The SmPL Grammar (version 0.2.3)

Research group on Coccinelle

August 20, 2010

This document presents the grammar of the SmPL language used by the Coccinelle tool. For the most part, the grammar is written using standard notation. In some rules, however, the left-hand side is in all uppercase letters. These are macros, which take one or more grammar rule right-hand-sides as arguments. The grammar also uses some unspecified nonterminals, such as id, const, etc. These refer to the sets suggested by the name, *i.e.*, id refers to the set of possible C-language identifiers, while const refers to the set of possible C-language constants. A HTML version of this documentation is available online at `http://coccinelle.lip6.fr/docs/main_grammar.html`.

1 Program

$$
\begin{array}{rcl}
program & ::= & include\_cocci^{*}\ changeset^{+} \\
include\_cocci & ::= & \texttt{using}\ string \\
 & | & \texttt{using}\ pathToIsoFile \\
 & | & \texttt{virtual}\ \mathsf{id}\ (\texttt{,}\ \mathsf{id})^{*} \\
changeset & ::= & metavariables\ transformation \\
 & | & script\_metavariables\ \mathsf{script\_code}
\end{array}
$$

script_code is any code in the chosen scripting language. Parsing of the semantic patch does not check the validity of this code; any errors are first detected when the code is executed. Furthermore, @ should not be use in this code. Spatch scans the script code for the next @ and considers that to be the beginning of the next rule, even if @ occurs within e.g., a string or a comment.

virtual keyword is used to declare virtual rules. Virtual rules may be subsequently used as a dependency for the rules in the SmPL file. Whether a virtual rule is defined or not is controlled by the -D option on the command line.

2 Metavariables for transformations

The *rulename* portion of the metavariable declaration can specify properties of a rule such as its name, the names of the rules that it depends on, the isomorphisms to be used in processing the rule, and whether quantification over paths should be universal or existential. The optional annotation expression indicates that the pattern is to be considered as matching an expression, and thus can be used to avoid some parsing problems.

The *metadecl* portion of the metavariable declaration defines various types of metavariables that will be used for matching in the transformation section.

1

metavariables	::=	`@@` *metadecl** `@@`
	|	`@` *rulename* `@` *metadecl** `@@`
rulename	::=	id [`extends` id] [`depends on` *dep*] [*iso*] [*disable-iso*] [*exists*] [`expression`]
dep	::=	*pnrule*
	|	*dep* `&&` *dep*
	|	*dep* `||` *dep*
pnrule	::=	id
	|	`!` id
	|	`ever` id
	|	`never` id
	|	(*dep*)
iso	::=	`using` *string* (`,` *string*)*
disable-iso	::=	`disable` *COMMA_LIST*(id)
exists	::=	`exists`
	|	`forall`
COMMA_LIST(*elem*)	::=	*elem* (`,` *elem*)*

The keyword `disable` is normally used with the names of isomorphisms defined in standard.iso or whatever isomorphism file has been included. There are, however, some other isomorphisms that are built into the implementation of Coccinelle and that can be disabled as well. Their names are given below. In each case, the text describes the standard behavior. Using *disable-iso* with the given name disables this behavior.

- `optional_storage`: A SmPL function definition that does not specify any visibility (i.e., static or extern), or a SmPL variable declaration that does not specify any storage (i.e., auto, static, register, or extern), matches a function declaration or variable declaration with any visibility or storage, respectively.

- `optional_qualifier`: This is similar to `optional_storage`, except that here is it the qualifier (i.e., const or volatile) that does not have to be specified in the SmPL code, but may be present in the C code.

- `value_format`: Integers in various formats, e.g., 1 and 0x1, are considered to be equivalent in the matching process.

- `comm_assoc`: An expression of the form *exp bin_op* `...`, where *bin_op* is commutative and associative, is considered to match any top-level sequence of *bin_op* operators containing *exp* as the top-level argument.

The possible types of metavariable declarations are defined by the grammar rule below. Metavariables should occur at least once in the transformation immediately following their declaration. Fresh metavariables must only be used in + code. These properties are not expressed in the grammar, but are checked by a subsequent analysis. The metavariables are designated according to the kind of terms they can match, such as a statement, an identifier, or an expression. An expression metavariable can be further constrained by its type. A declaration metavariable matches the declaration of one or more variables, all sharing the same type specification (*e.g.*, int a,b,c=3;). A field metavariable does the same, but for structure fields.

