- *Axiomatic specifications interpreted as assertions in operational specifications* — Preconditions and postconditions specified in the state based domain become assertions checked at the initiation and termination of an operationally specified process. Useful for mapping "black box" requirements onto detailed specifications.
- *Temporal specifications interpreted as temporal constraints in operational specifications* — Like axiomatic specifications, but checked at specific temporal instances. Useful for mapping real time constraints onto detailed specifications.

In the design flow, interactions provide information to designers whenever models are composed using the projection operators. When the model composition occurs is a matter of style, however the projection operators deliver back to the domain specific designers the implications of the interaction. Specifically, the designers working with the state based, functional model learn what impacts constraints have on their design without requiring access to the constraint model. Conversely, the constraints engineer understands the impact of the functional design on constraints issues.

The first interaction equation defines the impact of an interaction between logic and state based specifications on the logic domain:

```
M__I(f::logic,g::state-based)::set(term) is
      dom(t::M__terms(f) | dom(s::g.S | t@s))
```

This function states that every predicate defined in the `logic` facet is an invariant in the `state-based` facet. Specifically, the interaction is defined as the set of all terms from the `logic` facet asserted in every state of the `state-based` facet. The notation `t@s` is used to represent the term `t` asserted in state `s`. The set `S` is the set of all possible states as defined by the `state-based` domain. Assuming that:

```
forall(x::integer | P(x) is a term in the logic domain,
```

then:

```
forall(s::g.S | forall( x:integer | P(x)@s))
```

holds in the state based domain and the set:

```
dom(s::S | forall(x::integer | P(x))@s)
```

defines the interaction. Specifically, because `forall(x::integer | P(x))` is monotonic, it must hold in every state.

This interaction is extremely useful in defining system constraints such as power consumption. Many constraints are defined in a monotonic fashion. By composing a constraints model for a component and a functional model using a state based semantics, the interaction asserts that the constraint is true in all states.

The second interaction equation is the dual of the first:

```
M__I(g::state-based,f::logic)::set(term) is
      sel(t::M__terms(g) | forall(s::g.S | t@s));
```

The interaction set is defined as all those terms in the `state-based` facet that are true in every state. The property expressed is that if a term is true in every state, that term is invariant over all states. It is, in effect, invariant and can be stated without reference to state. The generation of the interaction set is analogous to the previous example.

Although similar in nature to its dual, this interaction discovers invariant properties in a state-based specification. Although discovering a constraint is an interesting concept, its usefulness arises typically when modeling properties such as safety and liveness conditions.

As noted, this is the simplest of the interactions defining relationships between domains using different temporal semantics. Other currently defined domains provide pair-wise relationships between monontonic, state based, finite state, infinite state, discrete time and continuous time domains. Not all of these domains are isomorphic, thus many interactions define only partial transformations of information.

Just a few interesting interactions useful for defining constraints and requirements include:

- *Monotonic constraints interpreted as moving averages* — Rather than treating monotonic specifications as absolute limits, check moving averages over time. Useful for specifying constraints whose instantaneous values are not as important as values over time.

```
M__pi(D::domain,F::facet)::facet;
```

The projection function retrieves domain aspects of a composite facet specific to a domain. To find the projection, the composition is formed using the interaction function and the projection with respect to the domain in question is extracted. If there is no model in the composition associated with D, then the projection is undefined. In a design flow, taking projections in this way corresponds to assessing the results of putting a design in its operational environment from one particular perspective.

A special case of facet composition exists when $D_j = D_k$. In this case:

$$F = (D_j, T_j \cup T_k)$$

Specifically, the resulting model is an atomic model with the domain shared by the original models and term set equal to the union of the original term sets. Further, only π_{D_j} is defined. It is not possible to retrieve the original theories from the composition.

They syntax for an interaction is defined as:

```
interaction operator(domain1,domain2::domain) is
begin interaction
  l1: M__I(facet::domain1; facet::domain2) is term-expression;
  l2: M__I(facet::domain2; facet::domain1) is term-expression;
end operator;
```

In this definition, *operator* is the facet algebra operation associated with the interaction while the *domain1* and *domain2* values specify parameters representing the two interacting domains. Typically, interactions are defined only over conjunction operations. Other operations may introduce interactions in the future and the notation is flexible to allow such definitions. Note that the definition semantics differs from traditional facet definition in that the parameter types are the domains associated with facets, not types in the traditional sense.

The domain of an interaction definition is the special `interaction` domain where the interaction operations and the projection operation are defined. By building from the `interaction` domain, designers start with a basic set of interaction definition capabilities. The *terms* associated with an interaction are a set of traditional terms that define the interaction functions. To completely specify the interaction, the definition of M__I must be included for both perumtations of the domain parameters. If this is not done, then the interaction definition will not be complete. The interaction definition applies any time two facets of the appropriate types are specified in a facet algebraic operation involving *operator*.

