

QuLog: A flexibly typed logic and functional programming language with action rules - the *declarative* subset

Keith L. Clark and Peter J. Robinson

March 3, 2015

1 Introduction

QuLog is a declarative hybrid logic, functional, string processing language with an imperative action language sitting on top. Its declarative kernel was developed to complement our *guard* \rightarrow *action* rule based robotic agent programming language *TeleoR*. The latter is a major extension of Nilsson’s Teleo-Reactive robotic agent language. The primary use of declarative *QuLog* is to ‘program’ the agent’s *BeliefStore* and to inference evaluate the *TeleoR* rule guards, which are restricted *QuLog* queries. Its action layer is used to extend *TeleoR* rule actions with atomic *BeliefStore* update actions and inter-agent communication actions routed via our Pedro communications server. The multi-threaded agent architecture that concurrently executes *TeleoR* programmed tasks is implemented in imperative *QuLog*. This draft paper, which currently lacks references, describes the declarative subset of *QuLog*. A companion draft paper describes the layer of action rules and illustrates their use for communicating agent applications.

QuLog relations are *typed* and *moded*, like Mercury. Its functions are just typed as all their arguments must be given as ground terms (terms with no unbound variables). The types and modes are so that we can guarantee at compile time that there will be no runtime failures or errors due to wrongly typed arguments of relation or function calls, or due to arguments that *must* be given not having been computed before the call, particularly calls to *QuLog* primitives. They also enable us to guarantee at compile time that *TeleoR* rule actions will be fully ground and correctly typed when sent to a robotic resource. We need this to make *TeleoR* a serious robotic programming language. The result of these constraints is that all but a handful

of primitives of *QuLog*, and *most* program defined relations, ground all their arguments if they succeed. Functions always return ground term values.

To retain some of the flexibility of use of Prolog we have *union* (aka *disjunctive*) types. These allow us to say that a relation accepts either integers or atoms in some input argument position, with run-time tests to allow different rules to be given depending on which type of value is given in a particular call, as in Prolog. However in Prolog such run-time tests are limited to a small range of sub-types of the generic *term* type, such as *integer* and *atomic*. In *QuLog* run-time type tests can be done for all primitive types, any program defined type, as well as higher order types.

In *QuLog* sub-type values are also considered as valid values of all their super-types. This means that a sub-type value may be given as an argument where a super-type is specified, and returned as a value in lieu of a super-type value. The type checker ensures that all function calls are given ground arguments of their declared argument types, and will return a ground value of the declared value type. It checks that each relation call will have each argument an unbound variable, or a term of the declared type. It also checks that the argument will be a ground term *before* the call if moded ! (ground input), and will be ground *after* the successful evaluation of the call, if moded ? (ground output). Only if moded with a postfix ? (non-ground output) is there no check beyond the type check.

There is a sub-type lattice that includes every data value in *QuLog*, covered in Section 6.5. Relations and actions (defined by one or more action rules) have a sub-type relationship based on this data sub-type relationships and their modes of use, specified by their moded argument types. For functions the sub-type relation just depends on argument and value types as there is only one mode of use - all arguments ground and ground value returned. The sub-type relationship for functions and relations is covered in Section 9. As an example of the role of modes, suppose the higher order argument type specifies the relation given as an argument will only be used to check integer values. We can pass in a relation that checks *or generates number* values. A more flexible check or generate relation that handles numbers can be used where only integer values will be tested.

QuLog also has support for *meta-level programming*. Compound terms representing relation calls can be constructed from and mapped into lists, as in Prolog. These call terms can be passed around as data terms and then evaluated as the relation calls that they denote. But unlike Prolog, this is done in a type and mode safe way. No meta-call is allowed in *QuLog* that does not satisfy the type and mode constraints of the relation being called. This is discussed in Section 7.4.

QuLog is a fully integrated LP/FP language in that function calls can appear as or inside arguments to relation calls, and relational queries can be used as tests in function rules, and in set expression arguments of function calls.

The compiler does type *checking* of function and relation definitions, assisted by type *inference* for data terms and variables. We believe that type declarations, linked with mode of use declarations for relations and action procedures, are very useful *active* documentation of the program. Also, because we have union types and sub-types, type inference on code could be very complex in some cases. Type inference on data terms will assign the term the minimum type in the sub-type lattice.

This introduction to declarative *QuLog* assumes familiarity with Prolog and higher order functional programming, as in languages such as Haskell or Scala. Sections 2, 3, 4 introduce its use as a typed relational programming, with its pattern match string processing support being covered in Section 4. Sections 5 and 8 deal with the functional subset and higher order programming. Sections 6, 9 and 10 deal with types, modes, the compile-time type/mode checking, and the need for auxiliary run-time type testing to filter out sub-type values. Section 7 introduces program defined types and illustrates their use. Type declarations for types with finite extensions double as both testers and *generators* for instances of the type. We exemplify *QuLog*'s debugging support in Section 3, and give a glimpse of its imperative action rules in Section 7.

2 *QuLog* Relational Subset

2.1 Python influenced syntax

QuLog does *not* have an operator precedence syntax and its syntax is not extensible. As in Prolog, a functor or predicate is written immediately next to its tuple of arguments as in `p(...)`.

There is *no need* to use a full stop followed by a white space character, as in Prolog, to separate the relation rules (aka clauses), function and action rules, type definitions and type declarations. Instead we borrowed an idea from Python and made what is a normal program layout format for a Prolog program - each clause starting on a new line - a syntax requirement.

All the above program statements *must* begin at the left end of a new line. Each can be continued over several lines where all but the first line is *indented* by at least one space or tab. Starting rules at the left end of a newline, and indenting a continuation of a rule by several spaces, or

one tab, is normal Prolog program layout. Our indentation indicator for continuation of a rule encourages this, making programs more readable. However, as a gesture towards Prolog programmers, including ourselves, a full stop followed by a newline, or spaces and a newline, *may* also be used at the end of a statement. It is *ignored* by the *QuLog* parser. Even if you use fullstops as terminators, you *cannot* put two clauses or rules on the same line of the program file.

Apart from requiring a predicate or functor to be adjacent to its (...) bracketed arguments, and treating a space or tab at the beginning of a new line as a continuation marker, *QuLog* is tolerant of spaces. They should be used to aid readability of the program.

As in Prolog, alphanumeric names beginning with an upper case letter or underscore `_`, or underscore on its own, are variables. To make such a name an **atom** (aka symbol), or the name of a relation or function, it can be singly quoted as in `'Peter'`, `'Father_of'`, `'Fact'`.

In contrast, an alphanumeric name that begins with a lower case letter, which can contain under-scores, is an **atom** and is used to denote individual things or names of relations, functions or actions. Surrounding such a name with single quotes has no effect at all, and they will be dropped when the atom is displayed. However if we want to use the atom name of a relation, function or actions as an atomic data value, it must be preceded by a single back-quote.

2.1.1 Need for back quoting

In the main, it is advisable to keep the atoms used for function, relation and action names disjoint from the names of things, like people. The use of the name of a program defined relation or function as an argument of a call or inside some term data structure such as a list, is assumed to denote that function of relation. It is not of type **atom**. To use such a name as an atom it must be back-quoted using ```. So, suppose we have a definition in our *QuLog* program of the `parent:(atom,atom)` relation. `[parent]` is a list of type `[(atom,atom)<=]` (the `<=` indicates that we have a relation over pairs of atoms), and its single element is the `parent` relation. However, `[`parent]` is a list of type `[atom]`, and its single element is the atom ``parent`.

This back quoting to map the name of a relation, denoting that relation, into an atom data value is also needed to map a *QuLog* type name into an atom data value. This is because type names, indeed any ground type expression, is a valid data value in *QuLog*. The type of the data value `nat`, is `typeE(nat)` and in general the type of any type expression term

TE is `typeE(TE)`. The list `[`int`,`num]` has type `[atom]` - list of atoms. `[int,num]` has type `[typeE(int) || typeE(num)]` - a list of specific type expressions. `||` is the type union operator.

There is no need to back quote a parameterised type name like the `tree(T)` type of Section 6.2, as `tree` without a type argument is not a type expression, `tree(int)` is. Similarly, `typeE` does not need to be quoted to be used as an atom. This also applies to the constructor functors of a type. The `tree(T)` type has two constructors named `empty` and `tr`. Both can be used as atoms without back quoting.

To summarise, all relation, function and action names need to be prefixed with a ``` to be used as atoms. All non-parameterised type names need to be prefixed with a ``` to be used as atoms. The sets of relation names, function names, action names, type names, and type constructor names need to be disjoint. A name cannot be overloaded and used, say, for two relations differing with respect to their number of arguments.

2.2 Moded type declarations

All relations (other than those with no argument), and all functions *should* have their types declared. If not declared, a default declaration is assumed where all arguments have *term* type, and for a function the value has *term* type. `term` is the top of the *QuLog* data type lattice (there is a separate type lattice for relation, action and function types). If an argument has type `term`, run-time type tests then have to be inserted in clauses and rules, before any *QuLog* primitive is used that requires a value of particular type, such as the arithmetic inequalities `<`, `>` etc. We exemplify this in Section 6.6.

The default mode for each argument of an undeclared relation is ground input, so an undeclared relation will be assumed to be a test only relation that needs to be called with ground arguments (terms not containing variables) but of unconstrained type. For Prolog relations, the default moded type of each argument relation would be specified using the *QuLog* moded type `term?`.

`term` says that the argument can be *any* Prolog term: a variable, `atomic` value - a `nat` (non-negative integer), an `int` (integer), a `num` (decimal number), an `atom` or a `string` (a double quoted sequence of any characters), or a `c(...)` compound term, which may contain variables. A postfix `?`, says a successful evaluation of the call *may not* result in a ground value for the argument term - may not generate any value, or may not fully instantiate a given compound term. Compound terms that may contain variables

in argument terms we call template terms. A *QuLog* primitive `template:term?` will test for being a template term.

This contrasts with use of a prefix `?`. This tells us that the argument, even if given as an unbound variable, or template term, will be fully instantiated to a ground term on success of the call. You should use the `?` mode whenever a ground value is guaranteed to be returned if not given as the argument value. Use of a variable in such an argument position, allows the variable to be later passed into a function call - for all function call arguments must be ground - or to be used as the value of an argument of a relation or action call that requires a ground value. If used as an argument of a robotic action of a *TeleoR* rule, it guarantees this action argument will be ground. If you want to emphasise that an argument for a call to a relation or action **must** be ground, you can use the prefix annotation `!`. If you mode an argument with a prefix `?`, the *QuLog* type/mode checker will ensure that a ground value will be returned and gives an error message if this is not the case.

To summarise, a 3-argument relation `rel` that does not have a type declaration will have an assumed, test only use, type declaration

```
rel: (!term,!term,!term) <=
```

The `<=` tells us `rel` is a relation. To indicate that a relation `rel2` has arguments of type `term` where the first must be given as a ground term, the second can be a variable or any term, but will become ground on success of the call, and the third may likewise be a variable or any term and may not be ground on success of the call, you can use a declaration

```
rel2: (!,?,term?) <=
```

Finally, to indicate that a relation `rel3`, had to be given a number as first argument, could be given but may also return an atomic value as second argument, and may be given a variable or list term containing only strings and variables and which may not become ground on success, use the declaration

```
rel3: (!num,?atomic,[string]?) <=
```

An unannotated type is implicitly `!` annotated - ground argument only. We could therefore drop the prefix `!` from the `num` type. Surrounding a type expression with `[...]` indicates a list of terms of that type.

A relation can have more than one type declaration. For example, suppose we know that when `rely` is called with an integer as first argument, and is given an integer second argument, it will always ground its last argument, we can use the disjunctive type declaration

```
rel3: (num,?atomic,[string]?) | (int,int,?[string]) <=
```

Here we have used the equivalent `[string]?`, with a `?` *after* the type, to indicate the list of strings argument may still contain variables *after* the call succeeds. We have also dropped the prefix `!` mode on `num` and `int` as this is the default for an un-moded argument.