$$
\begin{array}{rcl}
metadecl & ::= & \texttt{fresh identifier } \textit{ids} \texttt{ ;} \\
& \mid & \texttt{identifier } \textit{COMMA\_LIST}(\textit{pmid\_with\_regexp}) \texttt{ ;} \\
& \mid & \texttt{identifier } \textit{COMMA\_LIST}(\textit{pmid\_with\_virt\_or\_not\_eq}) \texttt{ ;} \\
& \mid & \texttt{parameter } [\texttt{list}] \textit{ ids} \texttt{ ;} \\
& \mid & \texttt{parameter list [} \textit{id} \texttt{] } \textit{ids} \texttt{ ;} \\
& \mid & \texttt{parameter list [} \textit{const} \texttt{] } \textit{ids} \texttt{ ;} \\
& \mid & \texttt{type } \textit{ids} \texttt{ ;} \\
& \mid & \texttt{statement } [\texttt{list}] \textit{ ids} \texttt{ ;} \\
& \mid & \texttt{declaration } [\texttt{list}] \textit{ ids} \texttt{ ;} \\
& \mid & \texttt{field } [\texttt{list}] \textit{ ids} \texttt{ ;} \\
& \mid & \texttt{typedef } \textit{ids} \texttt{ ;} \\
& \mid & \texttt{declarer name } \textit{ids} \texttt{ ;} \\
& \mid & \texttt{declarer } \textit{COMMA\_LIST}(\textit{pmid\_with\_regexp}) \texttt{ ;} \\
& \mid & \texttt{declarer } \textit{COMMA\_LIST}(\textit{pmid\_with\_not\_eq}) \texttt{ ;} \\
& \mid & \texttt{iterator name } \textit{ids} \texttt{ ;} \\
& \mid & \texttt{iterator } \textit{COMMA\_LIST}(\textit{pmid\_with\_regexp}) \texttt{ ;} \\
& \mid & \texttt{iterator } \textit{COMMA\_LIST}(\textit{pmid\_with\_not\_eq}) \texttt{ ;} \\
& \mid & [\texttt{local}] \texttt{ idexpression } [\textit{ctype}] \textit{ COMMA\_LIST}(\textit{pmid\_with\_not\_eq}) \texttt{ ;} \\
& \mid & [\texttt{local}] \texttt{ idexpression } [\{\textit{ctypes}\} \star^*] \textit{ COMMA\_LIST}(\textit{pmid\_with\_not\_eq}) \texttt{ ;} \\
& \mid & [\texttt{local}] \texttt{ idexpression } \star^+ \textit{ COMMA\_LIST}(\textit{pmid\_with\_not\_eq}) \texttt{ ;} \\
& \mid & \texttt{expression list } \textit{ids} \texttt{ ;} \\
& \mid & \texttt{expression } \star^+ \textit{ COMMA\_LIST}(\textit{pmid\_with\_not\_eq}) \texttt{ ;} \\
& \mid & \texttt{expression } \textit{COMMA\_LIST}(\textit{pmid\_with\_not\_ceq}) \texttt{ ;} \\
& \mid & \texttt{expression list [} \textit{id} \texttt{] } \textit{ids} \texttt{ ;} \\
& \mid & \texttt{expression list [} \textit{const} \texttt{] } \textit{ids} \texttt{ ;} \\
& \mid & \textit{ctype} \texttt{ [ ] } \textit{COMMA\_LIST}(\textit{pmid\_with\_not\_eq}) \texttt{ ;} \\
& \mid & \textit{ctype } \textit{COMMA\_LIST}(\textit{pmid\_with\_not\_ceq}) \texttt{ ;} \\
& \mid & \{\textit{ctypes}\} \star^* \textit{ COMMA\_LIST}(\textit{pmid\_with\_not\_ceq}) \texttt{ ;} \\
& \mid & \{\textit{ctypes}\} \star^* \texttt{ [ ] } \textit{COMMA\_LIST}(\textit{pmid\_with\_not\_eq}) \texttt{ ;} \\
& \mid & \texttt{constant } [\textit{ctype}] \textit{ COMMA\_LIST}(\textit{pmid\_with\_not\_eq}) \texttt{ ;} \\
& \mid & \texttt{constant } [\{\textit{ctypes}\} \star^*] \textit{ COMMA\_LIST}(\textit{pmid\_with\_not\_eq}) \texttt{ ;} \\
& \mid & \texttt{position } [\texttt{any}] \textit{ COMMA\_LIST}(\textit{pmid\_with\_not\_eq\_mid}) \texttt{ ;}
\end{array}
$$

It is possible to specify that an expression list or a parameter list metavariable should match a specific number of expressions or parameters.

$$
\begin{array}{rcl}
\textit{ids} & ::= & \textit{COMMA\_LIST}(\textit{pmid}) \\
\textit{pmid} & ::= & \texttt{id} \\
& \mid & \textit{mid} \\
\textit{mid} & ::= & \texttt{rulename\_id.id} \\
\textit{pmid\_with\_regexp} & ::= & \textit{pmid } \tilde{}= \textit{ regexp} \\
\textit{pmid\_with\_not\_eq} & ::= & \textit{pmid } [\mathrel{!=} \textit{id\_or\_meta}] \\
& \mid & \textit{pmid } [\mathrel{!=} \{ \textit{ COMMA\_LIST}(\textit{id\_or\_meta}) \}] \\
\textit{pmid\_with\_virt\_or\_not\_eq} & ::= & \texttt{virtual.id} \\
& \mid & \textit{pmid\_with\_not\_eq} \\
\textit{pmid\_with\_not\_ceq} & ::= & \textit{pmid } [\mathrel{!=} \textit{id\_or\_cst}] \\
& \mid & \textit{pmid } [\mathrel{!=} \{ \textit{ COMMA\_LIST}(\textit{id\_or\_cst}) \}] \\
\textit{id\_or\_cst} & ::= & \texttt{id} \\
& \mid & \texttt{integer} \\
\textit{id\_or\_meta} & ::= & \texttt{id} \\
& \mid & \texttt{rulename\_id.id} \\
\textit{pmid\_with\_not\_eq\_mid} & ::= & \textit{pmid } [\mathrel{!=} \textit{mid}] \\
& \mid & \textit{pmid } [\mathrel{!=} \{ \textit{ COMMA\_LIST}(\textit{mid}) \}]
\end{array}
$$

Subsequently, we refer to arbitrary metavariables as metaid^{ty}, where *ty* indicates the *metakind* used in the declaration of the variable. For example, metaid^{Type} refers to a metavariable that was declared using `type` and stands for any type.

The *ctype* and *ctypes* nonterminals are used by both the grammar of metavariable declarations and the grammar of transformations, and are defined on page 7.

An identifier metavariable with `virtual` as its "rule name" is given a value on the command line. For example, if a semantic patch contains a rule that declares an identifier metavariable with the name `virtual.alloc`, then the command line could contain `-D alloc=kmalloc`. There should not be space around the `=`. An example is in `demos/vm.cocci` and `demos/vm.c`.

