghost must be semantically ghost.

³Not checked in the current implementation

<code>ghost-type-specifier</code>	<code>::=</code>	<code>C-type-specifier</code> <code>logic-type-name</code>	
<code>declaration</code>	<code>::=</code>	<code>C-declaration</code> <code>/*@ ghost</code> <code>ghost-declaration</code> <code>*/</code>	
<code>direct-declarator</code>	<code>::=</code>	<code>C-direct-declarator</code> <code>direct-declarator</code> <code>(parameter-type-list[?])</code> <code>/*@ ghost</code> <code>(parameter-list)</code> <code>*/</code>	ghost args
<code>postfix-expression</code>	<code>::=</code>	<code>C-postfix-expression</code> <code>postfix-expression</code> <code>(argument-expression-list[?])</code> <code>/*@ ghost</code> <code>(argument-expression-list)</code> <code>*/</code>	call with ghosts
<code>statement</code>	<code>::=</code>	<code>C-statement</code> <code>statements-ghost</code>	
<code>statements-ghost</code>	<code>::=</code>	<code>/*@ ghost</code> <code>ghost-statement⁺</code> <code>*/</code>	
<code>ghost-selection-statement</code>	<code>::=</code>	<code>C-selection-statement</code> <code>if (expression)</code> <code>statement</code> <code>/*@ ghost else</code> <code>C-statement⁺</code> <code>*/</code>	
<code>struct-declaration</code>	<code>::=</code>	<code>C-struct-declaration</code> <code>/*@ ghost</code> <code>C-struct-declaration</code> <code>*/</code>	ghost field

Figure 2.23: Grammar for ghost statements

Example 2.55 *The following example shows some invalid assignments of ghost pointers:*

```

1 void f(int x, int *q) {
2     /*@ ghost int *p = q;
3     /*@ ghost *p = 0;
4     // above assignment is wrong: it modifies *q which lies
5     // in regular memory heap
6
7     /*@ ghost p = &x;
8     /*@ ghost *p = 0;
9     // above assignment is wrong: it modifies x which lies
10    // in regular memory stack
11
12 }
13
```

Example 2.56 *The following example shows some invalid ghost statements:*


```

1 | int f (int x, int y) {
2 |   /*@ ghost int z = x + y;
3 |   switch (x) {
4 |     case 0: return y;
5 |     /*@ ghost case 1: z=y;
6 |     // above statement is correct.
7 |     /*@ ghost case 2: { z++; break; }
8 |     // invalid, would bypass the non-ghost default
9 |     default: y++;
10 |   }
11 |   return y;
12 | }
13 |
14 | int g(int x) {
15 |   /*@ ghost int z = x;
16 |   if (x > 0) { return x; }
17 |   /*@ ghost else { z++; return x; }
18 |   // invalid, would bypass the non-ghost return
19 |   return x+1;
20 | }

```

Differences between model variables and ghost variables A ghost variable is an additional specification variable which is assigned in ghost code like any C variable. On the other hand, a model variable cannot be assigned, but one can state it is modified and can express properties about the new value, in a non-deterministic way, using logic assertions and invariants. In other words, specifications using ghost variable assignments are executable.

Example 2.57 *The example 2.54 can also be specified with a ghost variable instead of a model variable:*

```

1 | /*@ ghost set<integer> forbidden = \empty;
2 |
3 | /*@ assigns forbidden;
4 | @ ensures ! \subset(\result, \old(forbidden))
5 | @   && \subset(\result, forbidden)
6 | @   && \subset(\old(forbidden), forbidden);
7 | @*/
8 | int gen() {
9 |   static int x = 0;
10 |  /*@ global invariant I: \forall integer k;
11 |   @   \subset(k, forbidden) ==> x > k;
12 |   @*/
13 |   x++;
14 |   /*@ ghost forbidden = \union(x, forbidden);
15 |   return x;
16 | }

```

2.12.1 Volatile variables

EXPERIMENTAL

Volatile variables can not be used in logic terms, since reading such a variable may have a side effect, in particular two successive reads may return different values.

<pre> declaration ::= /*@ volatile tset reads id writes id </pre>

Figure 2.24: Grammar for volatile constructs

Specifying properties of a volatile variable may be done via a specific construct to attach two ghost functions to it. This construct, described by the grammar of Figure 2.24, has the following shape:

```

1 | volatile  $\tau$  x;
2 | /*@ volatile x reads f writes g;

```

where f and g are ghost functions with the following prototypes:

```

3 |  $\tau$  f(volatile  $\tau$ * p);
4 |  $\tau$  g(volatile  $\tau$ * p,  $\tau$  v);

```

This must be understood as a special construct to instrument the C code, where each access to the variable x is replaced by a call to $f(\&x)$, and each assignment to x of a value v is replaced by $g(\&x,v)$.

Example 2.58 *The following code is instrumented in order to inject fixed values at each read of variable x , and collect written values.*

```

1 | volatile int x;
2 |
3 | /*@ ghost /*@ requires p == &x;
4 | @ int reads_x(volatile int *p) {
5 | @ static int injector_x[] = { 1, 2, 3 };
6 | @ static int injector_count = 0;
7 | @ if (p == &x)
8 | @ return injector_x[injector_count++];
9 | @ else
10 | @ return 0; // should not happen
11 | @ }
12 | */
13 |
14 | /*@ ghost int collector_x[3];
15 | /*@ ghost int collector_count = 0;
16 |
17 | /*@ ghost /*@ requires p == &x;
18 | @ int writes_x(volatile int *p, int v) {
19 | @ if (p == &x)
20 | @ return collector_x[collector_count++] = v;
21 | @ else
22 | @ return 0; // should not happen
23 | @ }
24 | */
25 |
26 | /*@ volatile x reads reads_x writes writes_x;
27 |
28 | /*@ ensures collector_count == 3 && collector_x[2] == 2;
29 | @ ensures \result == 6;
30 | */
31 | int main () {
32 |     int i, sum = 0;
33 |     for (i=0 ; i < 3; i++) {
34 |         sum += x;
35 |         x = i;
36 |     }
37 |     return sum;
38 | }

```

2.13 Undefined values, dangling pointers

2.13.1 Initialization

`\initialized` is a predicate taking a set of l-values as argument and means that each l-value in this set is initialized.

Example 2.59 *In the following, the assertion is true.*

```

1 | int f(int n) {
2 |     int x;
3 |
4 |     if (n > 0) x = n ; else x = -n;
5 |     //@ assert \initialized(x);
6 |     return x;
7 | }

```

2.13.2 Unspecified values

`\specified` is a predicate taking a set of l-values as argument and means that each l-value in this set has a *specified value*: its value is not a dangling pointer (that is, the value is not the address of a local variable referred to outside of its scope)

Example 2.60 *In the following, the assertion is not true.*

```

1 |
2 | int* f() {
3 |     int a;
4 |     return &a;
5 | }
6 |
7 | int* g() {
8 |     int* p = f();
9 |     //@ assert \specified(p);
10 |    return p+1;
11 | }

```



Chapter 3

Libraries

Disclaimer: this chapter is unfinished, it is left here to give an idea of what it will look like in the final document.