2.3 Relation definitions

```
child_of: (?atom,?atom)<=                                % Type and mode declaration
child_of(mary, peter)
child_of(bill,peter)
child_of(peter,john)
.
.
age_is: (?atom,?atom)<=
age_is(peter,40)
age_is(mary,12)
.
.
male: (?atom)<=
male(peter)
.
.
female: (?atom)<=
female(mary)
.
.
descendant_is: (!atom,?atom)<=    % First arg. of a call must be a given
descendant_is(A,D) <= child_of(D,A)
descendant_is(A,D) <= child_of(C,A) & descendant_is(C,D)

ancestor_of: (?atom,?atom)<=    % Test or generate use of both arguments
ancestor_of(A,D) <= child_of(D,A)
ancestor_of(A,D) <= child_of(D,P) & ancestor_of(A,P)
```

The *input* (`!`) mode annotation for the first argument of the type/mode declaration for `descendant_is` is because when the name of the ancestor `A` is given it is more efficient than the `ancestor_of` definition. It is worse than the other definition if `D` is given and `A` is to be found. The `descendant_is` recursive rule 'walks' down the `child_of` links starting at the given `A`. In

contrast, the `ancestor_of` recursive rule 'walks' up the `child_of` links from a given descendant D. If D is not given then it finds some child/parent pair D,P and then walks up from that P. This is acceptable if A is not given as well, and all `ancestor_of` pairs need to be found, but if A is given the logically equivalent `descendant_is` relation should be queried.

If the programmer forgets and calls the less efficient `descendant_is` relation in a position where the first argument will not be given, the mode checker will give a mode use error so that the programmer can change the called relation. However, it will not object if `ancestor_of` is used where A will be given with D to be found, where `descendant_is` is more efficient. We are considering adding an optional `equiv` declaration to *QuLog*. Then, in the presence of

```
equiv ancestor_of, descendant_is
```

the mode checker would replace a call to `descendant_of`, where the first argument is not given, by a mode correct call to the equivalent `ancestor_of`, rather than raise a mode error. Conversely, it would replace a call to `ancestor_of` by a call to the more mode constrained equivalent relation `descendant_of`, when the mode analysis has determined that the first argument will be known at the time of the call. It would issue an advisory when either kind of replacement is made.

A simple list processing example is

```
doubleL: (![num],?[num])<=          % [num] is type of a list of numbers
doubleL([],[])
doubleL([N,..Nums], [DN,..DNums]) <=
    DN =! N*N & doubleL(Nums,DNums)
```

The `[num]` type expression for both arguments indicates this is a relation over lists of floating point numbers. The `!` mode annotation on the first argument indicates that in any call this argument must be a ground (variable free) list of numbers or integers. It can contain integers, as `int` is a sub-type of `num`. The lack of an annotation on the second argument indicates this may be an unbound variable in the call, for generate use, or that it can be a ground list of numbers for test use, or even a list of numbers and variables for mixed test/generate use of the numbers on the list. For example, a call

The `, ..` should be read as *followed by*. `[N,..Nums]` is read as *N followed by list Nums*, which may be the empty list `[]`. After the `, ..` we usually have a variable which denotes the tail list of terms following the number of elements that precede the `, ...`. A list pattern that does not have a variable after the

,... can be rewritten as one that does. $[X1, X2, \dots [X3, \dots L]]$ can be written $[X1, X2, X3, \dots L]$. $[X1, X2, X3, \dots]$ is equivalent to $[X1, X2, X3, \dots_]$. We can use the Prolog infix `|` instead of `,...` except that $[X1, X2, X3|]$ cannot be used as shorthand for $[X1, X2, X3|_]$

`=!` is *QuLog*'s generalisation of the Prolog `is`, and is explained in Section 5.6. It is also the structure pattern match operator that will evaluate function calls in both the pattern expression and the value expression, and it is the operator for non-deterministic pattern matching against string and list values. This last use is to split strings into sub-strings and lists into sub-lists satisfying certain constraints. It is illustrated in Section 4. `=!` is multi-purpose operator but is always used for some variant of matching with the right hand side a ground value. Even the above use is a simple match of an variable, that may or may not have a number value, against a number value $N*N$. For unification we must use `=`, discussed in Section 2.6.

It will be a match against a number value if N is a number. If it were not, or is an unbound variable, we would have a run-time error when the multiplication $N*N$ is attempted. This is why we have type and mode of use declarations in *QuLog*, checked against the source program, so such a run-time error will not occur. All uses of `double` will be checked by the mode analyser to ensure that for every call the first argument will always be an ground list of numbers, or integers.

The postfix `<=` after a relation type declaration can be dropped when the relation is defined in the same program file. From now on we shall drop this optional `<=` from relation type declarations. We shall also drop the `!` annotation from input argument types unless no type is given - meaning **term** type. However, when a program clause is displayed using the `show` command which we will exemplify in Section 3, the default `!` annotations and postfix `<=` will be displayed on the type declaration.

2.4 General form of relation rules

The 'if' operator for a *QuLog* conditional rule is `<=`. The 'and' operator is `&`. We choose to use `<=` as it is often used as the 'if' of predicate logic, and we believe that comma is overloaded in Prolog. Its use as the associative 'and' for conjunction also means that tuples are not fully supported as a data term in Prolog, as explained in Section 6.1.

Typically, a conditional rule for a relation is written

$$Head \leq Cond_1 \ \& \ Cond_2 \ \& \ \dots \ \& \ Cond_n$$

where **Head** is a predication of the form $rel(Arg_1, \dots, Arg_k)$, $k \geq 0$, and each

$Cond_i$ is a predication, or a negated predication using the negation-as-failure operator **not**.

In *QuLog* **not** is restricted to a *single* positive condition, as in Datalog. In a negated condition all but the variables starting with an *_* *must* have ground values by the time it is evaluated. The relation modes are used to check this will be the case at compile time. If not, a mode error is signalled. If the programmer wants a negated test that requires a conjunction, an auxiliary relation must be defined with body that conjunction. The negation is then applied to a call to the auxiliary relation.

2.5 General form of moded type declarations

QuLog's relation that have arguments must all be mode and type declared. A type+mode declaration for a k argument declarative relation has the form

$rel: (mt_1, \dots, mt_k) \leq$

Here each mt_i is a type expression t_i (often just a type name) with an optional **!** or **?** prefix and an optional **?** postfix. The absence of a mode annotation is equivalent to having given a prefix **!** annotation. Clauses defining a no argument relation **rel** do not need a type declared, but may be declared as **rel**:() \leq . When several relations have the same type we can use $rel1, rel2, rel3: (mt_1, \dots, mt_k) \leq$.

The \leq postfix operator indicates that *rel* is a declarative relation, as against a *QuLog* action procedure declaration (or a *TeleoR* procedure declaration when used with *TeleoR*). It is intentionally the 'if' connective used in relation definitions so it is easy to remember. As already mentioned, extra information that is not needed for a relation defined in the same program file. That *rel* is a declarative relation can be inferred from the form of its rules. However, it does no harm and is useful 'documentation' if you want to separate the type declaration from the relation defining rules, say having all the type information at the beginning of a program file. The **:** is to signal this is a type declaration not a fact about *rel*. It allows the \leq to be optional without risk of confusing the compiler.

The \leq *is needed* if you want to specify that the type of an argument of a higher order relation or function is a relation value. In that case you *must* use an \leq annotated type expression $(mt_1, \dots, mt_k) \leq$. If the relation type for an argument has no arguments, $() \leq$ is used. So,

test: ($() \leq$, **atom**) \leq

tells us that **test** is a relation with a no-argument relation as its first argument and an **atom** as its second, both of which need to be given.

2.6 Non-grounding mode

A postfix `?` is a relaxation of the prefix `?` mode, most useful when the argument type covers non-atomic terms. As for a `?` prefix moded argument, an unbound variable, or a non-ground term, may be passed as an argument value for a postfix `?` moded argument, but the call may not ground that argument. It can return the same template term argument, or another template that is a partial instantiation of the given template argument, or a ground instantiation. If an unbound variable is given as the argument in the call, it may still be an unbound variable (perhaps a different one) after the call.

If for some particular call `C`, of a `rel`, with a `?` postfix moded argument `A`, it is important that the call returns a ground value because `A` will be passed as a `!` moded argument to a later call, then the primitive `ground(A)` test should be used after `C` to force backtracking until `C` returns a ground value for `A`. Without such a `ground(A)` test, the mode checker will raise an error. If it is just required that `A` be bound to a term of the form `con(...)`, then the primitive `template(A)` can be used after the call `C`, to force backtracking to generate such a partial structure.

The postfix `?` mode is not often used for user programs. We have an example in Section 4.2. The unification primitive has moded type `term? = term?`. `=` takes any pair of terms `trm1` and `trm2`, and succeeds if it can instantiate each to single term `trm` by binding variables, *providing* no variable is bound to a term that contains that variable (the *occurs* check of unification). The `trm` result of unification may not be ground, it will be `trm1` or `trm2` if the other is a variable. `trm` is always a least instantiated common instance of `trm1` and `trm2`.

As an example, `f(X,g(Y,a))=f(h(U),g(V,U))` instantiates both to the non-ground `f(h(a),g(V,a))`, or the non-ground `f(h(a),g(Y,a))`.

The ability to return and pass to a later call a template term, that the later call either grounds or just further instantiates leaving as a non-ground term, is a powerful programming technique unique to logic programming which is preserved in *QuLog*. It comes at the cost of having to support unification rather than just pattern matching, and having to treat variables as first class values. However, because of the mode information in *QuLog* we can, for the most part, safely avoid the costly occurs check. We discuss such implementation issues further in the companion *QuLog* paper on its action rules. The retention of unification, and the consequent need to have variables as first class values, is a major difference between *QuLog* and Mercury.

2.7 Iteration over all solutions to a query

Rule bodies can contain *forall* conditions which have *explicitly* quantified variables. Their simple form is

```
forall(V1, ..., Vj)(exists EVars1 Conditions1 =>
                                exists EVars2 Conditions2)
```

with $j \geq 1$. The spaces before and after the tuple of quantified variables, and the implication, are not essential. The (...) brackets around the implication are essential but can be dropped around the sequence of quantified variables. The variables V_1, \dots, V_j are said to be *universally* quantified.

Each variable of V_1, \dots, V_j *must* appear in *Conditions₁* and *Conditions₂*. They must be variables *not* used anywhere else in the query or rule body in which the **forall** appears, and all other variables of the *forall*, except the explicitly existentially quantified variables **EVars₁** and **EVars₂**, and **_** variables, must have ground values when it is evaluated. If there are no existentially quantified variable for either the antecedent or consequent, the **exists** is dropped.

Each existentially quantified variable must have a name different from that of any other variable in the query or clause body in which the **forall** appears. Nested existentially quantifications are *not* allowed. They are not necessary as all the inner existential quantified variables can be put into a single outermost existential quantification.

$$(\text{exists } X \ (g(X) \ \& \ \text{exists } Y \ h(X,Y,Y)) \Rightarrow \dots)$$

is equivalent to

$$(\text{exists } X,Y \ g(X) \ \& \ h(X,Y,Y) \Rightarrow \dots)$$

The *forall* holds if for all bindings for its *UVars* quantified variables, given by a solution θ to *Conditions₁*, in which the explicit existential quantifications are ignored, then *Conditions₂* θ has at least one inferable answer for its existentially quantified variables¹.

¹If there is an explicitly existentially quantified consequent **exists EVars₂ Conditions₂** this will only have the variables in *EVars₂* unbound by the time *Conditions₂* is checked, that is the existential quantification is being used as a test. The compiler always transforms the consequent to **once Conditions₂** (see Section 2.10)

2.7.1 Example uses of forall

```
only_has_adult_sons: ?atom      % arg. can be a var or given as an atom
only_has_adult_sons(P) <=
    person(P,_,_) &      % This checks P is a male or female, or finds one
    forall C (child_of(C,P) =>
        exists A male(C) & age_is(C,A) & A>20)
```

```
person: (?atom,?atom,?nat)<=
person(P,male,A) <= male(P) & age_is(P,A)
person(P,female,A) <= female(P) & age_is(P,A)
```

We need two rules for `person` as *QuLog* has no disjunction.