3 Metavariables for scripts

Metavariables for scripts can only be inherited from transformation rules. In the spirit of scripting languages such as Python that use dynamic typing, metavariables for scripts do not include type declarations.

$$
\begin{array}{lll}
\textit{script\_metavariables} & ::= & \texttt{@ script:}\textit{language}~[\textit{rulename}]~[\texttt{depends on}~\textit{dep}]~\texttt{@}~\textit{script\_metadecl}^*~\texttt{@@} \\
& | & \texttt{@ initialize:}\textit{language}~[\texttt{depends on}~\textit{dep}]~\texttt{@} \\
& | & \texttt{@ finalize:}\textit{language}~[\texttt{depends on}~\textit{dep}]~\texttt{@} \\
\textit{language} & ::= & \texttt{python} \\
& | & \texttt{ocaml} \\
\textit{script\_metadecl} & ::= & \textsf{id}~\texttt{<<}~\textsf{rulename\_id.id}~\texttt{;} \\
& | & \textsf{id}~\texttt{;}
\end{array}
$$

Currently, the only scripting languages that are supported are Python and OCaml, indicated using `python` and `ocaml`, respectively. The set of available scripting languages may be extended at some point.

Script rules declared with `initialize` are run before the treatment of any file. Script rules declared with `finalize` are run when the treatment of all of the files has completed. There can be at most one of each per scripting language (thus currently at most one of each). Initialize and finalize script rules do not have access to SmPL metavariables. Nevertheless, a finalize script rule can access any variables initialized by the other script rules, allowing information to be transmitted from the matching process to the finalize rule.

A script metavariable that does not specify an origin, using «, is newly declared by the script. This metavariable should be assigned to a string and can be inherited by subsequent rules as an identifier. In Python, the assignment of such a metavariable x should refer to the metavariable as `coccinelle.`x. Examples are in the files `demos/pythontococci.cocci` and `demos/camltococci.cocci`.

In an ocaml script, the following extended form of *script_metadecl* may be used:

$$
\begin{array}{lll}
\textit{script\_metadecl} & ::= & \texttt{(}\textsf{id}\texttt{,}\textsf{id}\texttt{)}~\texttt{<<}~\textsf{rulename\_id.id}~\texttt{;} \\
& | & \textsf{id}~\texttt{<<}~\textsf{rulename\_id.id}~\texttt{;} \\
& | & \textsf{id}~\texttt{;}
\end{array}
$$

In a declaration of the form `(`id`,`id`)` `<<` rulename_id.id `;`, the left component of `(`id`,`id`)` receives a string representation of the value of the inherited metavariable while the right component receives its abstract syntax tree. The file `parsing_c/ast_c.ml` in the Coccinelle implementation gives some information about the structure of the abstract syntax tree. Either the left or right component may be replaced by _, indicating that the string representation or abstract syntax trees representation is not wanted, respectively.

4 Transformation

The transformation specification essentially has the form of C code, except that lines to remove are annotated with −
in the first column, and lines to add are annotated with +. A transformation specification can also use *dots*, "`...`", describing an arbitrary sequence of function arguments or instructions within a control-flow path. Dots may be modified with a `when` clause, indicating a pattern that should not occur anywhere within the matched sequence. Finally, a

transformation can specify a disjunction of patterns, of the form (pat_1 | ... | pat_n) where each (, | or) is in column 0 or preceded by \.

The grammar that we present for the transformation is not actually the grammar of the SmPL code that can be written by the programmer, but is instead the grammar of the slice of this consisting of the − annotated and the unannotated code (the context of the transformed lines), or the + annotated code and the unannotated code. For example, for parsing purposes, the following transformation is split into the two variants shown below and each is parsed separately.

```
1   proc_info_func(...) {
2       <...
3   -     hostno
4   +     hostptr->host_no
5       ...>
6   }
```

```
1   proc_info_func(...) {
2       <...
3   -     hostno
4       ...>
5   }
```

```
1   proc_info_func(...) {
2       <...
3   +     hostptr->host_no
4       ...>
5   }
```

Requiring that both slices parse correctly ensures that the rule matches syntactically valid C code and that it produces syntactically valid C code. The generated parse trees are then merged for use in the subsequent matching and transformation process.

The grammar for the minus or plus slice of a transformation is as follows:

$$
\begin{array}{lll}
\textit{transformation} & ::= & \textit{include}^{+} \\
& | & \textit{OPTDOTSEQ}(\textit{expr},\ \textit{when}) \\
& | & \textit{OPTDOTSEQ}(\textit{decl\_stmt}^{+},\ \textit{when}) \\
& | & \textit{OPTDOTSEQ}(\textit{fundecl},\ \textit{when}) \\
\textit{include} & ::= & \texttt{\#include } \texttt{include\_string} \\
\textit{when} & ::= & \texttt{when != } \textit{when\_code} \\
& | & \texttt{when = } \textit{rule\_elem\_stmt} \\
& | & \texttt{when } \textit{COMMA\_LIST}(\textit{any\_strict}) \\
& | & \texttt{when true != } \textit{expr} \\
& | & \texttt{when false != } \textit{expr} \\
\textit{when\_code} & ::= & \textit{OPTDOTSEQ}(\textit{decl\_stmt}^{+},\ \textit{when}) \\
& | & \textit{OPTDOTSEQ}(\textit{expr},\ \textit{when}) \\
\textit{rule\_elem\_stmt} & ::= & \textit{one\_decl} \\
& | & \textit{expr}; \\
& | & \texttt{return } [\textit{expr}]; \\
& | & \texttt{break}; \\
& | & \texttt{continue}; \\
& | & \texttt{\textbackslash(} \textit{rule\_elem\_stmt} \ (\texttt{\textbackslash|}\ \textit{rule\_elem\_stmt})^{+}\texttt{\textbackslash)} \\
\textit{any\_strict} & ::= & \texttt{any} \\
& | & \texttt{strict} \\
& | & \texttt{forall} \\
& | & \texttt{exists}
\end{array}
$$

$$
\textit{OPTDOTSEQ}(\textit{grammar\_ds}, \textit{when\_ds})\ ::= \\
[\texttt{...}\ \ [\textit{when\_ds}]]\ \textit{grammar\_ds}\ (\texttt{...}\ \ [\textit{when\_ds}]\ \textit{grammar\_ds})^{*}\ [\texttt{...}\ \ [\textit{when\_ds}]]
$$

Lines may be annotated with an element of the set $\{-, +, \star\}$ or the singleton ?, or one of each set. ? represents at most one match of the given pattern. $\star$ is used for semantic match, *i.e.*, a pattern that highlights the fragments annotated with $\star$, but does not perform any modification of the matched code. $\star$ cannot be mixed with $-$ and $+$. There are some constraints on the use of these annotations:

- Dots, *i.e.* `...`, cannot occur on a line marked +.