This chapter is devoted to libraries of specification, built upon the ACSL specification language. Section 3.2 describes additional predicates introduced by the Jessie plugin of Frama-C, to propose a slightly higher level of annotation.

3.1 Libraries of logic specifications

A standard library is provided, in the spirit of the List module of Section 2.6.11

3.1.1 Real numbers

A library of general purpose functions and predicate over real numbers, floats and doubles.

Includes

- abs, exp, power, log, sin, cos, atan, etc. over reals
- isFinite predicate over floats and doubles (means not NaN nor infinity)
- rounding reals to floats or doubles with specific rounding modes.

3.1.2 Finite lists

- pure functions nil, cons, append, fold, etc.
- Path, Reachable, isFiniteList, isCyclic, etc. on C linked-lists.

3.1.3 Sets and Maps

Finite sets, finite maps, in ZB-style.

3.2 Jessie library: logical addressing of memory blocks

The Jessie library is a collection of logic specifications whose semantics is well-defined only on source codes free from architecture-dependent features. In particular it is currently incompatible with pointer casts or unions (although there is ongoing work to support some of them [21]). As a consequence, in this particular setting, a valid pointer of some type τ^* necessarily points to a memory block which contains values of type τ .

3.2.1 Abstract level of pointer validity

In the particular setting described above, it is possible to introduce the following logic functions:

```

1  /*@
2  @ logic integer \offset_min<a>(a *p);
3  @ logic integer \offset_max<a>(a *p);
4  @/

```

- $\text{\offset_min}(p)$ is the minimum integer i such that $(p+i)$ is a valid pointer.
- $\text{\offset_max}(p)$: the maximum integer i such that $(p+i)$ is a valid pointer

The following properties hold:

```

1  \offset_min(p+i) == \offset_min(p)-i
2  \offset_max(p+i) == \offset_max(p)-i
3  \end{eqnarray*}
4  It also introduce syntactic sugar:
5  \begin{listing}{1}
6  /*@
7  predicate \valid_range<a>(a *p, integer i, integer j) =
8  \offset_min(p) <= i && \offset_max(p) >= j;
9  */

```

and the ACSL built-in predicate $\text{\valid}(p)$ is now equivalent to $\text{\validrange}(p,0,0)$.

3.2.2 Strings

EXPERIMENTAL **The logic function**

```

| /*@ logic integer \strlen(char* p);

```

denotes the length of a 0-terminated C string. It is total function, whose value is non-negative if and only if the pointer in argument is really a string.

Example 3.1 *Here is a contract for the strcpy function:*

```

1  /*@ // src and dest cannot overlap
2  @ requires \base_addr(src) != \base_addr(dest);
3  @ // src is a valid C string
4  @ requires \strlen(src) >= 0 ;
5  @ // dest is large enough to store a copy of src up to the 0
6  @ requires \valid_range(dest,0,\strlen(src));
7  @ ensures
8  @ \forall integer k; 0 <= k <= \strlen(src) ==> dest[k] == src[k]
9  @*/
10 char* strcpy(char *dest, const char *src);

```

3.3 Memory leaks

EXPERIMENTAL

Verification of absence of memory leak is outside the scope of the specification language. On the other hand, various models could be set up, using for example ghost variables.



Chapter 4

Conclusion

This document presents a Behavioral Interface Specification Language for ANSI C source code. It provides a common basis that could be shared among several tools. The specification language described here is intended to evolve in the future, and remain open to additional constructions. One interesting possible extension regards “temporal” properties in a large sense, such as liveness properties, which can sometimes be simulated by regular specifications with ghost variables [12], or properties on evolution of data over the time, such as the history constraints of JML, or in the Lustre assertion language.



Appendix A

Appendices

A.1 Glossary

pure expressions In ACSL setting, a *pure* expression is a C expression which contains no assignments, no incrementation operator `++` or `--`, no function call, and no access to a volatile object. The set of pure expression is a subset of the set of C expressions without side effect (C standard [14, 13], §5.1.2.3, alinea 2).

left-values A *left-value* (*lvalue* for short) is an expression which denotes some place in the memory during program execution, either on the stack, on the heap, or in the static data segment. It can be either a variable identifier or an expression of the form `*e`, `e[e]`, `e.id` or `e->id`, where `e` is any expression and `id` a field name. See C standard, §6.3.2.1 for a more detailed description of lvalues.

A *modifiable lvalue* is an lvalue allowed in the left part of an assignment. In essence, all lvalues are modifiable except variables declared as `const` or of some array type with explicit length.

pre-state and post-state For a given function call, the *pre-state* denotes the program state at the beginning of the call, including the current values for the function parameters. The *post-state* denotes the program state at the return of the call.

function behavior A *function behavior* (*behavior* for short) is a set of properties relating the pre-state and the post-state for a possibly restricted set of pre-states (*behavior assumptions*).

function contract A *function contract* (*contract* for short) forms a specification of a function, consisting of the combination of a precondition (a requirement on the pre-state for any caller to that function), a collection of behaviors, and possibly a measure in case of a recursive function.

A.2 Comparison with JML

Although we took our inspiration in the Java Modeling Language (aka JML [16]), ACSL is notably different from JML in two crucial aspects:

- ACSL is a BISO for C, a low-level structured language, while JML is a BISO for Java, an object-oriented inheritance-based high-level language. Not only the language features

are not the same but the programming styles and idioms are very different, which entails also different ways of specifying behaviors. In particular, C has no inheritance nor exceptions, and no language support for the simplest properties on memory (*e.g.*, the size of an allocated memory block).

- JML relies on runtime assertion checking (RAC) when typing, static analysis and automatic deductive verification fail. The example of CCured [22, 7], that adds strong typing to C by relying on RAC too, shows that it is not possible to do it in a modular way. Indeed, it is necessary to modify the layout of C data structures for RAC, which is not modular. The follow-up project Deputy [8] thus reduces the checking power of annotations in order to preserve modularity. On the contrary, we choose not to restrain the power of annotations (*e.g.*, all first order logic formulas are allowed). To that end, we rely on manual deductive verification using an interactive theorem prover (*e.g.*, Coq) when every other technique failed.

In the remainder of this chapter, we describe these differences in further details.

A.2.1 Low-level language vs. inheritance-based one

No inherited specifications

JML has a core notion of inheritance of specifications, that duplicates in specifications the inheritance feature of Java. Inheritance combined with visibility and modularity account for a number of complex features in JML (*e.g.*, `spec_public` modifier, data groups, represents clauses, etc), that are necessary to express the desired inheritance-related specifications while respecting visibility and modularity. Since C has no inheritance, these intricacies are avoided in ACSL.

Error handling without exceptions

The usual way of signaling errors in Java is through exceptions. Therefore, JML specifications are tailored to express exceptional postconditions, depending on the exception raised. Since C has no exceptions, ACSL does not use exceptional specifications. Instead, C programmers are used to signal errors by returning special values, like mandated in various ways in the C standard.