2.8 Commitment to a rule

The infamous Prolog ! *cut* primitive is *not* allowed in *QuLog* rules or queries, but commitment to a rule can be expressed using a `::` commitment test *before* the `<=`, as in:

- *Head :: Commit <= Body* - do not use any later rule if *Commit* holds even if *Body* fails.
- *Head::Commit* - ditto, rule has no preconditions except *Commit*.
- *Head::* - do not use any later rule if the *Call* matches *Head*

Commit is a conjunction of predications and negated predications. It often only has one solution, and as such it is a query test extension of the head unification with a call. If it happens to have more than one solution only the first solution will be found. It may generate values for variables in non-input arguments of the *Head*, which may be used in *Body*

The recommended style in defining a relation *Rel* is to have mutually exclusive *Commit* conditions or rule *Heads* with perhaps a last default rule that implicitly has an initial negated call `not above_Rel(...)` to an auxiliary (but undefined) relation *above_Rel*. The *above_Rel* definition has a rule *above_Rell_R* constructed from the head term arguments and *Commit* tests of each rule *R* above the last rule, such that the use of the *above_Rell_R* rule would succeed exactly when the head unification and *Commit* test of *R* would succeed. The implicit negated condition `not above_Rel(...)` tests that none of the head term match/commit tests of the rules above can be inferred. We give an example below.

`::` can be read as *such that* and the *Commit* test following it is a ‘case’ condition for use of the rule augmenting the term matching tests given by any non-variable terms in the *input* and *in/out* argument positions in the rule head. Ideally the ‘case’ conditions are non-overlapping, but in practice they often are not.

Disjunction and `if ... then ... else ...` conditionals are *not* supported. Their purpose must be achieved by defining auxiliary relations with commit tests.

2.9 Example use of commit

```
max_of: num,num,?num                                % First 2 args implicitly !moded
max_of(M,N,Max)::M<N <= Max=N
max_of(M,N,Max)::M>N <= Max=M
```

The rules treat disjoint cases. As in Prolog the commit test of the second rule can be dropped and we would get the same behaviour. Implicitly the second rule then has the test `not above_max_of(M,N,Max)` where

```
above_max_of(M,N,Max) <= M =< N
```

i.e. in this case the implicit negated test is the equivalent of the given test `M>N`.

Note that the modes mean that the `max_of` relation will only be called if the first two arguments are given, thus preventing a runtime error on use of the arithmetic comparison tests. An attempt to use it when they will not both be number values - checked by the type and mode analyser - results in a compile time error.

2.10 Restricting a query to a single solution

QuLog has a `once` operator that can be used as a prefix to a single call or a bracketed conjunction of predications or negated predications (so no nested `once` calls or `forall`s), as in

```
once (Conditions)
```

When it is used only the first successful inference of *Conditions* is found. After it has been found, there is no backtracking inside the *Conditions* conjunction to find alternative inferences/solutions. Brackets are not needed if there is only one condition.

It is particularly useful if *Conditions* are being used to test and not to generate values. For example, the **once** in the body of a rule of the test only relation **has_son**.

```
has_son : ?atom
has_son(P) <= person(P,_,_) & once (child_of(C,P) & male(C))
```

If a call to the relation succeeds backtracking does not look for alternative proofs that the found P has an adult child.

2.11 Generating lists and sets from query conjunctions

QuLog has a list comprehension expressions that generate lists of terms, each one being the instantiation of a template term corresponding to a solution to a conjunction of conditions. It also has a set comprehension expression. Sets contain no duplicate terms. Sets and lists are distinct types in *QuLog*.

To generate a list of instances of some **Term** corresponding to solutions of some conjunction of conditions we can use an expression

```
[Term :: exists Vars Conditions]
```

Vars are all the variables in *Conditions* that do not appear in **Term**, or elsewhere in the clause or query in which the list expression appears. They are the local variables of *Conditions*. If there are no such variables, **exists Vars** is dropped.

All variables in *Conditions*, not in *Vars* or *Term*, must appear earlier in the clause or query and be such that they will have a ground binding by the time the expression is evaluated. This also applies to variables in *Term* not appearing in *Conditions*.

Conditions is a conjunction of predications and negated predications with the usual constraints on the negated predications. The predications may have list comprehension expressions and function call arguments. Finally, any variable in **Term** that will not have a ground value when the set expression is evaluated *must* appear in a grounding condition in *Conditions*. This ensures the generated list will only contain ground values.

The list comprehension expression denotes the list of all instances of *Term* such that *Conditions*, *as found*. That is, the instances of *Term* are generated by finding all the different inferences of *Conditions*, ordered as found by the backtracking search for all the different inferences. It can contain duplicates. It is like the Prolog **findall**.

2.11.1 Set comprehension

The set comprehension expression is very similar in structure but uses `{..}` braces instead of `[..]` brackets. The generated set will not contain any duplicates. However its internal representation is such that operations of `union`, `diiff` and `inter` (set union, difference and intersection) are efficiently implemented. Also a set will be displayed surrounded with `{..}` braces, with its elements ordered using the standard *QuLog* term ordering relation `@<`, which is the same as Prolog's.

A set expression has the form

`{Term :: exists Vars Conditions}`

and is subject to the same constraints regarding variable bindings as list comprehension expressions.

2.11.2 Conversion between lists and sets

To covert a list to a set we use the built in function `{}`. `{}([3,2,-6,3,-4])` is the set `{-6,-4,2,3}`. To convert a set to a list, that can be processed using list patterns, we use the built in function `[]`. So, `[]({2,9,2,-2,-4})` is the list `[-4,-2,2,9]`.

We can use `in` to access elements of both lists and sets. Used to access elements of a set the elements will be returned starting at the least element in size order. So, `E in {2,-3,7,2,0}` will successively bind `E` to -3, 0, 2 and 7. So we can think of a set as being ordered.

The set comprehension expression can be viewed as shorthand for

`{}([Term :: exists Vars Conditions])`

2.12 Example uses of comprehension expressions

```
children_are: ?atom,?{atom}           % {atom} is type for a set of atoms
children_are(P,L) <= person(P,_,_) & L =! {C :: child_of(C,P)}
                                           % L will be a set of children of P
```

The expression queries

```
| ?? {}([C1 :: child_of(C1,peter)] <> [C2 :: child_of(C2,mary)])
| ?? {C1 :: child_of(C1,peter)} union {C2 :: child_of(C2,mary)}
```


will both return the set of the children of either **peter** or **mary**. The set will be displayed with its elements ordered using the term order relation @< . Also, when the elements of a set **S** are accessed using **E in S** then **E** is bound to successive elements in @< term order.

For this example we could also have used the query

```
| ?? {C :: child_of_either(C,peter,mary)}
```

where `child_of_either` has been defined as

```
child_of_either: ?atom,?atom,?atom
child_of_either(C,P,_) <= child_of(C,P)
child_of_either(C,_,P) <= child_of(C,P)
```

<> is the list concatenation function of *QuLog*. It can be used for ‘glueing’ ground lists together, no matter how constructed. It can be used to concatenate lists of different types $[T1]$, $[T2]$. The result is a list of the union type of $[T1]$, $[T2]$, i.e. it is list of type $[(T1 \parallel T2)]$. \parallel is the type union operator. We will say more about union types in Sections 6.4, 6.9.

<> can also be used for splitting a given ground list subject to constraints. This is when used in the left hand side pattern of the =! operator, as illustrated in Section 4.2.

There is also an **append** relation primitive in *QuLog*. Its use covers both the appending and splitting use of <> , and more. It *must* be used, not <> , to concatenate or split lists containing variables or template terms. We shall illustrate its use in Section 4.2.

Here is an example that uses a necessary existential quantification.

```
adult_children_are: ?atom,{atom}
adult_children_are(P,L) <=
    person(P,_,_) &
    L =! {C :: exists A child_of(C,P) & age_is(C,A) & A>17}
    % L will be a set of name ordered adult children of P

children_in_age_order: (?atom,{(int,atom)})
children_in_age_order(P,L) <=
    person(P,_,_) &
    L =! {(A,C) :: child_of(C,P) & age_is(C,A)}
    % L will be a set of pairs ordered by increasing age
```

```

children_as_recorded: ?atom, ?[atom]
children_as_recorded(P,L) <=
    person(P,_,_) & L =! [C :: child_of(C,P)]

```

Note the change from {...} brackets to [...] brackets. L will be a list of names of children of P in the order that they are found by the `findall` Prolog primitive, i.e. in the fact order.

2.12.1 Ordering lists

The simplest way to order a list is to use the converter functions `{}`, `[]` in succession. The function

```

order: [T] -> [T]
order(L) -> []({}(L))

```

will convert L in to list ordered by increasing value with all duplicates removed.

2.13 Relational queries

You have probably inferred by now that `| ??` is the *QuLog* interpreter prompt to enter a query or command. In Section 3 we shall illustrate the entering of commands such as those needed to *consult* a program file, to *show* particular definitions, and to *watch* calls to a particular relation. Here we say something about *QuLog*'s relation queries. These are quite different in *QuLog* from Prolog.

```

| ?? person(P,male,_) & child_of(C,P)
    % <fullstop><return> or <return><return> is query end

P=bill:atom                                % type of each answer binding is given
C=mary:atom
...                                         % the ... indicates backtracking to find next solution
.                                         % more answers
.
...
P=cheryl:atom                             % The 5th answer
C=jake:atom

```

If the response is to enter `..<return>` after the 5 answers have been displayed, up to 5 more will be displayed. If `..3` is entered, up to 3 more answers are displayed and you must respond again to get more. In general, if the response is `..n`, where **n** is an integer, **n** more answers are displayed if there are **n** more.

For the above query we could use `_` (underscore) for the age argument of `person` to suppress seeing its bindings. However, if we had a minimum age condition we could not do this. The query

```
| ?? person(P,male,A) & A>40 & child_of(C,P)
```

will display the **A** answer bindings as well as those for **P** and **C**. To see just the **P**, **C** bindings, we can prefix the query with these variables.

```
| ?? P,C ? person(P,male,A) & A>40 & child_of(C,P)
```

This will give us up to 5 pairs of bindings for **P** and **C** in the first instance. Notice that this is very like a list comprehension. If we did put a pair of the list comprehension brackets around the query, and wrapped **P**, **C** into a single term such as `(P,C)`, we would get *all* the answers displayed as a single list.

```
| ?? [(P,C):: exists A person(P,male,A) & A>40 & child_of(C,P)]
```

will display a list of pairs `[(bill,mary),...,(cheryl,jake),...]`. We have had to explicitly existentially quantify **A** which is implicitly existentially quantified in the relational query.

If we only wanted to see 3 answers of a query in the first instance we can prefix the query with that number. We can also use an existential quantification for the variables for which we do not want to see answer bindings instead of enumerating the ones for which we want to see answer bindings.

```
| ?? 3 ? exists A (person(P,male,A) & A>40) & child_of(C,P)
```

After the first 3 answers have been displayed we must enter a `..` to see the next 3, or a `..n`, with **n** an integer, to see the next **n**. We can also prefix a query with an integer greater than the default up to 5 answers displayed. We can also change this default using a command, as given in Section 3.

We can just give the number of solutions bindings we want to see in the first instance without giving a sequence of variables. The query

```
| ?? 1 ? person(P,male,A) & A>40 & child_of(C,P)
```

behaves exactly like a Prolog query. You will need to enter `..` to see each successive answer.

With all the query forms, if you enter `<return>` when able to enter `..` to see more answers (if any), the query evaluation will terminate. You will again get the top level prompt `| ??` for a new expression, relation query or command to be entered. Expression queries and commands are discussed further in Sections 5.2 and 3, respectively.

```
| ?? child_of(_,P). % <fullstop><return> used to mark query end

C=peter:atom
...
C=peter:atom
...
C=john:atom
...
.           % More answers
.
```

Notice that we get the answer `C=peter:atom` twice as there are two ways of inferring that `peter` is a child of someone using the facts of Section 2.3. *QuLog* also attaches its minimal type of each answer binding.

3 Query line commands and debugging

```
| ?? consult file           % Double return used to mark end of command

success
```

will read in and type and mode check the *QuLog* program in *file* giving useful syntax, mode and type error messages, and perhaps suggestions for re-programming. `[file]` cannot be used as that is just a list expression to *QuLog* and will generate the answer `[file]:[atom]`.

```
| ?? show BeginName.      % For example show chi.
```

will complete the *BeginName* to names of defined functions and relations that start with *BeginName* and display them with their type declarations. `show` on its own will display them all. Unlike most Prologs they will have their programmer given variable names of the source file.

```
| ?? types.
```

will display all the type definitions and type declarations. You can also display a specific type definition using `type TypeNm`.

```
| ?? set_num_answers n.
```

will change the default number of answers displayed before a *continue* input `..` is required from 5 to *n*. *n* must be a positive integer.

```
| ?? watch relation.
```

will display every call to *relation*, say which rule it unifies with, and give the bindings for the variables of the call and the rule head. The rule head variables will be of the form *var_i* where *var* is the variable name used in the relation rule and *i* is an integer. If the call is retried this is also displayed along with the same information for the next rule with which the call unifies, if any.