- Nested dots, *i.e.*, dots enclosed in < and >, cannot occur on a line with any marking.

Each element of a disjunction must be a proper term like an expression, a statement, an identifier or a declaration. Thus, the rule on the left below is not a syntactically correct SmPL rule. One may use the rule on the right instead.

```
1  @@
2  type T;
3  T b;
4  @@
5
6  (
7   writeb(...,
8  |
9   readb(
10 )
11 - (T)
12   b)
```

```
1  @@
2  type T;
3  T b;
4  @@
5
6  (
7  read
8  |
9  write
10 )
11   (...,
12 -   (T)
13     b)
```

5 Types

ctypes	::=	*COMMA_LIST*(*ctype*)
ctype	::=	[*const_vol*] *generic_ctype* $\star^*$
	|	[*const_vol*] `void` $\star^+$
	|	(*ctype* (| *ctype*)*)
const_vol	::=	`const`
	|	`volatile`
generic_ctype	::=	*ctype_qualif*
	|	[*ctype_qualif*] `char`
	|	[*ctype_qualif*] `short`
	|	[*ctype_qualif*] `int`
	|	[*ctype_qualif*] `long`
	|	[*ctype_qualif*] `long long`
	|	`double`
	|	`float`
	|	[`struct`| `union`] id [{ *struct_decl_list** }]
ctype_qualif	::=	`unsigned`
	|	`signed`
struct_decl_list	::=	*struct_decl_list_start*
struct_decl_list_start	::=	*struct_decl*
	|	*struct_decl* *struct_decl_list_start*
	|	`...` [`when` `!=` *struct_decl*]† [*continue_struct_decl_list*]
continue_struct_decl_list	::=	*struct_decl* *struct_decl_list_start*
	|	*struct_decl*
struct_decl	::=	*ctype* *d_ident*`;`
	|	*fn_ctype* (`*` *d_ident*) (*PARAMSEQ*(*name_opt_decl*, ε))`;`)
	|	[*const_vol*] id *d_ident*`;`
d_ident	::=	*id* [[*expr*]]*
fn_ctype	::=	*generic_ctype* $\star^*$
	|	`void` $\star^*$
name_opt_decl	::=	*decl*
	|	*ctype*
	|	*fn_ctype*

† The optional `when` construct ends at the end of the line.

6 Function declarations

$$
\begin{array}{rcl}
\textit{fundecl} & ::= & [\textit{fn\_ctype}] \;\; \textit{funinfo}^* \;\; \textit{funid} \;\; (\,[\textit{PARAMSEQ}(\textit{param}, \;\; \varepsilon)]\,) \;\; \{ \;\; [\textit{stmt\_seq}] \;\; \} \\
\textit{funproto} & ::= & [\textit{fn\_ctype}] \;\; \textit{funinfo}^* \;\; \textit{funid} \;\; (\,[\textit{PARAMSEQ}(\textit{param}, \;\; \varepsilon)]\,)\,; \\
\textit{funinfo} & ::= & \texttt{inline} \\
 & | & \textit{storage} \\
\textit{storage} & ::= & \texttt{static} \\
 & | & \texttt{auto} \\
 & | & \texttt{register} \\
 & | & \texttt{extern} \\
\textit{funid} & ::= & \texttt{id} \\
 & | & \texttt{metaid}^{\text{Id}} \\
\textit{param} & ::= & \textit{type} \;\; \texttt{id} \\
 & | & \texttt{metaid}^{\text{Param}} \\
 & | & \texttt{metaid}^{\text{ParamList}} \\
\textit{decl} & ::= & \textit{ctype} \;\; \textit{id} \\
 & | & \textit{fn\_ctype} \;\; (\texttt{*} \;\; \textit{id}) \;\; (\textit{PARAMSEQ}(\textit{name\_opt\_decl}, \;\; \varepsilon)) \\
 & | & \texttt{void} \\
 & | & \texttt{metaid}^{\text{Param}}
\end{array}
$$

$$
\textit{PARAMSEQ}(\textit{gram\_p}, \textit{when\_p}) \;\; ::= \;\; \textit{COMMA\_LIST}(\textit{gram\_p} \;\; | \;\; \ldots [\textit{when\_p}])
$$

To match a function it is not necessary to provide all of the annotations that appear before the function name. For example, the following semantic patch:

```
1 @@
2 @@
3
4 foo() { ... }
```

matches a function declared as follows:

```
1 static int foo() { return 12; }
```

This behavior can be turned off by disabling the `optional_storage` isomorphism. If one adds code before a function declaration, then the effect depends on the kind of code that is added. If the added code is a function definition or CPP code, then the new code is placed before all information associated with the function definition, including any comments preceeding the function definition. On the other hand, if the new code is associated with the function, such as the addition of the keyword `static`, the new code is placed exactly where it appears with respect to the rest of the function definition in the semantic patch. For example,

```
1 @@
2 @@
3
4 + static
5 foo() { ... }
```

causes static to be placed just before the function name. The following causes it to be placed just before the type

```
1 @@
2 type T;
3 @@
4
5 + static
6 T foo() { ... }
```

It may be necessary to consider several cases to ensure that the added ode is placed in the right position. For example, one may need one pattern that considers that the function is declared `inline` and another that considers that it is not.

