

The Devil Language

release 0.4

Laurent Réveillère
F. Mérillon, C. Consel, R. Marlet, and G. Muller

August 24, 2000

Foreword

This document is intended as a reference manual for the Devil language. It lists the language constructs, gives their precise syntax and an informal semantics. It is by no means a tutorial introduction to the language.

Notations

The syntax of the language is given in BNF-like notation. Terminal symbols are set in a typewriter font (**like this**). Non-terminals are set in an italic font (*like that*). The vertical bar | denotes an alternative in a rule. Parentheses (...) denote grouping. Parentheses with a trailing star sign (...)* denote zero, one or several occurrences of the enclosed item. Parentheses with a trailing plus sign (...)+ denote one or several occurrences of the enclosed item. Parentheses with a trailing question mark sign (...)? denote an optional item.

Availability

The complete Devil distribution (including the taz compiler) is freely available on the World Wide Web, <http://www.irisa.fr/compose/devil>

1 Introduction

Devil is a domain-specific language for specifying the functional interface of a device. Concretely, a device can be described by three layers of abstractions: *ports*, *registers*, and *device variables*. The entry point of a Devil specification is the declaration of a device, parameterized by ports or range of ports, which abstract physical addresses. Ports then allow device registers to be declared; these define the granularity of interaction with the device. Finally, device variables are defined from registers, forming the functional interface to the device.

Only device variables are visible from outside a Devil description; ports and registers are hidden since these abstractions are not part of the functional interface of the device. In fact, the Devil compiler (named Taz) generates for each variable two C procedures which permit to write or read the variable by emitting the proper I/O operations.

2 Lexical conventions

Blanks

The following characters are considered as blanks : space, newline and horizontal tabulation. Blanks separate adjacent identifiers, literals and keywords that would be otherwise confused as a single identifier, literal or keyword. Apart from that, they are ignored.

Comments

Comments are C-like comments. They start with the two characters `/*` and end with the characters `*/`. C++ comments style can also be used; all characters from the two characters `//` till the end of the line are considered as comments too. Comments are treated as blanks.

Identifiers

Identifiers are a sequence of letters, digits and `_` (the underscore character) starting with a letter. A letter can be any of the 52 lowercase and uppercase letters from the ASCII set. The current implementation places no limit on the number of characters of an identifier.

$$\begin{aligned} \textit{ident} & ::= \textit{letter} (\textit{letter} \mid \textit{digit} \mid _)^* \\ \textit{letter} & ::= \text{A..Z} \mid \text{a..z} \\ \textit{digit} & ::= \text{0..9} \end{aligned}$$

Integer literals

An integer literal is a sequence of one or more digits, optionally preceded by a minus sign. By default, integers literals are in decimal (radix 10).

$$\begin{aligned} \textit{integer-literal} & ::= (0..9)^+ \\ & \mid 0\text{x} (0..9 \mid \text{A..F} \mid \text{a..f})^+ \\ & \mid 0\text{o} (0..7)^+ \end{aligned}$$

The following prefixes select a different radix:

Prefix	Radix
0x	hexadecimal (radix 16)
0o	octal (radix 8)

Note that the initial 0 is digit zero and the o for octal is letter o.

Boolean literals

The *boolean* type has two possible values, represented by the literals `true` and `false`. A boolean literal is always of type `boolean`.

$$\begin{aligned} \textit{boolean-literal} & ::= \text{true} \\ & \mid \text{false} \end{aligned}$$

Bit literals

Bit literals are delimited by ' (single quote) characters.

```
bits-literal ::= ' (bit)+ '  
bit          ::= 0 | 1 | * | .
```

Prefix and infix operators

The following tokens are the Devil operators :

```
@    ..  .  =>  <=  <=>  
=    #  
+    -    *    /  
==   !=   <   >   >=  
&&  ||   !
```

Note that sequences of “operator characters”, such as != or <=> are read as a single token.