A variant `watchC relation` will also display the partially instantiated body of the rule being tried. `unwatch relation` will turn the watching of *relation* off. Any number of relations can be watched at the same time. Watching a relation in *QuLog* is a little like the spying feature of Prolog except Prolog's spy invokes the interactive tracer at the first call of a spied relation. The *QuLog* watch is not interactive. It is a substitute for inserting write statements in Prolog rules that most Prolog programmers actually use for debugging. This is what `watch` and `watchC` do under the covers. It cannot be done at the source level as writing terms is an action and can only be done inside an action rule of *QuLog*.

We give an example use of `watch` and `watchC`. For this we assume we have a program defined relation

```
app: ([T]?, [T]?, [T]?) | ([T], [T], ?[T]) | (?[T], ?[T], [T])
app([], L, L)
app([U|L1], L2, [U|L3]) <= app(L1, L2, L3)
```

that has the same rules as the builtin `append`, which cannot be watched. This is a polymorphic relation that can be used to append or split lists of any type *T*. The most general type declaration is the first. It tells us the relation will handle list terms containing variables and that it may not instantiate all these variables. The second two give extra information to the

mode checker. The first says that if the first two arguments are ground in a call, the third will be ground after a successful evaluation. The last says that the first two arguments will be ground after a successful call, if the third argument is given as a ground list.

```
| ?? watch app.
success                                     % app is now being watched

| ?? app([1,2,3],[4],L).
1:app([1, 2, 3], [4], L)
  Call 1 unifies clause 2
  input U_0=1 L1_0=[2,3] L2_0=[4]  output L=[1,..L3_0]
2:app([2, 3], [4], L3_0)
  Call 2 unifies clause 2
  input U_1=2 L1_1=[3] L2_1=[4]  output L3_0=[2,..L3_1]
3:app([3], [4], L3_1)
  Call 3 unifies clause 2
  input U_2=3 L1_2=[] L2_2=[4]  output L3_1=[3,..L3_2]
4:app([], [4], L3_2)
  Call 4 unifies clause 1
  input L2_3=[4]  output L3_2=[4]
4:app([], [4], [4]) succeeded
3:app([3], [4], [3, 4]) succeeded
2:app([2, 3], [4], [2, 3, 4]) succeeded
1:app([1, 2, 3], [4], [1, 2, 3, 4]) succeeded
L = [1, 2, 3, 4]:[nat]
...                                     % Backtracking to look for more answers
1:app([1, 2, 3], [4], L) seeking another proof
2:app([2, 3], [4], L3_0) seeking another proof
3:app([3], [4], L3_1) seeking another proof
4:app([], [4], L3_2) seeking another proof
  checking clauses after 1 for call 4
4:app([], [4], L3_2) no (more) proofs
  checking clauses after 2 for call 3
3:app([3], [4], L3_1) no (more) proofs
  checking clauses after 2 for call 2
2:app([2, 3], [4], L3_0) no (more) proofs
  checking clauses after 2 for call 1
1:app([1, 2, 3], [4], L) no (more) proofs
```

```

| ?? watchC app.
success                                     % Next level of watching turned on for app

| ?? app([U,2,V],[9],[1,W,3,..L]).
1:app([U, 2, V], [9], [1, W, 3,..L])
  Call 1 unifies clause 2
  input U_0=1 L1_0=[2, V] L2_0=[9] L3_0=[W,3,..L] output U=1
  Clause body is: app([2, V], [9], [W,3,..L])
2:app([2, V], [9], [W, 3|L])
  Call 2 unifies clause 2
  input U_1=2 L1_1=[V] L2_1=[9] L3_1=[3,..L] output W=2
  Clause body is: app([V], [9], [3,..L])
3:app([V], [9], [3,..L])
  Call 3 unifies clause 2
  input U_2=3 L1_2=[] L2_2=[9] output V=3
  Clause body is: app([], [9], L)
4:app([], [9], L)
  Call 4 unifies clause 1
  input L2_3=[9] output L=[9]
  No clause body
4:app([], [9], [9]) succeeded
3:app([3], [9], [3, 9]) succeeded
2:app([2, 3], [9], [2, 3, 9]) succeeded
1:app([1, 2, 3], [9], [1, 2, 3, 9]) succeeded
U = 1:nat
V = 3:nat
W = 2:nat
L = 9:[nat]
...
                                     % Backtracking to look for more answers
1:app([U, 2, V], [9], [1,W,3,..L]) seeking another proof
2:app([2, V], [9], [W,3,..L]) seeking another proof
3:app([V], [9], [3,..L]) seeking another proof
4:app([], [9], L) seeking another proof
  checking clauses after 1 for call 4
4:app([], [9], L) no more proofs
  checking clauses after 2 for call 3
3:app([V], [9], [3,..L]) no more proofs
  checking clauses after 2 for call 2
2:app([2, V], [9], [W,3,..L]) no more proofs
  checking clauses after 2 for call 1

```

```
1:app([U, 2, V], [9], [1,W,3,..L]) no more proofs
```

```
| ?? unwatch app.
```

```
success
```

```
% Watching of app is turned off
```

A watched relation does not need to be called at the query level. It could be called deep inside some query evaluation.

3.1 Fact updating commands

QuLog action rules can update facts for relations defined solely by a sequence of facts for relations declared to be **belief** relations. This is similar to the Prolog **dynamic** except that **belief** relations can only be defined by a sequence of ground facts.

Suppose we had type declared the **child_of**, **male**, **female** and **age** relations using

```
belief child_of:(atom,atom), female:(atom), male:(atom)
```

The **belief** keyword tells the *QuLog* system that these four relations can be updated after the program file has been consulted. They can be updated both in action rules and command queries to the *QuLog* interpreter. Since belief relations always comprise a sequence of fact, the mode for each argument is ? so does not have to be given in the type declaration.

The belief updaters are

```
remember Fact % Add Fact as new last fact for its rel. if not present
```

```
rememberA Fact % Add Fact as new first fact for its rel. if not present
```

```
forget FactPtn % Remove first Fact matching FactPtn if present
```

```
replace FactPtn by Fact % Remove first Fact matching FactPtn and  
% add Fact as a new last fact for its relation  
% Fact and FactPtn can share variables
```

```
replaceA FactPtn by Fact % Remove first Fact matching FactPtn and  
% add Fact as a new first fact for its relation
```

For both replace actions the replacing fact will *not* be added if already present.

We can enter query commands such as


```
| ?? remember child_of(olive,sammy); remember female(sammy);
    remember female(olive)
```

to add three new dynamic facts. We have used ; to separate the actions of the command as this is what must be used in *QuLog* action rules.

Suppose we also had facts for the relation `spouse:(atom,atom)<=`. We could do an update recording a divorce and re-marriage using

```
| ?? replace married(bill,_) by married(bill,mary)
```

4 Non-deterministic String and List Pattern Matching

Strings in *QuLog* are the same as C strings - packed sequences of bytes. They are not lists of byte codes as in most Prologs.

`=!` can also be used as a string or a list pattern match operator to split strings into substrings, or lists into and sublists, satisfying certain conditions. Its uses are

StringPattern `=!` *StringExpression*

ListPattern `=!` *ListExpression*

4.1 String matching

For string matching we use `++` for string concatenation and splitting. We use `<>` for list appending and splitting. When one or more `++` or `<>` operators appear at the top level of the left hand side of `=!` it acts as a *non-deterministic* or multi-valued match operator. There may be several bindings of variables appearing in the left hand side pattern that make both the pattern and the expression have identical string or list values. As a simple example,

```
S1 ++ S2 =! "hello"
```

has seven different possible solutions for `S1` and `S2` starting with `S1=""`, `S2="hello"`, then `S1="h"`, `S2="ello"`, and ending with `S1="hello"`, `S2=""`.

In the pattern we can have

```
Var?Test(..,Var,..)
```

components that restrict the values assigned to `Var` to those that satisfy `Test`. Failure of the test will cause backtracking to find an alternative value for `Var`. When there are no more alternatives it causes an attempt to find alternative values for variables appearing to the left of this condition in the pattern, ultimately resulting in a failure of the entire `=!` non-deterministic pattern match. For list matching the variable `Var` can be replaced by a list pattern such as `[E]` or `[E|L]` where the associated `Test` can contain all the variables of the pattern. As an example, the pattern match

```
S1?S1\="" ++ S2?S2\="" =! "hello"
```

excludes the first and last solutions we gave above.

We will eventually have full regular expression pattern components of the form

```
V~"RegExp"?Test(...,V,...)
```

where `RegExp` is a regular expression filtering out the allowed string values of `V` before the `Test(...,V,...)` is applied.

For example `V~"(a | b)*c"` only allows `V` to match a string beginning with any length sequence of `a`'s and `b`'s, including none, followed by a `c`.

Here are some uses of string patterns to ‘parse’ strings representing sentences into lists of string words. Such string processing can be a precursor to Definite Clause Grammar (DCG) parsing of lists of string words.

```
spaces: string
spaces(S) :: " " ++ RS?spaces(RS) =! S    % S has more than 1 space
spaces(" ")                                % We have a single space string
```

```
sepchar,endchar,sep_endchar: ?string
sepchar(" ")
sepchar(",")
sepchar(";")
```

```
endchar(".")
endchar("?")
endchar("!=")
```

```
sep_endchar(C) <= sepchar(C)
sep_endchar(C) <= endchar(C)
```

```

wordchar,word,seps: string                                     % Test only use
wordchar(S) :: 1 == #S & not sep_endchar(S)
% ie a single character string which does not contain a sep or end character

word(S) :: C?wordchar(C) ++ RS?word(RS) != S
word(S) :: wordchar(S) % Single char alternative for a string being a word

seps(Seps) :: 0?sepchar(0) ++ RSeps?seps(RSeps) != Seps
seps(Sep) <= sepchar(Sep)

words: !string,[string]
words(Str,[W]) :: W?word(W) ++ E?endchar(E) != Str
words(Str,[W,..Words]) ::
    W?word(W) ++ Seps?seps(Seps) ++ RStr?words(RStr,Words) != Str

words converts a string sentence beginning with a word terminated with an
end char, and with spaces and/or any number of sep chars between each
word, into a list of word strings. An example use is

| ?? words("Hello    Keith,,    how are you?", Ws).

Ws = ["Hello","Keith","how","are","you"]:[string]

```

4.2 List pattern matching

```

| ?? L1<>L2 != [1,2,3].

L1 = []:[T]          % [] is a list of unspecified type T
L2 = [1,2,3]:[nat]
...
L1 = [1]:[nat]
L2 = [2,3]:[nat]
...
L1 = [1,2]:[nat]
L2 = [3]:[nat]
...
L1 = [1,2,3]:[nat]
L2 = []:[T]

```

```

% or we can use
| ?? [U1,..L1]<>[U2,..L2] =! [1,2,3].

U1= 1:[nat]
L1 = []:[T]
U2 = 2:nat
L2 = [3]:[nat]
...
U1=1:nat
L1 = [2]:[nat]
U2 = 3:nat
L2 = []:[T]

| ?? [1,U]<>[3|L] =! [1,2,3,4].
U = 2:nat
L = [4]:[nat]

splitsAroundE: [T],[T],T,?[T],?[T]
splitsAroundE(L1,L2,E,OL1,OL2) <= OL1<>[E]<>OL2 =! L1<>L2

| ?? splitsAroundE([1,5,3],[8,5,7],5,OL1,OL2).

OL1 = [1]:[nat]
OL2 = [3,8,5,7]:[nat]
...
OL1 = [1,5,3,8]:[nat]
OL2 = [7]:[nat]

```

The right hand side of a list match can use set expressions to denote the component lists, e.g. *list_pattern* =! [E|cond(E)]<>L.