7 Declarations

decl_var   ::= *common_decl*
          | [*storage*] *ctype COMMA_LIST*(*d_ident*) ;
          | [*storage*] [*const_vol*] id *COMMA_LIST*(*d_ident*) ;
          | [*storage*] *fn_ctype* (* *d_ident*) (*PARAMSEQ*(*name_opt_decl*, ε)) = *initialize* ;
          | `typedef` *ctype typedef_ident* ;
one_decl   ::= *common_decl*
          | [*storage*] *ctype* id;
          | [*storage*] [*const_vol*] id *d_ident* ;
common_decl ::= *ctype*;
          | *funproto*
          | [*storage*] *ctype d_ident* = *initialize* ;
          | [*storage*] [*const_vol*] id *d_ident* = *initialize* ;
          | [*storage*] *fn_ctype* (* *d_ident*) (*PARAMSEQ*(*name_opt_decl*, ε)) ;
          | *decl_ident* ([*COMMA_LIST*(*expr*)]) ;
initialize   ::= *dot_expr*
          | { [*COMMA_LIST*(*dot_expr*)] }
decl_ident   ::= DeclarerId
          | metaid[Declarer]

8 Statements

The first rule *statement* describes the various forms of a statement. The remaining rules implement the constraints that are sensitive to the context in which the statement occurs: *single_statement* for a context in which only one statement is allowed, and *decl_statement* for a context in which a declaration, statement, or sequence thereof is allowed.

$$
\begin{array}{rcl}
\textit{stmt} & ::= & \textit{include}\\
& | & \text{metaid}^{\text{Stmt}}\\
& | & \textit{expr};\\
& | & \texttt{if (}\textit{dot\_expr}\texttt{)}\ \textit{single\_stmt}\ [\texttt{else}\ \textit{single\_stmt}]\\
& | & \texttt{for (}[\textit{dot\_expr}];\ [\textit{dot\_expr}];\ [\textit{dot\_expr}]\texttt{)}\ \textit{single\_stmt}\\
& | & \texttt{while (}\textit{dot\_expr}\texttt{)}\ \textit{single\_stmt}\\
& | & \texttt{do}\ \textit{single\_stmt}\ \texttt{while (}\textit{dot\_expr}\texttt{)};\\
& | & \textit{iter\_ident}\ (\textit{dot\_expr}^{*})\ \textit{single\_stmt}\\
& | & \texttt{switch (}[\textit{dot\_expr}]\texttt{) \{}\textit{case\_line}^{*}\ \texttt{\}}\\
& | & \texttt{return}\ [\textit{dot\_expr}];\\
& | & \texttt{\{}\ [\textit{stmt\_seq}]\ \texttt{\}}\\
& | & \textit{NEST}(\textit{decl\_stmt}^{+},\ \textit{when})\\
& | & \textit{NEST}(\textit{expr},\ \textit{when})\\
& | & \texttt{break};\\
& | & \texttt{continue};\\
& | & \textit{id}\texttt{:}\\
& | & \texttt{goto}\ \textit{id};\\
& | & \texttt{\{}\textit{stmt\_seq}\ \texttt{\}}\\[4pt]
\textit{single\_stmt} & ::= & \textit{stmt}\\
& | & \textit{OR}(\textit{stmt})\\
\textit{decl\_stmt} & ::= & \text{metaid}^{\text{StmtList}}\\
& | & \textit{decl\_var}\\
& | & \textit{stmt}\\
& | & \textit{OR}(\textit{stmt\_seq})\\
\textit{stmt\_seq} & ::= & \textit{decl\_stmt}^{*}\ [\textit{DOTSEQ}(\textit{decl\_stmt}^{+},\ \textit{when})\ \textit{decl\_stmt}^{*}]\\
& | & \textit{decl\_stmt}^{*}\ [\textit{DOTSEQ}(\textit{expr},\ \textit{when})\ \textit{decl\_stmt}^{*}]\\
\textit{case\_line} & ::= & \texttt{default : }\textit{stmt\_seq}\\
& | & \texttt{case}\ \textit{dot\_expr}\ \texttt{: }\textit{stmt\_seq}\\
\textit{iter\_ident} & ::= & \text{IteratorId}\\
& | & \text{metaid}^{\text{Iterator}}
\end{array}
$$

$$
\begin{array}{rcl}
\textit{OR}(\textit{gram\_o}) & ::= & \texttt{(}\ \textit{gram\_o}\ (\texttt{|}\,\textit{gram\_o})^{*}\ \texttt{)}\\
\textit{DOTSEQ}(\textit{gram\_d, when\_d}) & ::= & \texttt{...}[\textit{when\_d}]\ (\textit{gram\_d}\ \texttt{...}[\textit{when\_d}])^{*}\\
\textit{NEST}(\textit{gram\_n, when\_n}) & ::= & \texttt{<...}[\textit{when\_n}]\ \textit{gram\_n}\ (\texttt{...}[\textit{when\_n}]\ \textit{gram\_n})^{*}\ \texttt{...>}\\
& | & \texttt{<+...}[\textit{when\_n}]\ \textit{gram\_n}\ (\texttt{...}[\textit{when\_n}]\ \textit{gram\_n})^{*}\ \texttt{...+>}
\end{array}
$$

OR is a macro that generates a disjunction of patterns. The three tokens (, |, and) must appear in the leftmost column, to differentiate them from the parentheses and bit-or tokens that can appear within expressions (and cannot appear in the leftmost column). These token may also be preceded by \ when they are used in an other column. These tokens are furthermore different from (, |, and), which are part of the grammar metalanguage.