Keywords

The identifiers below are reserved keywords:

```
as      bit      bool      const    device  else  
except  false     for       if       int     mask  
port    post      pre      private  read    register  
serialized  set      structure trigger  true    type  
unused  variable  volatile write
```

The following character sequences are also reserved:

```
{ } [ ] ( )  
: , = ;
```

Ambiguities

Lexical ambiguities are resolved according to the “longest match” rule: when a character sequence can be decomposed into tokens in several different ways, the resulting decomposition is the one with the longest first token.

3 Device Specification

```
device-spec ::= device ident (device-params) { (object-def;) + }  
device-params ::= device-param (, device-param)*  
device-param ::= ident : type-expr port (@ { integer-ranges } )?  
integer-ranges ::= integer-range (, integer-range)*  
integer-range ::= integer-literal  
| integer-literal .. integer-literal  
object-def ::= objdef  
| if (expr) objdef-body (else objdef-body)?  
objdef-body ::= objdef  
| { (objdef;) + }  
objdef ::= register-definition  
| structure-definition  
| type-definition  
| variable-definition  
| unused-declaration
```

A device specification is introduced by the **device** keyword, and consists of an identifier, followed by a comma separated list of parameters, and then one or more Devil object definitions . Devil objects are register, structure, type or variable. The identifier is the name of the device being defined.

Device parameters

A device parameter named *base* specifies the number of bits that can be exchanged at the *base* address. The optional construct @ { *integer-ranges* } allows one to define a set of integer offsets that can be added to the *base* address in order to compute an other address with the same characteristics. Note that the device parameter definition

```
base : bit[8] port
```

is equivalent to

```
base : bit[8] port @ {0}
```

Conditional definitions

The expression **if** (*expr*) *def-body*₁ **else** *def-body*₂ allows one to conditionally define a Devil object. A Devil object cannot be re-defined; the same identifier cannot be used for several different Devil objects, even if its definition is conditional. A definition in *def-body*₁ (respectively *def-body*₂) can be accessed at a time *t* if and only if the expression *expr* can be evaluated to **true** (respectively **false**) in the environment present at time *t*. The **else** *def-body*₂ part can be omitted, in which case it defaults to an empty definition.

Expressions of conditional definitions

The expression of a conditional definition is an expression where identifiers can refer to device parameters or variables. The expression is evaluated each time an access to a variable conditionally defined is performed. The identifier name of a read variable is evaluated to the previous value read. The identifier name of a write or read-write variable is evaluated to the

previous value assigned to this variable. Moreover, code can be inserted in the generated interface in order to check if a variable has been previously read or written.

4 Expressions

$$\begin{array}{lcl}
 \textit{expr} & ::= & \textit{boolean-literal} \\
 & & | \textit{integer-literal} \\
 & & | \textit{ident} \\
 & & | (\textit{expr}) \\
 & & | \textit{uop expr} \\
 & & | \textit{expr binary-operator expr} \\
 \textit{unary-operator} & ::= & - \mid ! \\
 \textit{binary-operator} & ::= & + \mid - \mid * \mid / \mid \% \\
 & & | \&\& \mid || \\
 & & | == \mid != \mid < \mid > \mid <= \mid >=
 \end{array}$$

The table below shows the relative precedence and associativity of operators. The operators are ordered top-down from the highest to the lowest precedence.

Operators	Associativity
!	right to left
* / %	left to right
+ -	left to right
< <= > >=	left to right
== !=	left to right
&&	left to right
	left to right

Note that the unary operator - has higher precedence than binary operators.

5 Registers

This section describes the register declaration aspects of a Devil specification.

```

register-definition ::= register ident ((register-params)? = register-def
register-params ::= register-param (, register-param)*
register-param ::= ident : type-expr
register-def ::= (register-ports | reg-params) (, register-attribute)* : type-expr
register-ports ::= register-port (, register-port)?
reg-params ::= reg-param (, reg-param)?
register-port ::= (rw)? regport-def
exprs ::= expr (, expr)*
regport-def ::= ident
                | ident @ expr
reg-param ::= (rw)? ident (exprs)
register-attribute ::= (rw)? pre { (action-definition ;)+ }
                | (rw)? post { (action-definition ;)+ }
                | (rw)? mask bits-literal
rw ::= read
        | write
action-definition ::= ident = expr
                    | ident.ident = expr
                    | ident[expr] = expr
                    | ident = { (action-definition ;)+ }

```

A register identifies data that are present in the device as well as how these data can be accessed (read or written). If a register definition only has a read (resp. write) part, the register is said to be a read-only (resp. a write-only) register. Otherwise it is considered as a read-write register. The type of a register consists of a number of bits and is the same for both the read and write parts. In addition, the register definition may contain the description of other four components: port, bit-mask, pre-actions and post-actions.