We could also code the above examples using the built in **append** relation to split lists as in Prolog. We think the above are more transparent. Here is the rule for **splitsAroundE** using the *QuLog* builtin concatenation relation **append**.

```

splitsAroundE(L1,L2,E,OL1,OL2) <=
    append(L1,L2,L12) &
    append(OL1,[E|OL2],L12)

```

In this case it is not too difficult to decipher but it is less clear than the use of the split pattern using <>. You have to look at the use of the variables quite

closely to determine what it is doing. It does not wear its ‘semantics’ on its sleeve unless you are well practiced at the use of **append** for list splitting.

append’s library definition is the same as the earlier **app**.

```
append: ([T]?, [T]?, [T]?) | ([T], [T], ?[T]) | (?[T], ?[T], [T])
append([], L, L)
append([U|L1], L2, [U|L3]) <= append(L1, L2, L3)
```

Here is a query we cannot express using **<>** as it uses a template term for the appending result.

```
append([1, U], [3], [V, 2|L]).
```

```
U=2:nat
V=1:nat
L=[3]:[nat]
```

5 *QuLog* Functional Subset

5.1 Example function definitions

```
fact: nat -> nat           % The factorial function
fact(0) -> 1
fact(N)::N1 =! N-1 & type(N1, 'nat) -> N * fact(N1)
```

Note we have the rather indirect way of testing $N > 0$ in the second rule. This is because use of the test $N > 0$ would not enable the type checker to infer that the value of the argument of the recursive call $N-1$ was a natural number as required. However, the presence of the run-time test `type(N1, nat)` gives it this assurance. Input arguments of functions do not need to be moded. They are all implicitly **!** mode. That is, they must always be ground terms when the function is called.

type: (top?, typeE(T))<= is a primitive of *QuLog* that allows the run-time type checking of any *QuLog* data value, or defined relation, function or action. It tests if the first argument belongs to the sub-type **T** of type **top**, denoted by the `typeE(T)` argument - the back quoted name of type **T**.

```
max: (num, num) -> num
max(M, N)::M >= N -> M
max(M, N) -> N           % implicitly when not M >= N, i.e. when M < N
```

```

appBetween: ([E],E,[E]) -> [E] % E is a variable, a polymorphic function
appBetween(L1,X,L2) -> L1<>[X]<>L

father_of: atom -> atom
father_of(C)::child_of(C,P) & male(P) -> P
father_of(C)::not child_of(C,_) -> no_father_

number_of_children: atom -> nat % nat is type of non-negative integers
number_of_children(P) -> #{C :: child_of(C,P)}

a_child_of: atom -> atom
a_child_of(P):: child_of(_,P) & Children=[C:: child_of(C,P)] ->
    memberAt(choose_int(1,#Children),Children)
a_child_of(P):: not child_of(_,P) -> no_child_

memberAt: (int,[T]) -> T
memberAt(1,[E|_]) -> E
memberAt(Pos,[_|L])::Pos>1 -> memberAt(Pos-1,L)
memberAt(_,_) -> bottom_ % Only value of the bottom builtin type

now:() -> int
now()::realtime(Now) -> Now

```

is a *QuLog* operator for finding the length of a list or string.

Auxiliary functions such as `father_of` can be defined for any use pattern of a relation that returns a single value, or a single tuple of values for some of its arguments, when certain other arguments are given. When there are several output values for a relation for a given input we can define a function such as `a_child_of` that will randomly choose a child from the recorded children of a parent `P`. `choose_int` is a *QuLog* non-deterministic primitive function which returns an integer somewhere in the inclusive range of integers given by its two integer arguments.

Notice that `memberAt` returns a special value `bottom_` when the position is either given as a negative integer or it is greater than the length of the list first argument. `bottom_` is the sole atom of the `bottom` primitive enumerated type, which the type checker knows is a sub-type of every other type. It is the default value that should be used for polymorphic functions. For non-polymorphic functions, like `fact`, `father_of` and `a_child_of`, we can use a more meaningful default value for a function that is a legitimate value of the functions value type: integer 0 for `fact`, atoms `no_father_` and

`no_child_` for the other two. The application then deals with these function ‘completion’ values. Of course, `bottom_` can also be used as a default value instead.

`realtime: int <=` is a *QuLog* relation primitive that accesses the host’s clock and returns the current time as an integer valued Unix epoch time in seconds. The `now` function is a non-deterministic function that allows this varying time value to be found using a function call. This may make functional programmers aghast, but is useful for robotic applications where time is important and the result from an expression evaluation pertinent to a robot’s behaviour may well be time dependent. Like `choose_int`, `now` is a *QuLog* primitive function.

At the other extreme, *QuLog* has two primitive fixed value functions, `e_:()->num` and `pi_:()=num`. The first returns a numeric value for the constant `e`, the second a numeric value for π .

5.2 Expression queries

```
| ?? 1+2.
3:nat
```

```
| ?? {1,-5,6} union {apple,1,4}.           % union is the set union
{-5,1,4,6,apple}:{int||atom}              % A set of ints or atoms
```

We have sets as a separate type from list with `union`, `diff` (first set minus the elements in the second) and `inter` (set intersection) for combining sets, and the `in` membership test for accessing their elements. Because of the way they are stored, testing whether a given term is in a set is more efficient than testing if its in a list of the same elements.

5.3 General form of a function rule

Functions are defined by sequences of condition/expression rules of the form

Head `:: Commit -> Expression`

The `::Commit` is optional. As with relation definitions it is a simple conjunctive query that augments the head/call pattern match to determine if that rule should be used. Unlike in most functional languages *Commit* may generate values for variables in *Expression* making the returned function value partly depend on current facts. Even without a *Commit* test, there is an implicit *commit* at the `->` of each function rule.

Note the similarity between the form of conditional rules for relations and the form of a function rule. The difference is just whether `<=` or `->` is used followed by a conjunction or an expression.

5.4 General form of function type declaration

Every function must have its type declared. A function declaration has the form

fun: $(t_1, \dots, t_k) \rightarrow t$

where t is the value type. The arguments are always unannotated. Unlike relations, function arguments must always be given as ground values. The brackets around the sequence of argument types can be dropped when there will be no ambiguity (the last type t_k is not a function type).

Zero argument functions have a type $() \rightarrow \mathbf{t}$ and must be defined and called using `f()`. This is *different* from no argument relations. They are *not* type declared and *must* be defined and called using just their name, as in Prolog. E.g. `no_args_rel <=`

If a function `f` of type

f: $t_1, \dots, t_k \rightarrow t$

was instead defined as a $k+1$ relation `fr`, its type declaration would be

fr: $!t_1, \dots, !t_k, t \leq$

For example, if we defined `fact` as a binary relation `factr`, which we would need to do in Prolog, its type would be `factr: !int, int` and could be used to find or check factorial values for given integers.

There is a comprehensive set of primitive functions: the usual arithmetic operators of Prolog, trigonometric functions, integer division, integer division remainder, `sqrt`, `exp` etc. All these can be called inside or as arguments to relation calls reducing the need for `=!` calls to find the value of expressions.

5.5 Expression use, implicit and explicit evaluation

Set expressions, list and string concatenation expressions, arithmetic expressions and any function calls can appear as or inside the term arguments of relation calls in queries and rule bodies, but *not* in rule heads. They are fully evaluated before the relation call is made. There is no lazy evaluation in *QuLog*.

As an example, we could have a call of the form


```
rel([X::q(X,fact(Y)+2)])
```

where Y is a variable that must have been given an numeric value before the *rel* call. *q* will have a moded type such as $q:(?int,int)$ and *rel* a moded type $rel:([int])$. Suppose $Y=3$ at the time the *rel* call is made and *fact*(3) will return 6. *rel* is called with argument the integer list value of $[X::q(X,8)]$.

We can also explicitly evaluate expressions. Two expressions can be evaluated and checked for identical values using $Exp1==Exp2$, *providing* they have the same type. The complement test $Exp1 \neq Exp2$ is subject to the same type constraint. It can only be used to test non identity of values for primitive and program defined types. The argument expressions for $==$ and \neq must both be ground before the comparison test will be evaluated, a constraint checked by the compiler. After evaluation of all function calls $==$ and \neq respectively reduce to identity and non identity tests for the computed ground term values.

5.6 Expression patterns

A call $Exp1 \neq Exp2$ can be used to evaluate and match any pair of expressions in the language with the constraint that *Exp2* will be ground (variable free) when the \neq condition is evaluated. If both expressions are ground, their values must be identical for the \neq test to succeed so the call is then equivalent to $Exp1==Exp2$, which should be used instead.

However, the intended use of \neq is when *Exp1* has variables that will not be given values by the time that $Exp1 \neq Exp2$ is called, where such variables are *not* arguments to function calls inside *Exp1*. If they are this is a mode error. The intended pattern match use of \neq , where only the left hand side can be a template term or unbound variable and the right hand side must be ground, is indicated by the *!*, the ground argument mode annotation, on the right side of the $=$. The two expression must have the same type. The compiler will give an error message if this is not the case.

When called, function calls in both argument expressions are evaluated and replaced by their ground values. An exception is calls to functions that return a higher order value such as *curry* and *curryr*, which are left as function call expressions. The other exception is top level calls to *++* string concatenation or *<>* list concatenation in the left hand side expression *Exp1*. Their occurrence at the top level signals that the \neq call is being used for non-deterministic string or list pattern matching as discussed in Section 4. In that case only function calls inside test conditions in *Exp1* are evaluated.

After the function call evaluation, the `=!` call becomes the pattern match call

```
match_(Exp1',Exp2')
```

where `Exp1'`, `Exp2'` are the evaluated expressions. It succeeds if the resulting template value of `Exp1'` *matches* the computed ground value of `Exp2'`, instantiating all unbound variables in `Exp1` with ground terms from `Exp2`.

`=!` significantly generalises the *is* expression evaluator of Prolog. An example use is

```
g(X,mod(U,V)) =! f(W)
```

where `U`, `V`, `W` will all have been given ground values before the call and `f(W)` is a function call that will return a `g` term structure. `mod` is the *QuLog* primitive for the modulus of an integer divide. Let us say we have `U=4`, `V=3`, `W=5` at the time the `=!` call is evaluated and that `f(5)` will return `g(5,1)`. The `=!` test becomes `match_(g(X,1),g(5,1))` resulting in `X=5`.

The unification primitives `=` and `\=` do *no* prior evaluation and there must be no function calls in the arguments. Function calls in arguments of the unification primitives is often a programmer error in that they should have used `=!` instead. Whenever the mode analyser comes across a `=` or `\=` with a function call in either argument, it gives an error message identifying the suspect unification with a suggestion that `=!` or `not ..=!` be used.

If there are higher order values in either of `Exp1'`, `Exp2'` that need to be matched against one another they must be denoted by identical expressions. Usually each expression is just the program given name of the definition of the higher order value such as `child.of` or `fact`. But if one expression is a function call that returns a higher order value and the other is the name of that higher order value the match will fail. As an example, one expression could be `uncurryr(curryr(child.of))` and the other `child.of`. The match will fail even though both expressions denote the same relation. So beware.

6 Primitive and Defined Types

6.1 Primitive Types

The primitive *QuLog* types are:

```
bottom atom nat int num string atomic
```

`[T] {T} (T1,...,Tk) term top`

The `Ts` are any type expression or a variable (for parameterised types). `[T]` is the list of terms of type `T`. `{T}` is the set of terms of type `T`, i.e. no duplicate identical terms. `(T1,...,Tk)`, $k > 1$ is a tuple of k terms of the given, possibly mixed types. `bottom` is the special enumerated type that has one data value `bottom_` which can be used when a polymorphic function is undefined for some subset of values of its input type, as in Section 5.1. `top` covers every value that can be passed as an argument or returned as a value in *QuLog*. It includes every data term and every higher order value - functions, relations and actions. `nat` is the natural number (non-negative) sub-type of `int`, the integer type.

There are problems using tuples in Prolog because of associativity of `'.'`. This is needed because it is also used as the 'and' connective in rules. Prolog converts what should be a three element tuple `(2,big,(4.7,"hello"))`, containing a last element that is a pair, into `(2,big,4.7,"hello")`, a four element tuple. But `(2,big,(4.7,"hello"))` remains a three element tuple in *QuLog*, with a two element tuple as last component.

If we have a relation or function that has one argument, that is a tuple, in its type declaration we use `((T1,...,Tk)). ((atom,int))<=` is the type expression for a unary relation with argument a pair that can be given or generated by a call to the relation. The type expression `((atom,int))->int` is for a function of one argument that is a pair.