9 Expressions

A nest or a single ellipsis is allowed in some expression contexts, and causes ambiguity in others. For example, in a sequence ...*expr* ..., the nonterminal *expr* must be instantiated as an explicit C-language expression, while in an array reference, $expr_1$ [$expr_2$], the nonterminal $expr_2$, because it is delimited by brackets, can be also instantiated as ..., representing an arbitrary expression. To distinguish between the various possibilities, we define three nonterminals for expressions: *expr* does not allow either top-level nests or ellipses, *nest_expr* allows a nest but not an ellipsis, and *dot_expr* allows both. The EXPR macro is used to express these variants in a concise way.

$$
\begin{array}{lll}
\textit{expr} & ::= & \textit{EXPR}(\textit{expr}) \\
\textit{nest\_expr} & ::= & \textit{EXPR}(\textit{nest\_expr}) \\
& | & \textit{NEST}(\textit{nest\_expr},\ \textit{exp\_whencode}) \\
\textit{dot\_expr} & ::= & \textit{EXPR}(\textit{dot\_expr}) \\
& | & \textit{NEST}(\textit{dot\_expr},\ \textit{exp\_whencode}) \\
& | & \texttt{...}\ [\textit{exp\_whencode}] \\
\textit{EXPR}(\textit{exp}) & ::= & \textit{exp\ assign\_op\ exp} \\
& | & \textit{exp}\texttt{++} \\
& | & \textit{exp}\texttt{--} \\
& | & \textit{unary\_op\ exp} \\
& | & \textit{exp\ bin\_op\ exp} \\
& | & \textit{exp}\ \texttt{?}\ \textit{dot\_expr}\ \texttt{:}\ \textit{exp} \\
& | & \texttt{(}\textit{type}\texttt{)}\ \textit{exp} \\
& | & \textit{exp}\ \texttt{[}\textit{dot\_expr}\texttt{]} \\
& | & \textit{exp}\ \texttt{.}\ \textit{id} \\
& | & \textit{exp}\ \texttt{->}\ \textit{id} \\
& | & \textit{exp}\,\texttt{(}[\textit{PARAMSEQ}(\textit{arg},\ \textit{exp\_whencode})]\texttt{)} \\
& | & \textit{id} \\
& | & \text{metaid}^{\text{Exp}} \\
& | & \text{metaid}^{\text{Const}} \\
& | & \textit{const} \\
& | & \texttt{(}\textit{dot\_expr}\texttt{)} \\
& | & \textit{OR}(\textit{exp}) \\
\textit{arg} & ::= & \textit{nest\_expr} \\
& | & \text{metaid}^{\text{ExpList}} \\
\textit{exp\_whencode} & ::= & \texttt{when != }\textit{expr} \\
\textit{assign\_op} & ::= & \texttt{=}\ |\ \texttt{-=}\ |\ \texttt{+=}\ |\ \texttt{*=}\ |\ \texttt{/=}\ |\ \texttt{\%=} \\
& | & \texttt{\&=}\ |\ \texttt{|=}\ |\ \texttt{\textasciicircum=}\ |\ \texttt{<<=}\ |\ \texttt{>>=} \\
\textit{bin\_op} & ::= & \texttt{*}\ |\ \texttt{/}\ |\ \texttt{\%}\ |\ \texttt{+}\ |\ \texttt{-} \\
& | & \texttt{<<}|\ \texttt{>>}|\ \texttt{\textasciicircum}\ |\ \texttt{\&}\ |\ \texttt{|} \\
& | & \texttt{<}\ |\ \texttt{>}\ |\ \texttt{<=}\ |\ \texttt{>=}\ |\ \texttt{==}\ |\ \texttt{!=}\ |\ \texttt{\&\&}\ |\ \texttt{||} \\
\textit{unary\_op} & ::= & \texttt{++}\ |\ \texttt{-}\ |\ \texttt{\&}\ |\ \texttt{*}\ |\ \texttt{+}\ |\ \texttt{-}\ |\ \texttt{!} \\
\end{array}
$$

10 Constant, Identifiers and Types for Transformations

$$
\begin{array}{lll}
\textit{const} & ::= & \textit{string} \\
& | & \texttt{[0-9]+} \\
& | & \texttt{...} \\
\textit{string} & ::= & \texttt{"[\textasciicircum"]*"} \\
\textit{id} & ::= & \text{id}\ |\ \text{metaid}^{\text{Id}} \\
\textit{typedef\_ident} & ::= & \text{id}\ |\ \text{metaid}^{\text{Type}} \\
\textit{type} & ::= & \textit{ctype}\ |\ \text{metaid}^{\text{Type}} \\
\textit{pathToIsoFile} & ::= & \texttt{<.*>} \\
\textit{regexp} & ::= & \texttt{"[\textasciicircum"]*"} \\
\end{array}
$$

11 Examples

This section presents a range of examples. Each example is presented along with some C code to which it is applied. The description explains the rules and the matching process.

11.1 Function renaming

One of the primary goals of Coccinelle is to perform software evolution. For instance, Coccinelle could be used to perform function renaming. In the following example, every occurrence of a call to the function foo is replaced by a call to the function bar.

Before	Semantic patch	After

```
1 #DEFINE TEST "foo"
2
3 printf("foo");
4
5 int main(int i) {
6 //Test
7   int k = foo();
8
9   if(1) {
10     foo();
11   } else {
12     foo();
13   }
14
15   foo();
16 }
```

```
1 @@
2
3 @@
4
5
6 - foo()
7 + bar()
```

```
1 #DEFINE TEST "foo"
2
3 printf("foo");
4
5 int main(int i) {
6 //Test
7   int k = bar();
8
9   if(1) {
10     bar();
11   } else {
12     bar();
13   }
14
15   bar();
16 }
```