One of the simplest and most powerful interactions defined is the interaction between monotonic logic (mathematics) and a state-based semantics. Two interactions define the relationship defined in the *logic* domain and temporal claims defined in the `state-based` domain.. Assume that logic simply provides a mathematical domain that is monotonic, *i.e.* unchanging. In contrast the state based domain provides the concepts of current and next state. Intuitively, if both models describe the same system, each assertion made in the monotonic (logic) model must be true in every state of the temporal (state based) domain.

```
interaction and(state-based,logic::domain) is
begin interaction
  l1: M__I(f::logic,g::state-based)::set(term) is
      dom(t::M__terms(f) | dom(s::g.S | t@s));
  l2: M__I(g::state-based,f::logic)::set(term) is
      sel(t::M__terms(g) | forall(s::g.S | t@s));
end and;
```

6.2   Interactions

An *interaction* is a special purpose facet that defines situations where two domains interact. Unlike domain and traditional facet definitions, signficant syntactic sugar is added to the interaction definition syntax to simplify the special characteristics of the interaction definition.

Facets, Domains and Terms

An atomic facet, F_k, is a pair, (D_k, T_k), where D_k is the *domain* of F_k and T_k is the *term set* of F_k. The domain of a model is its semantic basis. The term set of a model is a set of terms that extend its domain to describe a more specific system. Thus, D_k provides meaning and inference capabilities to terms expressed in T_k. We say that a term, t, is a consequence of a model if $M_k \vdash_{D_k} t$, where $\vdash_{D_k}$ is inference as defined by domain D_k. Specifically, a term follows from a model if it can be inferred using the model's inference mechanism. The theory, Θ_k, of a model, M_k, is defined as the closure with respect to domain specific inference:

$$\Theta_k = \{t \mid M_k \vdash_{D_k} t\}$$

The complete calculus for models is given in the semantics guide and is taken largely from existing model theoretic research. It is sufficient for this effort to understand that each model consists of a semantic domain model and a set of terms extending that domain.

Interaction Semantics

A *composite facet* is a set of facets that are simultaneously true. Generally, a composite facet is defined using facet conjunction. Given two facets F_j and F_k, we define $F = F_j and F_k$ as:

$$F = \{(D_j, T_j \cup M\_I(F_j, F_k)), (D_k, T_k \cup M\_I(F_k, F_j))\}$$

where M_I is an *interaction function* defining the domain interaction. $M\_I(F_j, F_k)$ defines a set of terms, called the *interaction term set* or simply *interaction set*, in the semantic domain of F_j. The interaction set defines the impact of F_k on F_j using terms defined in the semantic domain D_j. Under composition, the terms of F_j are unioned with the interaction term set to augment the original model with interaction results. Again, the key to the approach is that the interaction term set is expressed in the affected domain. In a design flow, composing models in this way corresponds to putting a design in its operational environment.

The Rosetta syntax for the default domain interaction function is:

```
  M__I(F1::domain1; F2::domain2)::set(term) is empty;
```

Interaction definitions overload `M__I` for various domains. The default interaction defined by the function above is the empty term set. Specifically, if no interaction has been defined between domains, then the interaction is assumed to be null.

The projection of a composite model into a domain, π_D, is the atomic facet assocated with domain D resulting from the interaction. While composition combines models, projection pulls them back apart maintaining the effect of the interaction. Specifically:

$$\pi_D(F) = F_k \Leftrightarrow F_k \in F \land D = D_k$$

The Rosetta syntax for the projection function is:

The function `next` is defined as is `x'` for any variable `x`. Using these concepts, the time derivative, or instantaneous change associated with `x` is defined as `deriv(x@t)` or simply `deriv(x)` by viewing `x` as a function of time. An n_{th} order time derivative can be referenced by recursive application of `deriv`. The second derivative is defined `deriv(deriv(x))`, the third derivative `deriv(deriv(deriv(x)))`, and so forth.

If `x` is defined as a function over time, `x=f(t)`, the derivative is defined in the canonical fashion as:

$$\frac{dx}{dt} = \frac{df(t)}{dt} = \texttt{deriv(f)}$$

It is interesting to note that the definition `x=f(t)` expands to `x@t=f(t)@t`. Although this notation may be a bit awkward, it is consistent with the definition of the state of a Rosetta specification at a particular time `t`. It should be noted that `deriv(f)` defined in this context is identical to the derivative `deriv(f,t)` using the general derivative structure provide in the `logic` domain. The same applies for any of the derivative and integral functions provided in the time domain.

The indefinite integral with respect to `t` is defined as `antideriv(x)` and behaves similarly when `x` is a function over time. Note that the antiderivative with respect to time assumes an integration constant of zero. Making the integration constaint different is a simple matter of adding or subtracting a real value from the indenfinite integral. As with the standard indefinite integral, the following is defined:

```
antideriv(deriv(x)) == x
```

The indefinite integral of the derivative of any function over time is the original function.

The definite integral is provided as `integ(x,l,u)` and is defined as:

```
integ(x,l,u) == antideriv(x)(u) - antideriv(x)(l)
```

This is the canonical definition of the definite integral over a specified time period.