`term` is the type of every data structure, including those of program defined types. `top` also includes higher order values: functions, relations and actions. These are the sub-type relationships

```
bottom < nat < int < num < atomic < term < top
bottom < atom < atomic
bottom < string < atomic
[T] ≤ [T'] if T ≤ T'
{T} ≤ {T'} if T ≤ T'
bottom < [T] < term
bottom < {T} < term
bottom < (T1,...,Tk) < term
(T1,...,Tk) ≤ (T'1,...,T'k) if T1 ≤ T'1 & ... & Tk ≤ T'k
```

If `term` is used as the only data type, the *QuLog* program is essentially untyped like Prolog. In fact, as mentioned in Section 2.2, a mode annotation `m` with no type is an abbreviation for `m term`. The programmer must then

do run-time type checking using the **type** primitive, as described in Section 6.5, before calling any builtin relation or function such as an arithmetic primitive that requires a more specific type.

6.2 Program defined types

The primitive types can be extended with program defined types including parameterised and recursive *constructor* types, *enumerated* types (sets of atoms), *range* types (sets of successive integers), *union* or *disjunctive* types (the values can be from any one of a set of alternative types). All defined types are sub-types of **term**. Examples are given below.

Different enumerated types must either be disjoint, or one must be a subset, hence a sub-type, of the other. Different range types must either be disjoint, or one must be a subset, hence a sub-type, of the other. The constraint regarding not being partially overlapping is checked by the type checker and it issues an error message if flaunted. This is because in *QuLog* every term must belong to just one minimal type in the sub-type lattice.

```
tree(T) ::= empty() | tr(tree(T),T,tree(T))
```

This defines a new type parameterised constructor type **tree(T)**, with labels which are any primitive or program defined data type T. [T] denoting the type of a list of elements of type T, is the only primitive parameterised constructor type. **empty** is a 0 argument constructor of **tree** terms, and **tr** is a 3 argument constructor.

```
gender ::= male | female
          % enumerated type gender of two atoms, is a sub-type of atom
          % it is not a constructor type as the alternatives are atoms
age ::= 0..110
          % defines age type as a range of nat values, is a sub-type of nat
digit ::= 0..9      % is a sub-type of age, hence of nat, int, atomic, term
genage ::= gender || age
          % defines genage as a union type of age nats and gender atoms
          % [male,67,female,89] has type [genage], a list of genage terms
          % whereas [(male,57),(female,89)] has type [(gender,age)]
```

We use **|** to separate the alternatives of an enumerated type and a constructor type as both are alternatives between data *values*. We use **||** for separating the alternatives of a type union as they are alternative data *types*.

6.3 Data constructors do not have a function type

We do not view `tr` as a function of type $(\text{tree}(T), T, \text{tree}(T)) \rightarrow \text{tree}(T)$ which Haskell does, nor do we view `empty` as a function of type $() \rightarrow \text{tree}(T)$. `tr` is not a function in the normal sense of returning a value *different from* the call term. `tr` is what is termed a *free* function. Its role is just to wrap its name around its arguments. We do not want `tr` to be passed as an argument value where what is required is a non-wrapper function of type $(\text{tree}(T), T, \text{tree}(T)) \rightarrow \text{tree}(T)$. In *QuLog* `tr` has type `atom`.

There is another more important reason for not assigning any type to names of data constructors. In order to do meta-programming in *QuLog* we need to have data terms ‘naming’ relation and action procedure calls. We can do this in a type safe way by having a double role for the program given names of the relations and action procedures as constructors of two special meta-types `relcall` and `actcall`. If constructors had function types, this simple way of ‘naming’ calls with template data terms that use the relation and action procedure names as data constructors would not be possible. Section 7.4, shows a use of the `relcall` type to program a very simple query evaluator that evaluates a list of `relcall` terms. It is example of what might be used inside an agent, with its behaviour programmed using *QuLog* action rules, to answer queries from other agents, as will be shown in the companion paper.

As an example, since `children_are` is the name of a relation of type $(?atom, ?\{atom\})$, there will a constructor `chid_is(atom, {atom})` of the special meta-type `relcall`. If, because of this, `children_are` had to have the type $(atom, \{atom\}) \rightarrow \text{relcall}$ we would have a clash with its already declared program type as a relation of moded type $(?atom, ?\{atom\})$. By not assigning a function type to each of a constructor type’s constructor functors we avoid this clash. Since relation and action names automatically become constructors of the compiler generated `relcall`, `actcall` types respectively, their names *must not* be used in any programmer defined types. If used, there will be a type error.

6.4 Sub-type relation for defined types

For a parameterised type such as `tree` we have

$$\text{tree}(T) < \text{term} \quad \text{tree}(T) < \text{tree}(T') \text{ if } T < T'$$

The general sub-type relationships for defined types $Type$, $Type'$ is

$$Type(T) \leq Type(T') \text{ if } T \leq T'$$

$$\begin{aligned}
Type \leq Type' \text{ if } & Type ::= atom_1 \mid \dots \mid atom_m \ \& \\
& Type' ::= atom'_1 \mid \dots \mid atom'_n \ \& \\
& \{atom_1, \dots, atom_m\} \subseteq \{atom'_1, \dots, atom'_n\}
\end{aligned}$$

$$Type < \mathbf{atom} \text{ if } Type ::= atom_1 \mid \dots \mid atom_m$$

$$Type \leq Type' \text{ if } Type ::= I..J \ \& \ Type' ::= M..N \ \& \ M \leq I \ \& \ J \leq N$$

$$Type < \mathbf{int} \text{ if } Type ::= M..N$$

$$\begin{aligned}
Type \leq Type' \text{ if } & Type ::= T_1 \mid \dots \mid T_m \ \& \\
& Type' ::= T'_1 \mid \dots \mid T'_n \ \& \\
& \forall i (1 \leq i \leq m) \exists j (1 \leq j \leq n) \ T_i \leq T'_j
\end{aligned}$$

In the first rule *Type* is the name of a parameterised constructor type with *n* argument types.

6.5 Run-time type checking

A run-time type checking primitive can be used to check if an expression, which could be a variable, has a *ground* value of a specific type - either primitive or programmer defined. It is useful for checking if a value is a particular sub-type of an expected argument type.

`type(Exp,!Type)`

in a program clause or function rule will check that the *Exp* has a ground value of type *Type*, where *Type* is a ground type expression - so no type variables.

One extreme of *QuLog* programming has no user defined types and makes much use of the `term` type and run-time type checks. This style is not recommended as it trades run-time failures of type testing conditions for compile time indications of a type error. If `term` is used as a catch all type, the programmer *must* insert type tests before the use of any *QuLog* primitive that requires a specific type, such as a `num`. For example, any call to an arithmetic function or arithmetic comparison relation must be preceded by runtime tests that ensure that the arguments will be numbers.

The type checker will take into account the occurrence of the run-time type tests acting as filters of all values that are not of the tested type to make

sure that *QuLog* primitives can only be passed arguments of their required types. If they are not present it will give a type error. This trades the compile time checking of programs that use more constrained type declarations, for run-time checks that will fail if an incorrect type of value would have been passed into a call to a primitive. The programmer then has to have extra rules that handle these run-time failures of clauses due to type test failures. Far better to let the compiler do all the work of type validation. The best approach is to declare as constrained a type as possible for each relation and function argument, and to avoid uses of run-time type tests except to select out sub-types of declared argument types, either before a call that requires a sub-type, or after a call that might generate a super-type of the type that must be returned from a clause. These type filtering uses of run-time type tests are discussed in Section 10. A test

```
type(Exp, Type)
```

without the prefix `!` on *Type*, will succeed even if the `Exp` argument is not ground providing its variables can be assigned types consistent with the given *Type*.

If `type` is used in a single condition top level query in the interpreter, *Type* may be a variable which will be bound to the inferred type of *Exp*. This is useful for learning about the type system.

There is also a three argument `typeC`, which can only be used as a single condition interpreter query. This will tell us the types that any variables in *Exp* must have for *Exp* to have type *T*. The third argument of a `typeC` call *must* be a variable. It will be bound to a list of type constraints for any variables in *Exp*, as exemplified below.

6.6 Some example uses of run-time type checks

As an example, this is the program for the factorial function defined as a relation `factr`, where we do not use the `nat` type in its declaration. By giving no type the default type `term` for both arguments is assumed.

```
factr: !,?
factr(0,1)::
factr(N,FN)::type(N,nat) & N1 =! N-1 & type(N1,nat) <=
                    factr(N1,FN1) & FN =! FN1*N
factr(N,0)           % The rule for N not a non-negative integer
```

Without the run-time `type(N,nat)` test, the type checker will reject this definition as not being type safe. Note the last rule returns the value 0 for any non-nat first argument. If we had used the type declaration `factr:!nat,?nat` the relation would never be called with a non-nat value as that would be a type error picked up at compile time. We then only need the `type(N1,nat)` type test as for the `fact` function definition.

If a relation has no type declaration at all, it is assumed to have been declared with `term` as each argument type with implicit mode `?`. Typically the programmer must then also use the `type` primitive with a `!` prefixed type that checks that the argument is ground as well as of the given type.

Using run-time type tests is similar to what one would do in Prolog before using a primitive to do addition, if the argument might not be a number. However, in the case of *QuLog* the use of the type check is *not optional*. It has to be in the rule else a type error will be flagged. The type checker must know that a variable holds a number value before arithmetic can be done using its value. The presence of the explicit `type(N,num)` test in the second rule below enables it to infer that `N` will have the `num` sub-type required for addition operation of this second rule.

```
add_nums_of_list_of_any_term: [term],?num
    % term is any data term so first arg can be a list of very mixed types
add_nums_of_list_of_any_term([],0)
add_nums_of_list_of_any_term([N|Rest],Total) ::
    type(N,num) <= % run-time check that N is a number
    add_nums_of_list_of_any_term(Rest,RTotal) &
    Total =! RTotal+N
add_nums_of_list_of_any_term([NonN|Rest],Total) <=
    not type(N,num) &
    add_nums_of_list_of_any_term(Rest,Total)
```

The use of the `::` test in the second rule does not affect the logical correctness of the last rule because it has the complement to the commit test of the second rule as a precondition. Of course, as in Prolog, one is tempted to drop complement tests in rules following a commit rule if these would require significant search to show they are not inferable.

An example query is

```
| ?? add_nums_of_list_of_any_term(
    [3,bill,"sally",[2,-19],0.6,empty()], Sum).
Sum = -16.6:num
```


Here the given list will have inferred type

```
[atomic||string||[int]||tree(T)]
```

as a list of any of the above types of terms. This is a sub-type of `[term]` so the call is type correct.

6.7 Type expressions as first class values

As *QuLog* has the built-in meta-type `typeExp`, we can pass in the type expression to be used for a run-time `type` test as an argument. Here is an example of a filter function that removes all but values of some given type from a list of terms.

```
filter: (typeE(T),[term]) -> [T]
filter(_TypeE,[]) -> []
filter(Type,[Trm,..Trms]) :: type(Trm,Type) ->
                                [Trm,..filter(Type,Trms)]
filter(Type,[_,..Trms] -> filter(Type,Trms)

| ?? filter(int,[3.4,hello,-6,"happy",2]).
```

```
[-6,2]:[int]
```

The role of the parameter `T` of the type expression `typeE(T)` is to allow us to inform the type checker that the value returned by the function will be a list of the type given as the first argument of a call to the function.

We give more examples of the use of run-time type checks in Section 10.1.

6.8 Generating instances of finite types

`digit` and `gender` are special program defined data types in that they each comprise a finite set of data values. In the case of `gender` it is the set of atoms `{male, female}` and in the case of `digit` it is the set `{0,1,...,9}`. For enumerated and range types, and any other types that have a finite set of values, for example unions of enumerated types, we can use another primitive `isa(E,T)`. It will generate, one at a time as successive bindings for an unbound variable `E`, all the values of the type `T`, in the order that they appear in the enumerated or range type definition of `T`. If `E` is given, it will also test that it is a value of type `T`.

For example,

```
[D | isa(D,digit)] evaluates to [0,...,9]
[G | isa(G,gender)] evaluates to [male,female].
[G | isa(G,genage)] evaluates to [male,female,0,...,110].
```

If $L=[23,8,-4,5,61]$,

```
[D | D in L & isa(D,digit)] evaluates to [8,5].
```

In the last expression, `isa(D,digit)` could be replaced by `type(D,digit)`. `isa` can be used for testing or generating instances of finite types. `type` can only be used for testing but it can test for any type, including higher order types.