11.2   Removing a function argument

Another important kind of evolution is the introduction or deletion of a function argument. In the following example, the rule `rule1` looks for definitions of functions having return type `irqreturn_t` and two parameters. A second *anonymous* rule then looks for calls to the previously matched functions that have three arguments. The third argument is then removed to correspond to the new function prototype.

```
1  @ rule1 @
2  identifier fn;
3  identifier irq, dev_id;
4  typedef irqreturn_t;
5  @@
6
7  static irqreturn_t fn (int irq, void *dev_id)
8  {
9     ...
10 }
11
12 @@
13 identifier rule1.fn;
14 expression E1, E2, E3;
15 @@
16
17   fn(E1, E2
18 -   ,E3
19     )
```

drivers/atm/firestream.c at line 1653 before transformation

```
1  static void fs_poll (unsigned long data)
2  {
3          struct fs_dev *dev = (struct fs_dev *) data;
4
5          fs_irq (0, dev, NULL);
6          dev->timer.expires = jiffies + FS_POLL_FREQ;
7          add_timer (&dev->timer);
8  }
```

drivers/atm/firestream.c at line 1653 after transformation

```
1  static void fs_poll (unsigned long data)
2  {
3          struct fs_dev *dev = (struct fs_dev *) data;
4
5          fs_irq (0, dev);
6          dev->timer.expires = jiffies + FS_POLL_FREQ;
7          add_timer (&dev->timer);
8  }
```

11.3 Introduction of a macro

To avoid code duplication or error prone code, the kernel provides macros such as BUG_ON, DIV_ROUND_UP and FIELD_SIZE. In these cases, the semantic patches look for the old code pattern and replace it by the new code.

A semantic patch to introduce uses of the DIV_ROUND_UP macro looks for the corresponding expression, *i.e.*, $(n + d - 1)/d$. When some code is matched, the metavariables n and d are bound to their corresponding expressions. Finally, Coccinelle rewrites the code with the DIV_ROUND_UP macro using the values bound to n and d, as illustrated in the patch that follows.

Semantic patch to introduce uses of the DIV_ROUND_UP macro

```
1  @ haskernel @
2  @@
3
4  #include <linux/kernel.h>
5
6  @ depends on haskernel @
7  expression n,d;
8  @@
9
10 (
11 - (((n) + (d)) - 1) / (d))
12 + DIV_ROUND_UP(n,d)
13 |
14 - (((n) + ((d) - 1)) / (d))
15 + DIV_ROUND_UP(n,d)
16 )
```

Example of a generated patch hunk

```
1  --- a/drivers/atm/horizon.c
2  +++ b/drivers/atm/horizon.c
3  @@ -698,7 +698,7 @@ got_it:
4                  if (bits)
5                          *bits = (div<<CLOCK_SELECT_SHIFT) | (pre-1);
6                  if (actual) {
7  -                        *actual = (br + (pre<<div) - 1) / (pre<<div);
8  +                        *actual = DIV_ROUND_UP(br, pre<<div);
9                          PRINTD (DBG_QOS, "actual_rate:_%u", *actual);
10                 }
11                 return 0;
```

The `BUG_ON` macro makes a assertion about the value of an expression. However, because some parts of the kernel define `BUG_ON` to be the empty statement when debugging is not wanted, care must be taken when the asserted expression may have some side-effects, as is the case of a function call. Thus, we create a rule introducing `BUG_ON` only in the case when the asserted expression does not perform a function call.

On particular piece of code that has the form of a function call is a use of `unlikely`, which informs the compiler that a particular expression is unlikely to be true. In this case, because `unlikely` does not perform any side effects, it is safe to use `BUG_ON`. The second rule takes care of this case. It furthermore disables the isomorphism that allows a call to `unlikely` be replaced with its argument, as then the second rule would be the same as the first one.

```
@@
expression E,f;
@@

(
   if (<+... f(...) ...+>) { BUG(); }
|
-  if (E) { BUG(); }
+  BUG_ON(E);
)

@ disable unlikely @
expression E,f;
@@

(
   if (<+... f(...) ...+>) { BUG(); }
|
-  if (unlikely(E)) { BUG(); }
+  BUG_ON(E);
)
```

For instance, using the semantic patch above, Coccinelle generates patches like the following one.

```
--- a/fs/ext3/balloc.c
+++ b/fs/ext3/balloc.c
@@ -232,8 +232,7 @@ restart:
                     prev = rsv;
            }
         printk("Window_map_complete.\n");
-        if (bad)
-                 BUG();
+        BUG_ON(bad);
   }
 #define rsv_window_dump(root, verbose) \
         __rsv_window_dump((root), (verbose), __FUNCTION__)
```

11.4   Look for `NULL` dereference

This SmPL match looks for NULL dereferences. Once an expression has been compared to NULL, a dereference to this expression is prohibited unless the pointer variable is reassigned.

Original

```
1 foo = kmalloc(1024);
2 if (!foo) {
3    printk ("Error %s", foo->here);
4    return;
5 }
6 foo->ok = 1;
7 return;
```

Semantic match

```
1 @@
2 expression E, E1;
3 identifier f;
4 statement S1,S2,S3;
5 @@
6
7 * if (E == NULL)
8 {
9    ... when != if (E == NULL) S1 else S2
10        when != E = E1
11 * E->f
12    ... when any
13    return ...;
14 }
15 else S3
```

Matched lines

```
1 foo = kmalloc(1024);
2 if (!foo) {
3    printk ("Error %s", foo->here);
4    return;
5 }
6 foo->ok = 1;
7 return;
```

11.5 Reference counter: the of_xxx API

Coccinelle can embed Python code. Python code is used inside special SmPL rule annotated with `script:python`. Python rules inherit metavariables, such as identifier or token positions, from other SmPL rules. The inherited metavariables can then be manipulated by Python code.