Example A.1 *In §7.12.1 of the standard, it is said that functions in `<math.h>` signal errors as follows: “On a domain error, [...] the integer expression `errno` acquires the value `EDOM`.”*

Example A.2 *In §7.19.5.1 of the standard, it is said that function `fclose` signals errors as follows: “The `fclose` function returns [...] `EOF` if any errors were detected.”*

Example A.3 *In §7.19.6.1 of the standard, it is said that function `fprintf` signals errors as follows: “The `fprintf` function returns [...] a negative value if an output or encoding error occurred.”*

Example A.4 *In §7.20.3 of the standard, it is said that memory management functions signal errors as follows: “If the space cannot be allocated, a null pointer is returned.”*

As shown by these few examples, there is no unique way to signal errors in the C standard library, not mentioning user-defined functions. But since errors are signaled by returning special values, it is sufficient to write an appropriate postcondition:

```
| /*@ ensures \result == error_value || normal_postcondition; */
```

C contracts are not Java ones

In Java, the precondition of the following function that nullifies an array of characters is always true. Even if there was a precondition on the length of array `a`, it could easily be expressed using the Java expression `a.length` that gives the dynamic length of array `a`.

```
1 public static void Java_nullify(char[] a) {
2     if (a == null) return;
3     for (int i = 0; i < a.length; ++i) {
4         a[i] = 0;
5     }
6 }
```

On the contrary, the precondition of the same function in C, whose definition follows, is more involved. First, remark that the C programmer has to add an extra argument for the size of the array, or rather a lower bound on this array size.

```
1 void C_nullify(char* a, unsigned int n) {
2     int i;
3     if (n == 0) return;
4     for (i = 0; i < n; ++i) {
5         a[i] = 0;
6     }
7 }
```

A correct precondition for this function is the following:

```
| /*@ requires \valid(a + 0..(n-1)); */
```

where predicate `\valid` is the one defined in Section 2.7.1. (note that `\valid(a + 0..(-1))` is the same as `\valid(\empty)` and thus is true regardless of the validity of `a` itself). When `n` is null, `a` does not need to be valid at all, and when `n` is strictly positive, `a` must point to an array of size at least `n`. To make it more obvious, the C programmer adopted a defensive programming style, which returns immediately when `n` is null. We can duplicate this in the specification:

```
| /*@ requires n == 0 || \valid(a + 0..(n-1)); */
```

Usually, many memory requirements are only necessary for some paths through the function, which correspond to some particular behaviors, selected according to some tests performed along the corresponding paths. Since C has no memory primitives, these tests involve other variables that the C programmer added to track additional information, like `n` in our example.

To make it easier, it is possible in ACSL to distinguish between the `assumes` part of a behavior, that specifies the tests that need to succeed for this behavior to apply, and the `requires` part that specifies the additional preconditions that must be true when a behavior applies. The specification for our example can then be translated into:

```
1 /*@ behavior n_is_null:
2     @ assumes n == 0;
3     @ behavior n_is_not_null:
4     @ assumes n > 0;
5     @ requires \valid(a + 0..(n-1));
6     @*/
```

This is equivalent to the previous requirement, except here behaviors can be completed with postconditions that belong to one behavior only. Contrary to JML, the set of behaviors for a function do not necessarily cover all cases of use for this function, as mentioned in Section 2.3.3. This allows for partial specifications, whereas JML behaviors cannot offer such flexibility. Here, Our two behaviors are clearly mutually exclusive, and, since `n` is an unsigned int, our they cover all the possible cases. We could have specified that as well, by adding the following lines in the contract (see Section 2.3.3).

```

1 | @ ...
2 | @ disjoint behaviors;
3 | @ complete behaviors;
4 | @*/

```

ACSL contracts vs. JML ones

To fully understand the difference between specifications in ACSL and JML, we detail in below the requirement on the pre-state and the guarantee on the post-state given by behaviors in JML and ACSL.

A JML contract is either *lightweight* or *heavyweight*. For the purpose of our comparison, it is sufficient to know that a lightweight contract has `requires` and `ensures` clauses all at the same level, while an heavyweight contract has multiple behaviors, each consisting of `requires` and `ensures` clauses. Although it is not possible in JML to mix both styles, we can define here what it would mean to have both, by conjoining the conditions on the pre- and the post-state. Here is an hypothetical JML contract mixing lightweight and heavyweight styles:

```

1 | /*@ requires P1;
2 | @ requires P2;
3 | @ ensures Q1;
4 | @ ensures Q2;
5 | @ behavior x1:
6 | @   requires A1;
7 | @   requires R1;
8 | @   ensures E1;
9 | @ behavior x2:
10 | @   requires A2;
11 | @   requires R2;
12 | @   ensures E2;
13 | @*/

```

It assumes from the pre-state the condition:

```

| P1 && P2 && ((A1 && R1) || (A2 && R2))

```

and guarantees that the following condition holds in post-state:

```

| Q1 && Q2 &&
| (\old(A1 && R1) ==> E1) && (\old(A2 && R2) ==> E2)

```

Here is now an ACSL specification:

```

1 | /*@ requires P1;
2 | @ requires P2;
3 | @ ensures Q1;
4 | @ ensures Q2;
5 | @ behavior x1:
6 | @   assumes A1;
7 | @   requires R1;
8 | @   ensures E1;
9 | @ behavior x2:
10 | @   assumes A2;
11 | @   requires R2;
12 | @   ensures E2;
13 | @*/

```

Syntactically, the only difference with the JML specification is the addition of the `assumes` clauses. Its translation to assume-guarantee is however quite different. It assumes from the pre-state the condition:

```
| P1 && P2 && (A1 ==> R1) && (A2 ==> R2)
```

and guarantees that the following condition holds in the post-state:

```
| Q1 && Q2 && (\old(A1) ==> E1) && (\old(A2) ==> E2)
```

Thus, ACSL allows to distinguish between the clauses that control which behavior is active (the `assumes` clauses) and the clauses that are preconditions for a particular behavior (the internal `requires` clauses). In addition, as mentioned above, there is by default no requirement in ACSL for the specification to be complete (The last part of the JML condition on the pre-state). If desired, this has to be precised explicitly with a `complete behaviors` clause as seen in Section 2.3.3.

A.2.2 Deductive verification vs. RAC

Sugar-free behaviors

As explained in details in [23], JML heavyweight behaviors can be viewed as syntactic sugar (however complex it is) that can be translated automatically into more basic contracts consisting mostly of pre- and postconditions and frame conditions. This allows complex nesting of behaviors from the user point of view, while tools only have to deal with basic contracts. In particular, the major tools on JML use this desugaring process, like the Common JML tools to do assertion checking, unit testing, etc. (see [19]) and the tool ESC/Java2 for automatic deductive verification of JML specifications (see [6]).

One issue with such a desugaring approach is the complexity of the transformations involved, as *e.g.* for desugaring assignable clauses between multiple *spec-cases* in JML [23]. Another issue is precisely that tools only see one global contract, instead of multiple independent behaviors, that could be analyzed separately in more detail. Instead, we favor the view that a function implements multiple behaviors, that can be analyzed separately if a tool feels like it. Therefore, we do not intend to provide a desugaring process.