6.9 Example type queries

```
| ?? type((tr(empty()),-3,empty()),!tree(int)).
yes
```

```
| ?? type(tr(empty()),-3,empty()),T).
T = tree(int)
```

```
| ?? typeC(tr(L,a,R),T,Cs).
T = tree(atom),
Cs = [L:tree(atom),R:tree(atom)]
```

```
| ?? type([1,2,male,11],T).
T = [genage]
```

We get `genage` rather than `[int||atom]` as all the given integers are in the `age` range of values, `male` is an atom of the `gender` type and `genage` is defined as union type `gender||age`.

However, a type checking query will *confirm* that `[1,2,male,11]` has type `[age||atom]`, `[int||atom]`, `[num||atom]`, `[atomic]`, `[term]`. A subtype can always be passed into a function or relation requiring a more covering type.

We can also use a **type** query to find or check the type of a function or relation

```
| ?? type(fact,T).  
T = (int->int)  
  
| ?? type(fact,T).           % fact is back-quoted making it an atom  
T = atom  
  
| ?? type(child_of,((atom,atom)<=)).  
yes  
  
| ?? type(child_of(C,mary),T).  
T = relcall           % we get the meta-type relcall as type of this call term  
  
| ?? type(child_of(C,mary),!relcall).   % Is it a ground relcall term  
no
```

7 Relations and functions over defined types

7.1 Relations defined by type validated fact sequences

Range types and enumerated types are useful for data checking. They allow us to declare a tight schema for data facts that is checked as the facts are read in from a consulted file.

Suppose that instead of defining the **person** relation in terms of **male**, **female** and **age_of** relations, it was instead an updateable belief relation. We can use the more the enumerated type **gender** and the range type **age** to have these facts data validated.

```
belief person:(atom,gender,age)  
person(bill,male,58)  
person(johnny,male,25)  
person(cheryl,female,23)  
...
```

A fact **person(bill,mle,34)** or **person(john,male,120)** will be rejected when the *QuLog* program file is consulted. **mle** is not a **gender** atom and **120** is outside the defined range for **age**.

If we query this relation the answers will be annotated with these new types.

```
| ?? 2 ? person(P,G,A) & A>=25.
P=bill:atom
G=male:gender
A=58:age
...
P=johnny:atom
G=male:gender
A=25:age
```

Suppose we also try to update a fact for this `belief` relation to record a birthday.

```
| ??  replace person(bill,male,Age) by person(bill,male,Age+1)
```

On the surface this looks fine but it will produce a type error. The reason is that the last argument of the `person` relation must be of type `age` which is a range type with maximum value 110. The type/mode checker cannot know the value of `Age` stored in the `person` fact for `bill`, and this could be the maximum age 110. If it were then `Age+1` is not a valid `age` integer.

We must use a runtime type test and enter the update command in a more verbose way.

```
| ?? person(bill,_,Age) & NewAge =! Age+1 & type(Age,age) ;
   replace person(bill,male,Age) by person(bill,male,NewAge)
```

To give you a taste of the use of action rules to define new actions, here are two action rules that will define a new action for updating the age of a person as recorded by a `person` fact.

```
update_age: atom
update_age(P) :: person(P,G,A) & NewA =! A+1 & type(Age,age)  ~>>
    replace person(P,G,_) by person(P,G,NewA)
update_age(P) ~>>>
    writeLine([P,' has no recorded age or has maximum age'])
```

`~>>>` can be read as *do*. The commit test of the first rule checks if the given `P` has a recorded age that can be incremented to a valid age. Only if that test fails will the second rule be used.

7.2 Recursion over a defined recursive data type

The label membership relation for the recursive `tree(T)` data type is the parameterised recursive relation

```

on_tree: ?T,tree(T)
on_tree(E,tr(_,E,_))
on_tree(E,tr(Left,_,_)) <= on_tree(E,Left)
on_tree(E,tr(_,_,Right)) <= on_tree(E,Right)

```

Note that the variable `E` of the first rule that appears in an `?` argument position also appears in the term `tr(_,E,_)` in the ground input argument position. This enables the mode checker to confirm that any variable given as the first argument of a call will be bound to a ground value - the root label of some sub-tree of the ground tree argument - if the call succeeds.

```

flatten: tree(T) -> [T]
flatten(empty()) -> []
flatten(tr(Left,E,Right)) -> flatten(Left)<>[E]<>flatten(Right)

```

Remember `<>` is the list concatenation function.

7.3 Complete skeleton structures containing variables

In Section 2.3 we defined the `doubleL` relation with type `([num],?[num])` and gave the example use

```
doubleL([1,3.5,-2.1,4], [2,N2,N3,..R])
```

producing the bindings `N2=7.0, N3=-4, R=[8]`.

If we wanted to insist that the second argument was always a complete list of variables or numbers, such as `[2,N2,N3,8]`, we can type the relation using

```
doubleL: ([num],[?num])
```

`[?num]` (implicitly `![?num]`) is the moded type expression for a complete list of numbers, *or variables*, of type `num`. We can have this sort of qualified grounded mode for any parameterised type. Where the type is a recursive type as here, the implicit outer `!` tells us that the recursive structure - the skeleton of the term - is fully given. If the type of its components - elements on a list or labels on a tree - is unannotated then they must all be given as ground values. The element type is also implicitly `!` moded. However, the element type may be `?` prefixed, or, if it is also a structure type, it can be have a postfix `?` mode. So, `[?tree(int)]` allows the argument to be a complete list of variables or partially given `tree` terms, all of which will become ground by a call to the relation. `[tree(int)?]` allows the call argument to be the same, and says it may remain a list of partial integer tree terms even after a successful evaluation of the call.

7.4 Meta-programming with lists of `relcall` terms

Another use of a skeleton list of non-ground terms is a list of moded type `[relcall]?`. This is a list of `relcall` terms denoting calls to program defined and primitive relations that may contain and share variables. An example is

```
[person(F,male,_), child_of(C,F), child_of(C,M),  
                                     person(M,female,_)]
```

Here is a recursive definition of an interpreter for such a list of `relcall` terms. It is what might be used by an agent implemented using *QuLog* action rules that has received such a list in a query message. We will give an example of such an agent in the companion paper on *QuLog* action rules.

```
eval: [relcall?]  
eval([])  
eval([Cond,..Conds]) <=  
    call Cond & eval(Conds)
```

`call:relcall?` is a meta-call primitive of *QuLog* very similar to Prolog's `call` primitive. The difference is that the *QuLog* `call` only accepts arguments of type `relcall`, which guarantees that each term denoted call has type correct arguments, and will only attempt to evaluate the call if the argument values satisfy the mode constraints of the relation's type declaration, which are retained after a *QuLog* program has been compiled. If it does not, the meta-call fails.

As the evaluator recurses down the above list of `relcall` terms it will first bind `F` to the name some male, then check if there is a recorded child of the male, binding `C` to the child's name, and so on. If any meta-call fails it will back-track to find an alternative solution to an earlier meta-call. It behaves like the *QuLog* query evaluator. The difference is that it will fail a meta-call that does not have its required ground arguments (if any), whereas the *QuLog* query evaluator will raise a mode error, and not even start the query evaluation.

7.4.1 Constructing `relcall` terms

There is a template term constructor operator `@` in *QuLog* that may be used to map between lists and template terms. It is similar to Prolog's `==` primitive. As a function its type is `atom@[term] -> term`. But it can also be used on the left hand side of `=!` to decompose a term into its constructor

functor and arguments, as a list. @ is like <> and ++ in being useable as a function and a pattern operator.

```
| ?? f@[2,3].

f(2,3)

| ?? F@Args =! g(1,X,"hello",X)
F = g
Args = [1,X,"hello",X]
```

This primitive can be used to construct `relcall` terms before using the `call` meta-call primitive. However, a type test will need to be used to check that the constructed template term is of type `recall`.

Here is a definition of a relation `evalListCall:[term?]` that will construct a `template` term from the list argument and evaluate it, if the result is a valid `relcall` term.

```
evalListCall: [term?]
evalListCall([Rel,..Args] <=
    type(R,atom) & Call =! Rel@Args &
    type(Call,relcall) & call Call
```

Here is an example use

```
| ?? evalListCall([`age_of,peter,A]).

A = 40
```

Note we have had to back-quote the relation name `age_of` to ensure the type checker views it as an atom rather than a relation in the list argument of `evalListCall`.

8 Higher order functions and relations

Here are the *QuLog* definitions of `mapF`, `curry` and `uncurry` for functional programming enthusiasts with example expression queries. As `->` is right associative not all the brackets are necessary in the type declarations, but using them resolves any ambiguity.

```
mapF: ((T1->T2), [T1]) -> [T2]
mapF(F, []) -> []
mapF(F, [U|L]) -> [F(U)|mapF(F,L)]
```

```
curry: ((T1,T2)->T3) -> (T1->(T2->T3))
curry(F)(A)(B) -> F(A,B)
```

```
uncurry: (T1->(T2->T3)) -> (T1,T2)->T3
uncurry(F)(A,B) -> F(A)(B)
```

```
| ?? mapF(curry(+)(2), [3,4,5]).
[5,6,7]:[digit]
```

curry(+) is a function of type

```
nat -> (nat -> nat) | int -> (int -> int) | num -> (num -> num)
```

curry(+)(2) is a single argument function that adds 2 to any supplied number. mapF takes this function and applies it to each value on the list of numbers second argument, so producing the list in which each number is 2 greater than the corresponding number on the argument list.

uncurry(curry(+)) gives back the original two argument function +. But uncurry will convert any function of type $T1 \rightarrow (T2 \rightarrow T3)$ to a function of type $(T1, T2) \rightarrow T3$, even if that function is not the result of currying.

For comparison, we also define mapR, curryR and uncurryR. Currying a relation is a little odd because it generates a function that returns a one argument relation. So its type is that of a higher order function, but its definition is a relation rule telling is how to evaluate a predication with a curryR function call value as its relation.

```
mapR: ((?T1,?T2)<=, ?[T1], ?[T2])
mapR(R, [], [])
mapR(R, [U|L], [RU|RL]) <= R(U,RU) & mapR(R,L,RL)
```

```
curryR: (T1,?T2)<= -> (T1->(T2<=))
curryR(R)(A)(B) <= R(A,B)
```

```
uncurryR: (T1->(T2<=)) -> (T1,?T2)<=
uncurryR(F2R)(A,B) <= F2R(A)(B)
```


`mapR`, as typed, can be used either to generate or test a pair of lists of the same length such that corresponding elements are in the `R` relation. Either list may be given, or neither list.

An example use of `mapR` is

```
mapR(child_of,[bill,C],[P,peter])
```

A call

```
mapR(child_of,[bill,mary],[june,peter])
```

will check that each person in the first list is a child of the corresponding person in the second.

A relation similar to `mapR` can also be defined, as it would need to be in Prolog, which takes the atom back-quoted name ``R` of a defined binary relation `R` and uses `evalListCall` to construct a type correct call to that named relation. This constructed call is the used to relate corresponding elements of the two lists. But this has a run-time failure if the constructed call is not type or mode correct for the named relation. Each use of the above higher order relation is guaranteed to evaluate only type and mode correct `R(U,RU)` calls for the mapping relation `R` as only relations of type that match the list element types can be passed into a `mapR` call. Run-time failures due to incorrect use of defined relations must be avoided whenever possible, for the serious use of inference in applications.

`curryR(child_of)` is a function that we can apply to a person `P` that returns a unary relation `curryR(child_of)(P)` that can be used to test if another person is a child of `P`. For example `curryR(child_of)(mary)` is a unary relation for checking if some person is a child of `mary`. This unary relation can be passed around as a value. For the sample facts we gave in Section 2.3, `curryR(child_of)(mary)(peter)` is inferable.

`uncurryR(curryR(child_of))` denotes the binary relation `child_of`. `uncurryR(curryR(child_of))(mary,peter)` holds.

Note that `curryR` is defined as only requiring a relation argument of type `(T1,?T2)<=`. It is used with a relation of type `(?atom,?atom)<=`, which is not an instance of the polymorphic type expression `(T1,?T2)<=`. This is allowed as the `?` mode for the first argument allows the required `!` mode of use. The relation type `(?atom,?atom)<=` is a sub-type of the instance `(atom,?atom)<=` of the `(T1,?T2)<=` type pattern. Hence `child_of` may be given as the relation argument of `curryR`. The sub-type relationship for functions and relations is defined in Section 9.