Ports

A port is defined by the expression *ident* @ *expr* which can be optionally preceded by a read or write modifier. *ident* is the identifier name of a device parameter. *expr* is a static expression that must be evaluable at compile time. The @ *expr* can be omitted, in which case it defaults to @ 0. One or two ports must be provided in a register definition. The read (resp. write) part of a register is defined if and only if a read (resp. write) port definition is provided. If only one port is provided without a modifier, its definition is used both for the read and write parts.

```

device logitech_busmouse (base : bit[8] port @ {0..3})
{ ... }

```

Bit-masks

A bit-mask is introduced by the **mask** keyword, optionally preceded by a read or write modifier and followed by a *bits-literal*. If a bit-mask is preceded by a read modifier (respectively write modifier), its definition applies only to the read part (respectively write part) of the register definition. The *bits-literal* of a bit-mask is a sequence of *n* mask elements (0,1, * or .) where *n* stands for the size (number of bits) of the register. Mask elements are numbered

from right to left so that the right hand side element of a bit-mask masks the lowest bit of the corresponding register. The table below shows, for each mask element, the constraint associated to its corresponding bit.

mask	bit mapped	bit value
0	no	0
1	no	1
*	no	any
.	yes	bit variable value

If no bit-mask is provided, it defaults to ' $\underbrace{\dots}_n$ ' where n stands for the size of the register.

Unused declarations

The bit of a register masked with . (dot sign) has to be mapped by a device variable. The unused declaration can be used when the variable is conditionally defined to declare bits that are conditionnaly mapped.

unused-declaration ::= **unused bit** *register-bits*

The unused declaration is introduced by the **unused bit** keywords, optionally preceded by a read or write modifier, and followed by a *register-bits* expression (see section 7). An unused declaration has to be conditionally defined since a bit which is always irrelevant must be masked with the * sign.

Pre and post actions

Pre-actions specify a list of actions that have to be performed before reading or writing the register. Post-actions specify a list of actions that have to be performed after reading or writing the register. Pre-actions and post-actions are introduced by the **pre** and **post** keywords, optionally preceded by a read or write modifier and followed by a list of actions. An action is the assignment of an identifier to a specific value. The identifier designates a variable, the field of a structure or the element of an array. If no pre-actions or post-actions are provided, they default to **pre** {} or **post** {} where no actions have to be performed. If pre-actions or post-actions are preceded by a modifier, their definitions apply only to the read or write part of the register definition.

5.1 Register Extensions

This section describes some language extensions that introduce some short-hand notations for register definitions.

Parameterized registers

A parameterized register is introduced by appending a list of parameters to the identifier name of the register being defined. The right-hand side of the definition is identical to a standard register definition except that the parameters can appear in the expressions. Parameterized registers can be seen as functions from a set of parameters to a register definition.

```
register reg(i : int(2)) = read base @ 0, pre {index = i} : bit[8];
```

Parameterized register applications

The parameterized register application of a register definition is replaced by the definition of the parameterized register where parameters have been replaced by specific values. If a component is specified in the parameterized register, it can not be redefined. That is, a component can not be specified both in the parameterized register definition and in the *register-attribute* part of the register definition. If the parameterized register application is preceded by a read or write modifier, the modifier is applied to all components of the parameterized register definition. Components that have the read and write modifiers are ignored.

```
register x_low = reg(0), mask '****....';
```