Axiomatized functions in specifications

JML only allows pure Java methods to be called in specifications [17]. This is certainly understandable when relying on RAC: methods called should be defined so that the runtime can call them, and they should not have side-effects in order not to pollute the program they are supposed to annotate.

In our setting, it is desirable to allow calls to logical functions in specifications. These functions may be defined, like program ones, but they may also be only declared (with a suitable declaration of `reads` clause) and defined through an axiomatization. This makes for richer specifications that may be useful either in automatic or in manual deductive verification.

A.2.3 Syntactic differences

The following table summarizes the difference between JML keywords and ACSL ones, when the intent is the same, although minor differences might exist.

JML	ACSL
modifiable,assignable	assigns
measured_by	decreases
loop_invariant	loop invariant
decreases	loop variant
(\forall x ; P ; Q)	(\forall x ; P ==> Q)
(\exists x ; P ; Q)	(\exists x ; P && Q)
\max x ; a <= x <= b ; f)	\max(a,b,\lambda x ; f)

A.3 Typing rules

Disclaimer: this section is unfinished, it is left here just to give an idea of what it will look like in the final document.

A.3.1 Rules for terms

Integer promotion:

$$\frac{\Gamma \vdash e : \tau}{\Gamma \vdash e : \text{integer}}$$

if τ is any C integer type `char`, `short`, `int`, or `long`, whatever attribute they have, in particular signed or unsigned

Variables:

$$\frac{}{\Gamma \vdash id : \tau} \text{ if } id : \tau \in \Gamma$$

Unary integer operations:

$$\frac{\Gamma \vdash t : \text{integer}}{\Gamma \vdash op\ t : \text{integer}} \text{ if } op \in \{+, -, \sim\}$$

Boolean negation:

$$\frac{\Gamma \vdash t : \text{boolean}}{\Gamma \vdash !t : \text{boolean}}$$

Pointer dereferencing:

$$\frac{\Gamma \vdash t : \tau*}{\Gamma \vdash *t : \tau}$$

Address operator:

$$\frac{\Gamma \vdash t : \tau}{\Gamma \vdash \&t : \tau*}$$

Binary

$$\frac{\Gamma \vdash t_1 : \text{integer} \quad \Gamma \vdash t_2 : \text{integer}}{\Gamma \vdash t_1\ op\ t_2 : \text{integer}} \text{ if } op \in \{+, -, *, /, \%\}$$

$$\frac{\Gamma \vdash t_1 : \text{real} \quad \Gamma \vdash t_2 : \text{real}}{\Gamma \vdash t_1\ op\ t_2 : \text{real}} \text{ if } op \in \{+, -, *, /\}$$

$$\frac{\Gamma \vdash t_1 : \text{integer} \quad \Gamma \vdash t_2 : \text{integer}}{\Gamma \vdash t_1\ op\ t_2 : \text{boolean}} \text{ if } op \in \{==, !=, <=, <, >=, >\}$$

$$\frac{\Gamma \vdash t_1 : \text{real} \quad \Gamma \vdash t_2 : \text{real}}{\Gamma \vdash t_1\ op\ t_2 : \text{boolean}} \text{ if } op \in \{==, !=, <=, <, >=, >\}$$

$$\frac{\Gamma \vdash t_1 : \tau^* \quad \Gamma \vdash t_2 : \tau^*}{\Gamma \vdash t_1 \text{ op } t_2 : \text{boolean}} \text{ if } \text{op} \in \{==, !=, <=, <, >=, >\}$$

(to be continued)

A.3.2 Typing rules for sets

We consider the typing judgement $\Gamma, \Lambda \vdash s : \tau, b$ meaning that s is a set of terms of type τ , which is moreover a set of locations if the boolean b is true. Γ is the C environment and Λ is the logic environment.

Rules:

$$\frac{}{\Gamma, \Lambda \vdash id : \tau, true} \text{ if } id : \tau \in \Gamma$$

$$\frac{}{\Gamma, \Lambda \vdash id : \tau, true} \text{ if } id : \tau \in \Lambda$$

$$\frac{\Gamma, \Lambda \vdash s : \tau^*, b}{\Gamma, \Lambda \vdash *s : \tau, true}$$

$$\frac{id : \tau \quad s : \text{set} < \text{struct } S^* >}{\vdash s - > id : \text{set} < \tau >}$$

$$\frac{\Gamma, b \cup \Lambda \vdash e : tset\tau}{\Gamma, \Lambda \vdash \{e \mid b; P\} : tset\tau}$$

$$\frac{\Gamma, \Lambda \vdash e_1 : \tau, b \quad \Gamma, \Lambda \vdash e_2 : \tau, b}{\Gamma, \Lambda \vdash e_1, e_2 : \tau, b}$$

A.4 Specification Templates

This section describes some common issues that may occur when writing an ACSL specification and proposes some solution to overcome them

A.4.1 Accessing a C variable that is masked

The situation may happen where it is necessary to refer in an annotation to a C variable that is masked at that point. For instance, a function contract may need to refer to a global variable that has the same name as a function parameter, as in the following code:

```

1 | int x;
2 | //@ assigns x;
3 | int g();
4 |
5 | int f(int x) {
6 |     // ...
7 |     return g();
8 | }
```

In order to write the `assigns` clause for `f`, we must access the global variable `x`, since `f` calls `g`, which can modify `x`. This is not possible with C scoping rules, as `x` refers to the parameter of `f` in the scope of the function.

A solution is to use a ghost pointer to `x`, as shown in the following code:

```

1 int x;
2
3 //@ ghost int* const ghost_ptr_x = &x;
4
5 //@ assigns x;
6 int g();
7
8 //@ assigns *ghost_ptr_x;
9 int f(int x) {
10     // ...
11     return g();
12 }

```

A.5 Illustrative example

This is an attempt to define an example for ACSL, much as the Purse example in JML description papers. It is a memory allocator, whose main functions are `memory_alloc` and `memory_free`, to respectively allocate and deallocate memory. The goal is to exercise as much as possible of ACSL.

```

1
2 #include <stdlib.h>
3
4 #define DEFAULT_BLOCK_SIZE 1000
5
6 typedef enum _bool { false = 0, true = 1 } bool;
7
8 /* predicate finite_list<A>((A* -> A*) next_elem, A* ptr) =
9  @ ptr == \null ||
10  @ (\valid(ptr) && finite_list(next_elem,next_elem(ptr))) ;
11  @
12  @ logic integer list_length<A>((A* -> A*) next_elem, A* ptr) =
13  @ (ptr == \null) ? 0 :
14  @ 1 + list_length(next_elem,next_elem(ptr)) ;
15  @
16  @
17  @ predicate lower_length<A>((A* -> A*) next_elem,
18  @ A* ptr1, A* ptr2) =
19  @ finite_list(next_elem, ptr1) && finite_list(next_elem, ptr2)
20  @ && list_length(next_elem, ptr1) < list_length(next_elem, ptr2) ;
21  @*/
22
23 // forward reference
24 struct _memory_slice;
25
26 /* A memory block holds a pointer to a raw block of memory allocated by
27 * calling [malloc]. It is sliced into chunks, which are maintained by
28 * the [slice] structure. It maintains additional information such as
29 * the [size] of the memory block, the number of bytes [used] and the [next]
30 * index at which to put a chunk.
31 */
32 typedef struct _memory_block {
33     //@ ghost boolean packed;
34     // ghost field [packed] is meant to be used as a guard that tells when
35     // the invariant of a structure of type [memory_block] holds
36     unsigned int size;
37     // size of the array [data]
38     unsigned int next;
39     // next index in [data] at which to put a chunk
40     unsigned int used;
41     // how many bytes are used in [data], not necessarily contiguous ones
42     char* data;
43     // raw memory block allocated by [malloc]
44     struct _memory_slice* slice;
45     // structure that describes the slicing of a block into chunks
46 } memory_block;