8.1 Run-time handling of higher order values

Calls to functions that return a function or relation are special in that they do not get evaluated in the way that a call to a normal function gets evaluated. A call such as `fact(3)` gets replaced by its value `6`. However, an application of a function such as `curry` or `curryR` that return a higher order value is left as a wrapper term around its argument, in the same way that application of a data constructor such as `tree` is left as a wrapper around its arguments. In both cases the value is a functor term that is passed around. So, the function *value* of `curryR(child_of)` is actually passed around as the function application term `curryR(child_of)`.

When a curried function or relation gets called, as in

```
curryR(child_of)(mary)
```

this is still not ‘evaluated’. What is then passed around, denoting the unary relation of type `atom<=`, is the above unevaluated complex term. It is only when this term denoting the one argument relation that tests if a person is a child of `mary` is called, say as in an `R(C)` condition of a rule in which `R`, typed `(atom)<=` has been assigned the ‘term’ value `curryR(child_of)(mary)`, that the resulting higher order call expression

```
curryR(child_of)(mary)(C)
```

which `R(C)` becomes, actually call some code.

First the rule `curryR(R)(A)(B) <= R(A,B)` of the `curryR` definition is applied to reduce this complex call to the call `child_of(mary,C)`. Then the relation `child_of` is queried.

8.2 Example higher order expression queries

When the expression value is a function or relation its type is returned as the value.

```
| ?? curry.
((T1,T2)->T3) -> T1->T2->T3  % RHS is T1->(T2->T3)
```

```
| ?? curry(+).
num->num->num | int->int->int | nat->nat->nat
% i.e. num->(num->num) etc
```

```

| ?? curry(+)(1).
num->num | int->int | nat->nat

| ?? curry(+)(1)(12).
13:age                                     % Again minimum inferable type for value is given

| ?? mapF(curry(+)(1), [1,2.3,3]).
[2,3.3,4]:[num]

| ?? curryR(child_of).
curryR(child_of): atom->((atom)<=)
                                     % A function from an atom to a relation over atoms

| ?? curry(<>).
curry(<>): [A] -> [B] -> [A||B]

| ?? curry(<>)([1,2]).
curry(<>)([1,2]): [A] -> [digit||A]

| ?? curry(<>)([1,2])([a,b]).
[1,2,a,b]:[digit||atom]

| ?? curry(<>)([1,2])([-3]).
[1,2,-3]:[int]

```

The first of the of the last four expression queries tells us that `curry(<>)` is a polymorphic function from a list of any type `A` to a function from a list of the same *or* different type `B`, to a list of type `A||B` elements. The second tells us that calling this function with a list of `digit` integers will return a function from a list of any type `A` to a list with elements of type the union of `digit` and `A`. This allows us to *later* append to `[1,2]` a list of atoms, as in the expression `[1,2]<>[a,b]`. We can do this because, although `[digit]` is the minimal type of `[1,2]`, and `[atom]` is the minimal type of `[a,b]`, both are sub-types `[digit||atom]`. In the last query we append a list of integers so we have a list of integers value. `[digit||int]` is simplified to the equivalent `[int]` type.

9 Sub-type relationships for functions and relations

When we have a higher order relation or function which has arguments that are themselves relations or functions, the type and mode checker makes sure that a relation or function with a type compatible with the required type is always passed in as an argument, or returned as a value. The higher order value will be type compatible if its type is in the following higher order \leq relation to the required type.

$$(t_1, \dots, t_k) \rightarrow t \leq (t'_1, \dots, t'_k) \rightarrow t' \text{ if } t_1 \geq t'_1 \ \& \dots \& \ t_k \geq t'_k \ \& \ t \leq t'$$

That is, the compatible function must be able to be given the same or a greater type for each of its arguments, and return a value of the same or lesser type.

For relation types we give the sub-type relation where every type is explicitly moded with a prefix mode.

$$(mt_1, \dots, mt_k) \leq (mt'_1, \dots, mt'_k) \text{ if } mt_1 \leq mt'_1 \ \& \dots \& \ mt_k \leq mt'_k$$

$!t \leq !t' \text{ if } t \geq t'$ *% test only use of super-type ok for test only use*
 $?t \leq !t' \text{ if } t \geq t'$ *% test/gen, for super type, ok for test only use*
 $t? \leq !t' \text{ if } t \geq t'$ *% even if a non-grounding generator*
 $?t \leq t? \text{ % grounding rel., same type, may be used if non-grounding is ok}$
 $!t(mt) \leq !t(mt') \text{ if } mt \leq mt'$ *% t a parameterised structure type*
% moded element type met of structure tester should cover all met' uses
 $?t(et) \leq ?t(et?) \text{ % may use a grounding rel for same element type if only}$
% non-grounding use for elements required, for test/gen of structure

The above tells us that the compatible relation or action procedure must have the same or a lesser *moded* type for each of its arguments. This *lesser* relation depends on both the type and the mode. The sub-type relationship between the argument types can flip from $=$ to $>$ to $<$, depending upon whether the higher order argument value may be used to both generate and test ($=$), or is only used to test ($>$), or is only used to generate ($<$), a particular argument.

Suppose a function argument is typed `(atom,num) -> num`. A function typed `(atomic,atomic)->int` has a compatible type, so could be passed as an argument. If a relation argument is typed `(!string,!int,!atom)`,

a relation moded (`?string,?num,!atomic`) can be used in the way the relation argument will be used, and may be passed as an argument value.

The fourth inequality is for parameterised structure types like list and trees, where the outer structure must be complete, but values for structure's elements, such as the elements of a list or the labels of a tree, may be given as variables or non-ground terms. The moded element types must satisfy the moded sub-type relation.

As an example, if an argument of a higher order relation is declared as being of type `([int])<=` (implicitly `(![!int])<=`) then we can pass a relation of type `([?num])<=`, (implicitly `(![?int])<=`), as the value of that argument. However, if the argument type is `([?int])<=`, we cannot pass in a relation of type `([?num])<=` in case the passed in relation is used to generate values for the list elements, which must be of type no greater than `int`.

The last rule is for relations which may be used to test, or generate, a complete structured term. It says that a relation typed `(?[?tree(num)])<=` is an allowed value for an argument of type `(?[tree(num)?])<=`, as it does not matter that the passed in relation will always generate a ground structure for the list elements - trees of numbers.

10 Moded type checks, declarative reading and sound inference

A variable *will have a ground* value if it appears in the head of the function rule, or in a `!` argument position in the head of the clause, or immediately after a call in which it appears in a prefix `?` argument position.

It *may have* a ground if it only appears in the head of a clause in prefix or postfix `?` argument terms, and only appears in earlier relation calls in postfix `?` moded argument positions.

The *initial* type of a variable is the of the type its occurrences in a rule head, or, if not in a rule head, the type of its occurrences in the first relation call in which it appears. This **initial type** then becomes its *current* type for its next call use.

The current type of variable V changes to be a sub-type t_V of its current type whenever it appears in a type test `type(V,tV)`. A warning is issued that the type test will always fail if t_V is not a sub-type of V 's current type immediately before the type test.

The type of the occurrences of a variable in a predication is the minimum of all its occurrence types. The minimum is only defined if they can

be linearly ordered. For example, if the types are `atomic` and `atom` the minimum is `atom`. For the types `int` and `atom` there is no minimum type. When a minimum type for occurrences of a variable cannot be computed, it is a type error.

Using the above rules for when a variable is ground, or may be ground, and for determining its current type, the *QuLog* type checker does the following checks as it moves left to right along a conjunction of relation calls.

1. If a variable appears in `!` argument positions in a clause head, with type t_I for these ground input occurrences, and it also occurs in other argument positions, the type t_{IO} of each non-input occurrence must be such that $t_I \leq t_{IO}$.
2. Each variable will have a ground value before it is used in a function call or a `!` argument of a relation call.
3. After it is determined a variable V has a ground value, all its later call uses have a type t_O such that $t_O \leq t_V$, where t_V is its current type just before the use.
4. After it is determined a variable V may have a ground value, all its later call uses have a type t_V , where t_V is its current type just before the use.
5. Each variable V in a prefix `?` argument position in the head of a clause will have a ground value at the end of the clause body.
6. Each `not C` condition will have ground values for all but anonymous `_` variables when evaluated
7. In each *forall*, all except the quantified variables, `_` variables and any `exists` explicitly quantified variables will have ground values when it is evaluated. In addition, the two conditions of the *forall* satisfy the normal type and mode constraints when viewed as queries.
8. In a comprehension containing `Term::exists Vars Conds`, all variables in `Conds`, except `_` variables and variables in `Term` and `Vars`, and any variables in `Term` not in `Conds`, will have ground values before the set expression is evaluated. In addition, `Conds` satisfies type and mode constraints when viewed as a query.

All these checks are made by a single pass algorithm over each clause, function rule and query. If any check fails, a type or mode use error is signalled giving the reason for the failure and the compiler terminates.

The last three checks ensure that *QuLog* rules for relations have a straightforward declarative reading as sentences of predicate logic enhanced with functions and set expressions. These rules are implicitly universally quantified with respect to all their variables except: the inner universally and existentially quantified variables of a *forall*, the local variables in the template terms of set comprehension expressions and any existentially quantified variables of its generator condition, and any `_` variables. The `_` variables are implicitly existentially quantified immediately before the condition in which they appear.

A *QuLog* query evaluation of relational and expression queries is an inference that is sound with respect to that declarative reading, providing:

- A `::` commitment test in a rule is only used to prevent unnecessary attempts to use later rules because their use would definitely fail. (In Prolog terms, each `::` is a so called ‘green’ cut).
- Equality (`=`) of higher order values is understood as identity of the expressions denoting the higher order values *not* equality of their extensions as sets of tuples of data values. It is identifier unification not denotation equality.

10.1 Need for runtime type checks to filter out sub-types

Ideally, the programmer has designed his program so that all the above checks are passed without the need for the use of run-time `type` or `ground` tests. Sometimes this is not possible.

Here is an example of the need for a type test before a call. Suppose we have

```
produce: ?atomic
use: !num
```

The conjunction of calls

```
produce(A) & use(A) ...
```

will generate a compile time type error message. However,

```
produce(A) & type(A,num) & use(A) ...
```

is acceptable, as is

```
produce(A) & type(A,int) & use(A) ...
```

The second conjunction would be used if a later call in the conjunction required `A` to have type `int`. This only makes sense if `produce` will generate some values of the sub-types `num`, `int` of `atomic`, which is a union of `nat`, `int`, `num`, `atom` and `string`. The type test then forces backtracking until such a value is generated, as this is what must be passed to the later call(s). The type test acts as a filter of the values of the super-type that might be generated by the `produce` call.

As an example of the need for a type test where the variable only may have a ground value, suppose we have

```
r: ?num
q: ?int
```

The clause

```
r(N) <= q(N)
```

must be written as

```
r(N) <= type(N,int) & q(N)
```

else there will be a type error signalled for this clause identifying the `q` call as the culprit. We need the `type` test in case `N` is given in the call to check that it is an integer. If not given, the test will just succeed and its effect is simply to satisfy the type checker that only `int` values will be passed into the `q` call. The integer binding for `N` that will then be generated by `q` satisfies the promise that a `num` value will be returned by a call to `r`.

An example related to check 4, is

```
p(...,?tree(int),...)
q(...,tree(int)?,...)
```

```
p(...,Tr,...) <= q(...,Tr,...)
```

The clause is not mode correct as `q` may not generate a ground `tree`. We must write the clause as

```
p(...,Tr,...) <= q(...,Tr,...) & ground(Tr)
```

The `ground` test will cause backtracking to retry the `q` call until it returns a ground `tree` of integer values. This only makes sense if `q` may generate several answer bindings for `Tr` some of which will be ground. If `Tr` was given as a ground tree when `p` was called, the test does no harm.

Here is a final example where only postfix `?` modes are involved.


```
p: ..,tree(num)?,..
q: ..,tree(int)?,..
```

```
... p(..,Tr,..) & type(Tr,tree(int)) & q(..,Tr,..) ...
```

is type and mode correct as only a ground or partial tree of integers can be passed through to the call to `q`. There will be backtracking to generate an alternative `Tr` value using the `p` call, should the type test fail. Without the `type` test it is not type correct as `q` could be given a tree containing non-integer values. If the `tree` argument of `q` had `!` ground input mode we would have to use

```
... p(..,Tr,..) & type(Tr,!tree(int)) & q(..,Tr,..) ...
```