6 Type Expressions

<i>type-expr</i>	::=	<i>typeexpr-def</i> <i>typeexpr-spec</i>
<i>typeexpr-def</i>	::=	{ <i>map</i> (, <i>map</i>)* } ([<i>integer-literal</i>])?
<i>typeexpr-spec</i>	::=	(signed)? int { <i>integer-ranges</i> } (signed)? int (<i>integer-literal</i>) bool <i>ident</i> <i>typeexpr-spec</i> [<i>integer-literal</i>]
<i>map</i>	::=	(private)? <i>ident</i> <i>arrow</i> <i>bit-pattern</i> (private)? <i>ident</i> (<i>type-params</i>) <i>arrow</i> <i>algtype-def</i>
<i>type-params</i>	::=	<i>type-param</i> (, <i>type-param</i>)*
<i>type-param</i>	::=	<i>ident</i> : <i>typeexpr-spec</i>
<i>algtype-def</i>	::=	<i>bit-pattern</i> <i>bit-pattern</i> # <i>algtype-def</i>
<i>arrow</i>	::=	<= => <=>
<i>bit-pattern</i>	::=	<i>bits-literal</i> <i>integer-literal</i>
<i>type-definition</i>	::=	(private)? type <i>ident</i> = <i>type-expr</i>

A *map-type* definition consists of a list of *map-type equations* and is declared by using the {*map* (, *map*)*} construct. It specifies an enumeration of identifiers that form the abstract values of the type being defined. A *map-type* can be thought of as enumerated type.

Map-type equation

A *map-type equation* is introduced by the identifier name of an abstract value, followed by whether a left (<=) or a right (=>) arrow and then a bits-literal or an integer-literal. The identifier on the left-hand side is the name of a value (*abstract value*) of the type being defined. The value on the right-hand side denotes the corresponding *concrete value* that is read or written to the corresponding variable. If the identifier name on the left-hand side is preceded by the **private** keyword, then the defined value is declared to be private to the Devil specification and will not be exported into the generated interface.

```

{
  ROTATE_IN_AEOI_CLEAR    => '000',
  NON_SPECIFIC            => '001',
  private NOP             => '010',
  SPECIFIC_EOI            => '011',
  ROTATE_IN_AEOI_SET      => '100',
  ROTATE_NON_SPECIFIC_EOI => '101',
  SET_PRIORITY            => '110',
  ROTATE_SPECIFIC_EOI     => '111'
};

```

One of the private value of a *map-type* must be used at least once. Abstract values denote values of the types used in the generated interface. Concrete values are the corresponding bit-strings that are read or written in the register of a device. A *map-type equation* specifies both abstract and concrete values and defines if the abstract value can be read or written. The left arrow (\leftarrow) is used to specify that the concrete value can be read, and that its interpretation must be the abstract value defined on the left-hand side. The right arrow (\rightarrow) specifies that the concrete value can be written. The left-right arrow (\leftrightarrow) can be used as a short hand notation when a concrete value, associated to the same abstract value, can be read or written. As an example, the following *map-type equations*

```

X  $\leftarrow$  '10',
X  $\rightarrow$  '10'

```

are equivalent to

```

X  $\leftrightarrow$  '10'

```

Bit Patterns

If different concrete values are associated to the same abstract value in read map-type equations, then one can use a *bit-pattern* as a short hand notation. The * (star sign) in a bits-literal denotes any bit. As an example, the following map-type equations

```

X  $\leftarrow$  '10',
X  $\leftarrow$  '11'

```

are equivalent to

```

X  $\leftarrow$  '1*'

```

Note that patterns are allowed only for read map-type equations.

Type abbreviations

The *type-definition* rule defines the type identifier name as an abbreviation for the type expression on the right-hand side of = (equal sign). If the **type** keyword is preceded by the **private** keyword, the type is not exported to the generated interface. A type declared as private must be used at least once in the Devil specification where it has been defined.

Algebraic types

An algebraic type equation denotes an abstract value which has a concrete representation that depends on values of other types. As an example, assuming that type t_1 is defined as follows :

```

type t1 = {
  X => '10',
  Y => '11',
  Z => '01'
};

```

one can define type t_2 by

```

type t2 = {
  A (x : t1) => '10' # x # '00',
  B          => '110000'
};

```

Any Devil type can be used as an algebraic type parameter. The # (sharp sign) denotes a concatenation of bits that forms the concrete value. Note that there is no cost penalty in using algebraic types for writing values. The cost for reading depends on the type representation complexity.