```

A.5. ILLUSTRATIVE EXAMPLE

```

47
48 /*@ strong type invariant inv_memory_block(memory_block mb) =
49 @   mb.packed ==>
50 @   (0 < mb.size && mb.used <= mb.next <= mb.size
51 @   && \offset(mb.data) == 0
52 @   && \block_length(mb.data) == mb.size) ;
53 @
54 @ predicate valid_memory_block(memory_block* mb) =
55 @   \valid(mb) && mb->packed ;
56 @*/
57
58 /* A memory chunk holds a pointer [data] to some part of a memory block
59 * [block]. It maintains the [offset] at which it points in the block, as well
60 * as the [size] of the block it is allowed to access. A field [free] tells
61 * whether the chunk is used or not.
62 */
63 typedef struct _memory_chunk {
64   //@ ghost boolean packed;
65   // ghost field [packed] is meant to be used as a guard that tells when
66   // the invariant of a structure of type [memory_chunk] holds
67   unsigned int   offset;
68   // offset at which [data] points into [block->data]
69   unsigned int   size;
70   // size of the chunk
71   bool           free;
72   // true if the chunk is not used, false otherwise
73   memory_block*  block;
74   // block of memory into which the chunk points
75   char*          data;
76   // shortcut for [block->data + offset]
77 } memory_chunk;
78
79 /*@ strong type invariant inv_memory_chunk(memory_chunk mc) =
80 @   mc.packed ==>
81 @   (0 < mc.size && valid_memory_block(mc.block)
82 @   && mc.offset + mc.size <= mc.block->next) ;
83 @
84 @ predicate valid_memory_chunk(memory_chunk* mc, int s) =
85 @   \valid(mc) && mc->packed && mc->size == s ;
86 @
87 @ predicate used_memory_chunk(memory_chunk mc) =
88 @   mc.free == false ;
89 @
90 @ predicate freed_memory_chunk(memory_chunk mc) =
91 @   mc.free == true ;
92 @*/
93
94 /* A memory chunk list links memory chunks in the same memory block.
95 * Newly allocated chunks are put first, so that the offset of chunks
96 * decreases when following the [next] pointer. Allocated chunks should
97 * fill the memory block up to its own [next] index.
98 */
99 typedef struct _memory_chunk_list {
100   memory_chunk*  chunk;
101   // current list element
102   struct _memory_chunk_list* next;
103   // tail of the list
104 } memory_chunk_list;
105
106 /*@ logic memory_chunk_list* next_chunk(memory_chunk_list* ptr) =
107 @   ptr->next ;
108 @
109 @ predicate valid_memory_chunk_list
110 @   (memory_chunk_list* mcl, memory_block* mb) =
111 @   \valid(mcl) && valid_memory_chunk(mcl->chunk, mcl->chunk->size)
112 @   && mcl->chunk->block == mb
113 @   && (mcl->next == \null ||
114 @   valid_memory_chunk_list(mcl->next, mb))
115 @   && mcl->offset == mcl->chunk->offset
116 @   && (
117 @   // it is the last chunk in the list

```

APPENDIX A. APPENDICES

```

118 | @      (mcl->next == \null && mcl->chunk->offset == 0)
119 | @      ||
120 | @      // it is a chunk in the middle of the list
121 | @      (mcl->next != \null
122 | @      && mcl->next->chunk->offset + mcl->next->chunk->size
123 | @      == mcl->chunk->offset)
124 | @      )
125 | @      && finite_list(next_chunk, mcl) ;
126 | @
127 | @ predicate valid_complete_chunk_list
128 | @      (memory_chunk_list* mcl, memory_block* mb) =
129 | @      valid_memory_chunk_list(mcl, mb)
130 | @      && mcl->next->chunk->offset +
131 | @      mcl->next->chunk->size == mb->next ;
132 | @
133 | @ predicate chunk_lower_length(memory_chunk_list* ptr1,
134 | @      memory_chunk_list* ptr2) =
135 | @      lower_length(next_chunk, ptr1, ptr2) ;
136 | @*/
137 |
138 | /* A memory slice holds together a memory block [block] and a list of chunks
139 | * [chunks] on this memory block.
140 | */
141 | typedef struct _memory_slice {
142 |     //@ ghost boolean      packed;
143 |     // ghost field [packed] is meant to be used as a guard that tells when
144 |     // the invariant of a structure of type [memory_slice] holds
145 |     memory_block*      block;
146 |     memory_chunk_list* chunks;
147 | } memory_slice;
148 |
149 | /*@ strong type invariant inv_memory_slice(memory_slice* ms) =
150 | @      ms.packed ==>
151 | @      (valid_memory_block(ms->block) && ms->block->slice == ms
152 | @      && (ms->chunks == \null
153 | @      || valid_complete_chunk_list(ms->chunks, ms->block))) ;
154 | @
155 | @ predicate valid_memory_slice(memory_slice* ms) =
156 | @      \valid(ms) && ms->packed ;
157 | @*/
158 |
159 | /* A memory slice list links memory slices, to form a memory pool.
160 | */
161 | typedef struct _memory_slice_list {
162 |     //@ ghost boolean      packed;
163 |     // ghost field [packed] is meant to be used as a guard that tells when
164 |     // the invariant of a structure of type [memory_slice_list] holds
165 |     memory_slice*      slice;
166 |     // current list element
167 |     struct _memory_slice_list* next;
168 |     // tail of the list
169 | } memory_slice_list;
170 |
171 | /*@ logic memory_slice_list* next_slice(memory_slice_list* ptr) =
172 | @      ptr->next ;
173 | @
174 | @ strong type invariant inv_memory_slice_list(memory_slice_list* msl) =
175 | @      msl.packed ==>
176 | @      (valid_memory_slice(msl->slice)
177 | @      && (msl->next == \null ||
178 | @      valid_memory_slice_list(msl->next))
179 | @      && finite_list(next_slice, msl)) ;
180 | @
181 | @ predicate valid_memory_slice_list(memory_slice_list* msl) =
182 | @      \valid(msl) && msl->packed ;
183 | @
184 | @ predicate slice_lower_length(memory_slice_list* ptr1,
185 | @      memory_slice_list* ptr2) =
186 | @      lower_length(next_slice, ptr1, ptr2)
187 | @ } */
188 |