The following semantic match looks for a call to the `of_find_node_by_name` function. This call increments a counter which must be decremented to release the resource. Then, when there is no call to `of_node_put`, no new assignment to the `device_node` variable `n` and a `return` statement is reached, a bug is detected and the position `p1` and `p2` are initialized. As the Python only depends on the positions `p1` and `p2`, it is evaluated. In the following case, some emacs Org mode data are produced. This example illustrates the various fields that can be accessed in the Python code from a position variable.

```
1  @ r exists @
2  local idexpression struct device_node *n;
3  position p1, p2;
4  statement S1,S2;
5  expression E,E1;
6  @@
7
8  (
9  if (!(n@p1 = of_find_node_by_name(...))) S1
10 |
11 n@p1 = of_find_node_by_name(...)
12 )
13 <... when != of_node_put(n)
14     when != if (...) { <+... of_node_put(n) ...+> }
15     when != true !n  || ...
16     when != n = E
17     when != E = n
18 if (!n || ...) S2
19 ...>
20 (
21    return <+...n...+>;
22 |
23 return@p2 ...;
24 |
25 n = E1
26 |
27 E1 = n
28 )
29
30 @ script:python @
31 p1 << r.p1;
32 p2 << r.p2;
33 @@
34
35 print "* TODO [[view:%s::face=ovl-face1::linb=%s::colb=%s::cole=%s][inc.
       counter:%s::%s]]" % (p1[0].file,p1[0].line,p1[0].column,p1[0].column_end,
       p1[0].file,p1[0].line)
36 print "[[view:%s::face=ovl-face2::linb=%s::colb=%s::cole=%s][return]]" % (p2
       [0].file,p2[0].line,p2[0].column,p2[0].column_end)
```

Lines 13 to 17 list a variety of constructs that should not appear between a call to `of_find_node_by_name` and a buggy return site. Examples are a call to `of_node_put` (line 13) and a transition into the then branch of a conditional testing whether n is `NULL` (line 15). Any number of conditionals testing whether n is `NULL` are allowed as indicated by the use of a nest `<...   ...>` to describe the path between the call to `of_find_node_by_name`, the return and the conditional in the pattern on line 18.

The previously semantic match has been used to generate the following lines. They may be edited using the emacs Org mode to navigate in the code from a site to another.

```
1 * TODO [[view:/linux-next/arch/powerpc/platforms/pseries/setup.c::face=ovl-
    face1::linb=236::colb=18::cole=20][inc. counter:/linux-next/arch/powerpc/
    platforms/pseries/setup.c::236]]
2 [[view:/linux-next/arch/powerpc/platforms/pseries/setup.c::face=ovl-face2::
    linb=250::colb=3::cole=9][return]]
3 * TODO [[view:/linux-next/arch/powerpc/platforms/pseries/setup.c::face=ovl-
    face1::linb=236::colb=18::cole=20][inc. counter:/linux-next/arch/powerpc/
    platforms/pseries/setup.c::236]]
4 [[view:/linux-next/arch/powerpc/platforms/pseries/setup.c::face=ovl-face2::
    linb=245::colb=3::cole=9][return]]
```

Note : Coccinelle provides some predefined Python functions, *i.e.*, `cocci.print_main`, `cocci.print_sec` and `cocci.print_secs`. One could alternatively write the following SmPL rule instead of the previously presented one.

```
1 @ script:python @
2 p1 << r.p1;
3 p2 << r.p2;
4 @@
5
6 cocci.print_main(p1)
7 cocci.print_sec(p2,"return")
```

The function `cocci.print_secs` is used when there is several positions which are matched by a single position variable and that every matched position should be printed.

Any metavariable could be inherited in the Python code. However, accessible fields are not currently equally supported among them.

11.6 Filtering identifiers, declarers or iterators with regular expression

If you consider the following SmPL file which uses the regexp functionality to filter the identifiers that contain, begin or end by foo,

```
1  @anyid@
2  type t;
3  identifier id;
4  @@
5  t id () {...}
6
7  @script:python@
8  x << anyid.id;
9  @@
10 print "Identifier: %s" % x
11
12 @contains@
13 type t;
14 identifier foo ~= ".*foo";
15 @@
16 t foo () {...}
17
18 @script:python@
19 x << contains.foo;
20 @@
21 print "Contains foo: %s" % x
```

```
23 @endsby@
24 type t;
25 identifier foo ~= ".*foo$";
26 @@
27
28 t foo () {...}
29
30 @script:python@
31 x << endsby.foo;
32 @@
33 print "Ends by foo: %s" % x
34
35 @beginsby@
36 type t;
37 identifier foo ~= "^foo";
38 @@
39 t foo () {...}
40
41 @script:python@
42 x << beginsby.foo;
43 @@
44 print "Begins by foo: %s" % x
```

and the following C program, on the left, which defines the functions foo, bar, foobar, barfoobar and barfoo, you will get the result on the right.

```
1  int foo () { return 0; }
2  int bar () { return 0; }
3  int foobar () { return 0; }
4  int barfoobar () { return 0; }
5  int barfoo () { return 0; }
```

```
1  Identifier: foo
2  Identifier: bar
3  Identifier: foobar
4  Identifier: barfoobar
5  Identifier: barfoo
6  Contains foo: foo
7  Contains foo: foobar
8  Contains foo: barfoobar
9  Contains foo: barfoo
10 Ends by foo: foo
11 Ends by foo: barfoo
12 Begins by foo: foo
13 Begins by foo: foobar
```

12 Tips and Tricks

12.1 How to remove useless parentheses?

If you want to rewrite any access to a pointer value by a function call, you may use the following semantic patch.

```
1  - a = *b
2  + a = readb(b)
```

However, if for some reason your code looks like `bar = *(foo)`, you will end up with `bar = readb((foo))` as the extra parentheses around `foo` are capture by the metavariable `b`.

In order to generate better output code, you can use the following semantic patch instead.

```
1  - a = *(b)
2  + a = readb(b)
```

And rely on your standard.iso isomorphism file which should contain:

```
1  Expression
2  @ paren @
3  expression E;
4  @@
5
6    (E) => E
```

Coccinelle will then consider `bar = *(foo)` as equivalent to `bar = *foo` (but not the other way around) and capture both. Finally, it will generate `bar = readb(foo)` as expected.