7 Variables

<i>variable-definition</i>	::=	(private)? <i>variable ident</i> = <i>vardef</i>
<i>vardef</i>	::=	<i>var-registers</i> (, <i>var-attribute</i>)* : <i>type-expr</i>
<i>var-registers</i>	::=	<i>var-regs</i> (, <i>var-regs</i>)?
<i>var-regs</i>	::=	(<i>rw</i>)? <i>register-bits</i>
<i>register-bits</i>	::=	<i>ident</i>
		<i>ident</i> [<i>integer-ranges</i>]
		<i>ident</i> (<i>exprs</i>)
		<i>ident</i> (<i>exprs</i>)[<i>integer-ranges</i>]
		<i>register-bits</i> # <i>register-bits</i>
<i>var-attribute</i>	::=	volatile
		(<i>rw</i>)? trigger
		(<i>rw</i>)? trigger for <i>trigexpr-paren</i>
		(<i>rw</i>)? trigger except <i>trigexpr-paren</i>
		(<i>rw</i>)? serial-spec
<i>trigexpr-paren</i>	::=	<i>trigexpr</i>
		(<i>trigexpr</i> (, <i>trigexpr</i>)*)
<i>trigexpr</i>	::=	<i>integer-range</i>
		<i>ident</i>
		<i>boolean-literal</i>

A variable object provides an abstract interface to concrete register values and specifies the semantics of the data stored in registers. A variable definition has two optional parts: the read part and the write part. If a variable definition does not have a read (write) part, the register is said to be write-only (read-only), otherwise it is referred to be read-write. The read part of a variable definition contains the specification of four components: register-bits, serialization, volatile and trigger. The write part specifies three components: register-bits, serialization and trigger. The type specification is shared for both the read and write parts.

Visibility

A variable object is introduced by the `variable` keyword, followed by the identifier name of the variable being defined, and optionally preceded by the `private` keyword. When a variable is declared as private, its definition is not exported in the generated interface and must be used at least once within the devil program where its definition occurs. As an example, private variables are often declared to model indexes and used in pre-actions of indexed-registers.

```
private variable index = index_reg[6..5] : int(2);
```

Register-bits

The *register-bits* component specifies which bit of which register must be concatenated in order to obtain the bit-string that forms a device variable. The register-bits component is recursively defined as follows: (1) a register identifier name denotes all its bits; (2) a register identifier name followed by an integer range listed between square brackets denotes all bits which have a position number that appear in the list; (3) the expression $rb_1 \# rb_2$ denotes the concatenation (from left to right) of the bits specified by rb_1 and rb_2 . Note that the bit number 0 of register *reg* is the less significant bit. One or two register-bits must be provided in a variable definition. The read or write part of a variable is defined if and only if a read or write register-bits definition is provided. If only one register-bits is provided, without a modifier, its definition is used for both the read and write parts.

Volatile

The `volatile` attribute is only attached to the read part of a variable definition. When a variable is declared as volatile, each read operation may produce a different value. If the volatile attribute is not provided, it defaults to non volatile.

Trigger

The trigger attribute is introduced by the `trigger` keyword, optionally preceded by a read or write modifier. When a variable is declared as trigger, each access to this variable triggers an action or a set of actions inside the device. For example, writing twice the same value does not produce the same result that one write operation. As a consequence, two trigger variables that map bits of the same register have to be grouped in a structure. When the trigger attribute is used for the write part of a variable definition, some specific values can sometimes cancel the trigger effect. These values can be specified by using the attributes `trigger for` and `trigger except`. If no trigger attribute is provided, it defaults to non trigger. If the trigger attribute is preceded by the read or write modifier, its definition applies only to the read or write part of the variable definition.

Trigger for ...

When the `trigger` keyword is followed by the `for` keyword and a list of expressions, the variable is declared as trigger only for values that match the specified list of expressions.

Trigger except ...

When the `trigger` keyword is followed by the `except` keyword and then by a list of expressions, the variable is declared as trigger only for values that do not match the specified list of expressions. Since all Devil types are finite, the set of values that cancel the trigger effect of a variable is also finite.