```

A.5. ILLUSTRATIVE EXAMPLE

```

189 typedef memory_slice_list* memory_pool;
190
191 /*@ type invariant valid_memory_pool(memory_pool *mp) =
192    @ \valid(mp) && valid_memory_slice_list(*mp) ;
193    @*/
194
195 /*@ behavior zero_size:
196    @ assumes s == 0;
197    @ assigns \nothing;
198    @ ensures \result == 0;
199    @
200    @ behavior positive_size:
201    @ assumes s > 0;
202    @ requires valid_memory_pool(arena);
203    @ ensures \result == 0
204    @ || (valid_memory_chunk(\result,s) &&
205    @      used_memory_chunk(*\result));
206    @ */
207 memory_chunk* memory_alloc(memory_pool* arena, unsigned int s) {
208     memory_slice_list *msl = *arena;
209     memory_chunk_list *mcl;
210     memory_slice *ms;
211     memory_block *mb;
212     memory_chunk *mc;
213     unsigned int mb_size;
214     /*@ ghost unsigned int mcl_offset;
215     char *mb_data;
216     // guard condition
217     if (s == 0) return 0;
218     // iterate through memory blocks (or slices)
219     /*@
220     @ loop invariant valid_memory_slice_list(msl);
221     @ loop variant msl for slice_lower_length;
222     @ */
223     while (msl != 0) {
224         ms = msl->slice;
225         mb = ms->block;
226         mcl = ms->chunks;
227         // does [mb] contain enough free space?
228         if (s <= mb->size - mb->next) {
229             /*@ ghost ms->ghost = false; // unpack the slice
230             // allocate a new chunk
231             mc = (memory_chunk*)malloc(sizeof(memory_chunk));
232             if (mc == 0) return 0;
233             mc->offset = mb->next;
234             mc->size = s;
235             mc->free = false;
236             mc->block = mb;
237             /*@ ghost mc->ghost = true; // pack the chunk
238             // update block accordingly
239             /*@ ghost mb->ghost = false; // unpack the block
240             mb->next += s;
241             mb->used += s;
242             /*@ ghost mb->ghost = true; // pack the block
243             // add the new chunk to the list
244             mcl = (memory_chunk_list*)malloc(sizeof(memory_chunk_list));
245             if (mcl == 0) return 0;
246             mcl->chunk = mc;
247             mcl->next = ms->chunks;
248             ms->chunks = mcl;
249             /*@ ghost ms->ghost = true; // pack the slice
250             return mc;
251         }
252         // iterate through memory chunks
253         /*@
254         @ loop invariant valid_memory_chunk_list(mcl,mb);
255         @ loop variant mcl for chunk_lower_length;
256         @ */
257         while (mcl != 0) {
258             mc = mcl->chunk;
259             // is [mc] free and large enough?

```

APPENDIX A. APPENDICES

```

260     if (mc->free && s <= mc->size) {
261         mc->free = false;
262         mb->used += mc->size;
263         return mc;
264     }
265     // try next chunk
266     mcl = mcl->next;
267 }
268     msl = msl->next;
269 }
270 // allocate a new block
271 mb_size = (DEFAULT_BLOCK_SIZE < s) ? s : DEFAULT_BLOCK_SIZE;
272 mb_data = (char*)malloc(mb_size);
273 if (mb_data == 0) return 0;
274 mb = (memory_block*)malloc(sizeof(memory_block));
275 if (mb == 0) return 0;
276 mb->size = mb_size;
277 mb->next = s;
278 mb->used = s;
279 mb->data = mb_data;
280 //@ ghost mb->ghost = true; // pack the block
281 // allocate a new chunk
282 mc = (memory_chunk*)malloc(sizeof(memory_chunk));
283 if (mc == 0) return 0;
284 mc->offset = 0;
285 mc->size = s;
286 mc->free = false;
287 mc->block = mb;
288 //@ ghost mc->ghost = true; // pack the chunk
289 // allocate a new chunk list
290 mcl = (memory_chunk_list*)malloc(sizeof(memory_chunk_list));
291 if (mcl == 0) return 0;
292 //@ ghost mcl->offset = 0;
293 mcl->chunk = mc;
294 mcl->next = 0;
295 // allocate a new slice
296 ms = (memory_slice*)malloc(sizeof(memory_slice));
297 if (ms == 0) return 0;
298 ms->block = mb;
299 ms->chunks = mcl;
300 //@ ghost ms->ghost = true; // pack the slice
301 // update the block accordingly
302 mb->slice = ms;
303 // add the new slice to the list
304 msl = (memory_slice_list*)malloc(sizeof(memory_slice_list));
305 if (msl == 0) return 0;
306 msl->slice = ms;
307 msl->next = *arena;
308 //@ ghost msl->ghost = true; // pack the slice list
309 *arena = msl;
310 return mc;
311 }
312
313 /*@ behavior null_chunk:
314  @ assumes chunk == \null;
315  @ assigns \nothing;
316  @
317  @ behavior valid_chunk:
318  @ assumes chunk != \null;
319  @ requires valid_memory_pool(arena);
320  @ requires valid_memory_chunk(chunk, chunk->size);
321  @ requires used_memory_chunk(chunk);
322  @ ensures
323  @     // if it is not the last chunk in the block, mark it as free
324  @     (valid_memory_chunk(chunk, chunk->size)
325  @     && freed_memory_chunk(chunk))
326  @     ||
327  @     // if it is the last chunk in the block, deallocate the block
328  @     ! \valid(chunk);
329  @ */
330 void memory_free(memory_pool* arena, memory_chunk* chunk) {

```

A.6. CHANGES

```
331     memory_slice_list *msl = *arena;
332     memory_block *mb = chunk->block;
333     memory_slice *ms = mb->slice;
334     memory_chunk_list *mcl;
335     memory_chunk *mc;
336     // is it the last chunk in use in the block?
337     if (mb->used == chunk->size) {
338         // remove the corresponding slice from the memory pool
339         // case it is the first slice
340         if (msl->slice == ms) {
341             *arena = msl->next;
342             //@ ghost msl->ghost = false;    // unpack the slice list
343             free(msl);
344         }
345         // case it is not the first slice
346         while (msl != 0) {
347             if (msl->next != 0 && msl->next->slice == ms) {
348                 memory_slice_list* msl_next = msl->next;
349                 msl->next = msl->next->next;
350                 // unpack the slice list
351                 //@ ghost msl_next->ghost = false;
352                 free(msl_next);
353                 break;
354             }
355             msl = msl->next;
356         }
357         //@ ghost ms->ghost = false;    // unpack the slice
358         // deallocate all chunks in the block
359         mcl = ms->chunks;
360         // iterate through memory chunks
361         /*@
362          @ loop invariant valid_memory_chunk_list(mcl,mb);
363          @ loop variant mcl for chunk_lower_length;
364          @ */
365         while (mcl != 0) {
366             memory_chunk_list *mcl_next = mcl->next;
367             mc = mcl->chunk;
368             //@ ghost mc->ghost = false;    // unpack the chunk
369             free(mc);
370             free(mcl);
371             mcl = mcl_next;
372         }
373         mb->next = 0;
374         mb->used = 0;
375         // deallocate the memory block and its data
376         //@ ghost mb->ghost = false;    // unpack the block
377         free(mb->data);
378         free(mb);
379         // deallocate the corresponding slice
380         free(ms);
381         return;
382     }
383     // mark the chunk as freed
384     chunk->free = true;
385     // update the block accordingly
386     mb->used -= chunk->size;
387     return;
388 }
```

A.6 Changes

A.6.1 Version 1.5

- Clarify the status of loop invariant in presence of break or side-effects in the loop test.
- Introduction of \with keyword for functional updates.