7.1 Structures

```
structure-definition ::= structure ident = structdef
structdef           ::= { (variable-definition;) + } (structure-attributes)?
structure-attributes ::= structure-attribute (, structure-attribute)?
structure-attribute ::= (rw)? serial-spec
```

A structure is a collection of variable definitions.

```
structure it_status = {
  variable it_status_rst = read isr[7], volatile : status;
  variable it_status_rdc = read isr[6], volatile : status;
  variable it_status_cnt = read isr[5], volatile : status;
  variable it_status_ovw = read isr[4], volatile : status;
  variable it_status_txe = read isr[3], volatile : status;
  variable it_status_rxe = read isr[2], volatile : status;
  variable it_status_ptx = read isr[1], volatile : status;
  variable it_status_prx = read isr[0], volatile : status;
};
```

When more than one register is used for defining variables contained in the structure, the `serialized as` expression has to be provided. When a structure is read or written, the serial construction specifies in which order read or write operations are executed. When bits of a register are mapped to variables defined in the same structure, only one read or write operation is executed for this register.

```
structure init = {
  variable sngl = icw1[1] : { SINGLE => '1', CASCADED => '0'};
  variable ic4 = icw1[0] : bool;
  ...
  variable microprocessor = icw4[0] : { X8086 => '1', MCS80_85 => '0'};
} serialized as {
  icw1; icw2;
  if (sngl == SINGLE) icw3;
  if (ic4 == true) icw4;
};
```

7.2 Serialization

```
serial-spec ::= serialized as serial-desc
serial-desc ::= ident
                | { serial-def (; serial-def)* }
serial-def  ::= ident
                | if (expr) serial-desc (else serial-desc)?
```

The `serialized as` construct is used to specify in which order registers of a variable or a structure have to be accessed. The `serialized` keyword is optionally preceded by a read

or write modifier, in which case its definition is restricted to the read or write part of the variable or structure being defined. If no modifier is specified, the same serial definition is used for both the read and write part. If no serial definition is provided, one is deduced from the variable or structure definition. Identifiers used in the rules *serial-desc* ::= *ident* and *serial-def* ::= *ident* must refer to register names.

Sequence

The serial description $r_1; r_2$ specifies that register r_1 must be accessed first, followed by an access to register r_2 .

Conditional

The serial description `if (expr) s1 else s2` specifies conditional definitions (depending on runtime values of given variables) of register sequences.

Example: The 8259A interrupt controller possesses various execution modes that depend on the hardware configuration (processor type, cascaded/single controller). The initialization of the controller is performed by writing to configuration variables defined over four initialization registers. In fact, the initialization sequence varies with the actual values of configuration variables. Additionally, three of the configuration registers (e.g., `icw2`, `icw3`, `icw4`) are mapped on a single port and their addressing is implicitly done by previously written configuration values. The following example shows how such an addressing mode can be specified in Devil: configuration variables are grouped together within the `init` structure whose register write operations are ordered using tests on variable values.

```

register icw1 = write base@0, mask '***1***' : bit[8];
register icw2 = write base@1 : bit[8];
register icw3 = write base@1 : bit[8];
register icw4 = write base@1, mask '000***' : bit[8];

structure init = {
  variable sngl = icw1[1] : { SINGLE => '1', CASCADED => '0'};
  variable ic4 = icw1[0] : bool;
  ...
  variable microprocessor = icw4[0] : { X8086 => '1', MCS80_85 => '0'};
} serialized as {
  icw1; icw2;
  if (sngl == SINGLE) icw3;
  if (ic4 == true) icw4;
};

```

7.3 Variable Extensions

When all bits of a register are mapped to a single variable and when this register can be expressed with a parameterized register, one can write:

```
variable foo = reg(1) : int(8);
```

instead of

```
register r1 = reg(1);
variable foo = r1 : int(8);
```

This two constructions are strictly equivalent. A similar case is when all bits (of a register) that can be mapped are mapped to a single variable. As illustrated previously, it is possible to insert a parameterized register application in a variable definition, but it is also possible to extract bits of this defined register with the [...] construct. The Devil code below illustrates this construction:

```
variable foo = reg(1)[4..0] : int(4);
```