- Added `bnf` entry for completeness of function behaviors.
- Order of clauses in statement contracts is now fixed.
- Requires clauses are allowed before behaviors of statement contracts.
- Added explicit singleton construct for sets.
- Introduction of logical arrays.
- Operations over pointers and arrays have been precised.

A.6.2 Version 1.4

- Added UTF-8 counterparts for built-in types (`integer`, `real`, `boolean`).
- Fixed typos in the examples corresponding to features implemented in Frama-C.
- Order of clauses in function contracts is now fixed.
- Introduction of abrupt termination clauses.
- Introduction of `axiomatic` to gather predicates, logic functions, and their defining axioms.
- Added specification templates appendix for common specification issues.
- Use of sets as first-class term has been precised.
- Fixed semantics of predicate `\separated`.

A.6.3 Version 1.3

- Functional update of structures.
- Terminates clause in function behaviors.
- Typos reported by David Mentré.

A.6.4 Version 1.2

This is the first public release of this document.

Bibliography

- [1] Jean-Raymond Abrial. *The B-Book: Assigning Programs to Meanings*. Cambridge University Press, 1996.
- [2] Ali Ayad and Claude Marché. Behavioral properties of floating-point programs. Hisseo publications, 2009. <http://hisseo.saclay.inria.fr/ayad09.pdf>.
- [3] Sylvie Boldo and Jean-Christophe Filliâtre. Formal Verification of Floating-Point Programs. In *18th IEEE International Symposium on Computer Arithmetic*, pages 187–194, Montpellier, France, June 2007.
- [4] Patrice Chalin. Reassessing JML’s logical foundation. In *Proceedings of the 7th Workshop on Formal Techniques for Java-like Programs (FTfJP’05)*, Glasgow, Scotland, July 2005.
- [5] Patrice Chalin. A sound assertion semantics for the dependable systems evolution verifying compiler. In *Proceedings of the International Conference on Software Engineering (ICSE’07)*, pages 23–33, Los Alamitos, CA, USA, 2007. IEEE Computer Society.
- [6] David R. Cok and Joseph R. Kiniry. ESC/Java2 implementation notes. Technical report, may 2007. <http://secure.ucd.ie/products/opensource/ESCJava2/ESCTools/docs/Escjava2-ImplementationNotes.pdf>.
- [7] Jeremy Condit, Matthew Harren, Scott McPeak, George C. Necula, and Westley Weimer. Ccured in the real world. In *PLDI ’03: Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation*, pages 232–244, 2003.
- [8] Jeremy Paul Condit, Matthew Thomas Harren, Zachary Ryan Anderson, David Gay, and George Necula. Dependent types for low-level programming. In *ESOP ’07: Proceedings of the 16th European Symposium on Programming*, Oct 2006.
- [9] Jean-Christophe Filliâtre and Claude Marché. Multi-prover verification of C programs. In Jim Davies, Wolfram Schulte, and Mike Barnett, editors, *6th International Conference on Formal Engineering Methods*, volume 3308 of *Lecture Notes in Computer Science*, pages 15–29, Seattle, WA, USA, November 2004. Springer.
- [10] Jean-Christophe Filliâtre and Claude Marché. The Why/Krakatoa/Caduceus platform for deductive program verification. In Werner Damm and Holger Hermanns, editors, *19th International Conference on Computer Aided Verification*, Lecture Notes in Computer Science, Berlin, Germany, July 2007. Springer.
- [11] The Frama-C framework for analysis of C code. <http://frama-c.cea.fr/>.
- [12] A. Giorgetti and J. Gros Lambert. JAG: JML Annotation Generation for verifying temporal properties. In *FASE’2006, Fundamental Approaches to Software Engineering*, volume 3922 of *LNCS*, pages 373–376, Vienna, Austria, March 2006. Springer.

BIBLIOGRAPHY

- [13] International Organization for Standardization (ISO). *The ANSI C standard (C99)*. <http://www.open-std.org/JTC1/SC22/WG14/www/docs/n1124.pdf>.
- [14] Brian Kernighan and Dennis Ritchie. *The C Programming Language (2nd Ed.)*. Prentice-Hall, 1988.
- [15] Joseph Kiniry. ESC/Java2. <http://kind.ucd.ie/products/opensource/ESCJava2/>.
- [16] Gary Leavens. Jml. <http://www.eecs.ucf.edu/~leavens/JML/>.
- [17] Gary T. Leavens, Albert L. Baker, and Clyde Ruby. Preliminary design of JML: A behavioral interface specification language for Java. Technical Report 98-06i, Iowa State University, 2000.
- [18] Gary T. Leavens, K. Rustan M. Leino, and Peter Müller. Specification and verification challenges for sequential object-oriented programs. *Form. Asp. Comput.*, 19(2):159–189, 2007.
- [19] Gary T. Leavens, K. Rustan M. Leino, Erik Poll, Clyde Ruby, and Bart Jacobs. JML: notations and tools supporting detailed design in Java. In *OOPSLA 2000 Companion, Minneapolis, Minnesota*, pages 105–106, 2000.
- [20] Claude Marché. Towards modular algebraic specifications for pointer programs: a case study. In H. Comon-Lundh, C. Kirchner, and H. Kirchner, editors, *Rewriting, Computation and Proof*, volume 4600 of *Lecture Notes in Computer Science*, pages 235–258. Springer-Verlag, 2007.
- [21] Yannick Moy. Union and cast in deductive verification. Technical Report ICIS-R07015, Radboud University Nijmegen, jul 2007. http://www.lri.fr/~moy/union_and_cast/union_and_cast.pdf.
- [22] George C. Necula, Scott McPeak, and Westley Weimer. CCured: Type-safe retrofitting of legacy code. In *Symposium on Principles of Programming Languages*, pages 128–139, 2002.
- [23] Arun D. Raghavan and Gary T. Leavens. Desugaring JML method specifications. Technical Report 00-03a, Iowa State University, 2000.
- [24] David Stevenson et al. An american national standard: IEEE standard for binary floating point arithmetic. *ACM SIGPLAN Notices*, 22(2):9–25, 1987.
- [25] Wikipedia. First order logic. http://en.wikipedia.org/wiki/First_order_logic.
- [26] Wikipedia. IEEE 754. http://en.wikipedia.org/wiki/IEEE_754-1985.

List of Figures

2.1	Grammar of terms	16
2.2	Grammar of predicates	17
2.3	Grammar of binders and type expressions	18
2.4	Operator precedence	19
2.5	Grammar of function contracts	28
2.6	\old and \result in terms	28
2.7	Grammar for sets of terms	33
2.8	Grammar for assertions	35
2.9	Grammar for loop annotations	36
2.10	Grammar for general inductive invariants	38
2.11	Grammar for statement contracts	42
2.12	Grammar for global logic definitions	45
2.13	Grammar for inductive definitions	46
2.14	Grammar for axiomatic declarations	47
2.15	Grammar for higher-order constructs	49
2.16	Grammar for concrete logic types and pattern-matching	50
2.17	Grammar for logic declarations with labels	51
2.18	Grammar for logic declarations with reads clauses	52
2.19	Grammar of contracts about abrupt terminations	56
2.20	Grammar for dependencies information	57
2.21	Grammar for declarations of data invariants	59
2.22	Grammar for declarations of model variables and fields	61
2.23	Grammar for ghost statements	64
2.24	Grammar for volatile constructs	65



Index

- ?, 16, 17
- _, 50
- abrupt clause, 55
- annotation, 35
 - loop, 35
- as, 50
- assert , 35
- assertion, 35
- assigns , 28, 29, 36, 57
- assumes, 28
- \at, 40
- axiom, 47
- axiomatic, 47
- axiomatic, 47
- \base_addr, 54
- behavior, 30, 75
- behavior , 28, 42
- behaviors , 28
- \block_length, 28, 54
- boolean, 18, 20
- breaks, 56
- case, 46, 50
- cast, 20–23
- complete, 28
- complete behaviors, 32
- comprehension, 34
- continues , 56
- contract, 28, 35, 41, 75
- data invariant, 58
- decreases , 28
- \decreases , 42
- dependency, 57
 - disjoint , 28
- disjoint behaviors, 32
- do, 36
- else , 64
- \empty, 33
- ensures , 28, 29, 41, 55
- \eq_double, 23
- \eq_float, 23
- \exists , 17
- \exit_status , 56
- \exit_status , 57
- exits, 56
- \false , 16, 17, 20
- for , 28, 35, 36, 38, 42
- \forall , 17
- formal parameter, 29, 30, 40, 59
- \freed , 54
- \fresh , 54
- \from, 57
- function behavior, 30, 75
- function contract, 28, 75
- functional expression, 57
- \ge_double, 23
- \ge_float, 23
- ghost, 62
- ghost, 64
- global , 59
- global invariant, 58
- grammar entries
 - C-global-decl*, 45
 - abrupt-clauses-fn*, 56
 - abrupt-clauses-stmt*, 56
 - assertion*, 35, 38
 - assigns-clause*, 28, 57
 - assumes-clause*, 28
 - axiom-decl*, 47
 - axiomatic-decl*, 47
 - behavior-body-stmt*, 42
 - behavior-body*, 28
 - bin-op*, 16
 - binders*, 18
 - binder*, 18
 - breaks-clause*, 56
 - built-in-logic-type*, 18

- completeness-clause*, 28
- compound-statement*, 35
- constructor*, 50
- continues-clause*, 56
- data-inv-decl*, 59
- data-invariant*, 59
- declaration*, 59, 61, 64, 65
- decreases-clause*, 28
- direct-declarator*, 64
- ensures-clause*, 28
- exits-clause*, 56
- ext-quantifier*, 49
- function-contract*, 28
- ghost-selection-statement*, 64
- ghost-type-specifier*, 64
- indcase*, 46
- inductive-def*, 46
- inv-strength*, 59
- label-binders*, 51
- lemma-decl*, 45
- locations*, 28
- logic-const-decl*, 47
- logic-const-def*, 45
- logic-decl*, 47
- logic-def*, 45–47, 50
- logic-function-decl*, 47, 52
- logic-function-def*, 45, 52
- logic-predicate-decl*, 47, 52
- logic-predicate-def*, 52
- logic-type-decl*, 47
- logic-type-def*, 50
- logic-type-expr*, 18
- logic-type*, 47
- loop-annot*, 36
- loop-assigns*, 36
- loop-behavior*, 36
- loop-clause*, 36
- loop-invariant*, 36
- loop-variant*, 36
- match-cases*, 50
- match-case*, 50
- named-behavior-stmt*, 42
- named-behavior*, 28
- parameters*, 45
- parameter*, 45
- pat*, 50
- poly-id*, 45, 51
- postfix-expression*, 64
- predicate-def*, 45
- pred*, 17, 28, 33
- reads-clause*, 52
- record-type*, 50
- rel-op*, 17
- requires-clause*, 28
- returns-clause*, 56
- simple-clause-stmt*, 42
- simple-clause*, 28
- statement-contract*, 42
- statements-ghost*, 64
- statement*, 35, 36, 42, 64
- struct-declaration*, 61, 64
- sum-type*, 50
- terminates-clause*, 28
- term*, 16, 28, 49, 50, 56
- tset*, 33
- type-expr*, 18, 45, 50
- type-invariant*, 59
- type-var-binders*, 45
- type-var*, 45
- unary-op*, 16
- variable-ident*, 18
- `\gt_double`, 23
- `\gt_float`, 23
- Here, 40, 56
- hybrid
 - function, 51
 - predicate, 51
- if, 64
- inductive, 46
- inductive definitions, 46
- inductive predicates, 46
 - `\initialized`, 66
- integer, 18, 20
- `\inter`, 33
- invariant, 35, 38
 - data, 58
 - global, 58
 - strong, 58
 - type, 58
 - weak, 58
- invariant, 36, 38, 59
- `\is_finite`, 24
- `\is_NaN`, 24
- l-value, 33
- `\lambda`, 49
- `\le_double`, 23
- `\le_float`, 23

- left-value, 75
- lemma, 45
- \let, 16, 17, 50
- library, 53
- location, 55
- logic, 45, 47, 52
- logic specification, 45
- loop, 36
- loop annotation, 35
- \lt_double, 23
- \lt_float, 23
- lvalue, 75
- \match, 50
- \max, 49
- \min, 49
- model, 61
- model, 61
- module, 53
- \ne_double, 23
- \ne_float, 23
- \nothing, 28
- \null, 54
- \numof, 49
- \offset, 54
- Old, 40, 56
- \old, 28, 29, 40, 56
- polymorphism, 48
- Post, 40, 56
- post-state, 75
- Pre, 40, 56
- pre-state, 75
- predicate, 16
- predicate, 45, 47, 52
- \product, 49
- pure expression, 75
- reads, 52, 65
- real, 18, 20
- real_of_double, 23
- real_of_float, 23
- record, 50
- recursion, 48
- requires, 28, 29
- \result, 28, 29, 56
- returns, 56
- \round_double, 23
- \round_float, 23
- \separated, 54
- set type, 55
- sizeof, 16, 22
- specification, 45
- \specified, 67
- statement contract, 35, 41
- strong, 59
- \subset, 33
- sum, 50
- \sum, 49
- term, 16
- terminates, 28, 44
- termination, 37, 42
- \true, 16, 17, 20
- type
 - concrete, 50
 - polymorphic, 48
 - record, 50
 - sum, 50
- type, 47, 50, 59
- type invariant, 58
- \union, 33
- \valid, 54
- variant, 37
- variant, 36
- \variant, 42
- volatile, 65
- volatile, 65
- weak, 59
- while, 36
- \with, 16
- writes, 65