Examples

Semantics

```
%% There are still numerous problems with the following definitions.

domain continuous-time(f::facet) is
  T::real;
  __@__(x::label; t::T)::univ is M__value(x,M__items(f,t));
  next(t::T)::T;
  __'(x::label)::item is x@next(s);
  deriv(x::real)::real is deriv(x,t);
  antideriv(x::real)::real is antideriv(x,t,0);
  integ(x,u,l::real)::real is integ(x,t,u,l);
begin logic
  t1: forall(<*(t::T)::boolean is
        forall(<*(d::M__type(x,M__items(f)))::boolean is
          forall(<*(e::T)::boolean is
            abs(x@(t+e)-x') < d *>)*>)*>)
  t2: forall(<*(t::T)::boolean is
        deriv(x@t) = lim((x@next(t) - x@t) / (t - next(t)),next(t),0))
end continuous-time;

%% Original definition of next from discussions earlier in the
%%  year.
%% deriv(x@t) = x@next(t) - x@t) / (t - next(t))
```

```
    M__parse(l::M__labels(f))::universal is M__value(M__item(l,t));
    __@__(l::label; s1::T)::universal is M__value(M__item(l,s1)));
    __'(x::label)::item is x@next(t);
 begin logic
   a1: forall(<*(s::T)::boolean is
          forall(<*(l::M__labels(f))::boolean is
            M__item(l,f) = M__item(l,s)*>)
   l1: forall(<*(s1::T)::boolean is next(s1) > t *>);
   l2: forall(<*(s1::T)::boolean is -(s1 < s1) *>);
   l3: forall(<*(s1::T,s2::T)::boolean is s1 < s2 => -(s2 < s1) *>);
   l4: forall(<*(s1::T,s2::T,s3::T)::boolean is
          s1 < s2 and s2 < s3 => (s1 < s1) *>);
 end discrete-time;
```

Of particular note is the axiom asserting that `next(s) > s` from the infinite-state domain. It is a simple matter to show that specializing `next(s)` with `t+delta` satisfies this requirement. Specifically, `delta` is constrained to be a non-zero natural number. Adding `delta` to any natural number `t` results in a value that is greater than `t` by the canonical definition of natural number addition.

Summary

The `discrete-time` domain is an extension of the `infinite-state` domain where the set of possible states is the set of natural numbers and the next state function is constrained to be the addition of a discrete value to the current state. Using the `discrete-time` domain is exactly the same as using the `infinite-state` domain with the addition of discrete time values to the definition of state. Note that all `discrete-time` specifications can be expressed as `infinite-state` and `state-based` definitions with appropriate added restrictions on state ordering.

The `discrete-time` domain is the workhorse of the state-based specification domain. It is exceptionally useful when defining digital systems of all types. Using the expression notation:

```
  x' = f(x,y,z)
```

provides a mechanism for constraining the value of variables in the next state. Furthermore, the notation:

```
  x@t+(n*delta) = f(x,y,z)
```

provides a mechanism for looking several discrete time units in the future. Such mechanisms are useful when defining delays in digital circuits.

The `discrete-time` domain should not be used when no fixed timing constraints are known. In such situations, the `infinite-state` or `state-based` domains may be more appropriate and will help avoid over-specification.

6.1.7 Continuous Time

Continous time specifications provide a mechanism for defining temporal specifications using a maximally general notion of time. Unlike discrete time specifications, continuous time specifications allow reference to any specific time. Time becomes real-valued.

The `continuous-time` facet provides a type `T` representing time that is real valued. This differs from the `discrete-time` domain where time values were restricted to the natural numbers, a countably infinite set.

Examples

Semantics

The `discrete-time` domain extends the `infinite-state` domain by: (i) refining state to be of type `natural`; and (ii) refining the `next`. These definitions are provided in the definition section while the association between state and time is made in term `l1` in the following definition:

```
domain discrete-time(f::facet) is
  T::natural;
  t::T;
  delta::natural--0;
  next(t::T)::T is t+delta;
begin infinite-state
  l1: T=S and t=s;
end discrete-time;
```

Expanding fully the definition of `infinite-state` results in the following specification:

```
domain discrete-time(f::facet) is
  S::bunch(items);
  s::S;
  T::natural;
  t::T;
  delta::natural--0;
  next(s::S)::S;
  next(t::T)::T is t+delta;
  __>__(s1::S,s2::S)::boolean;
  M__parse(l::M__labels(f))::universal is M__value(M__item(l,s));
  __@__(l::label; s1::S)::universal is M__value(M__item(l,s1)));
  __'(x::label)::item is x@next(s);
begin logic
  a1: forall(<*(s::S)::boolean is
         forall(<*(l::M__labels(f))::boolean is
           M__item(l,f) = M__item(l,s)*>)
  l1: forall(<*(s1::S)::boolean is next(s1) > s *>);
  l2: forall(<*(s1::S)::boolean is -(s1 < s1) *>);
  l3: forall(<*(s1::S,s2::S)::boolean is s1 < s2 => -(s2 < s1) *>);
  l4: forall(<*(s1::S,s2::S,s3::S)::boolean is
         s1 < s2 and s2 < s3 => (s1 < s1) *>);
  l5: T=S and t=s;
end discrete-time;
```

Simplifying and removing redundant terms results in: (Note that the simplification is achieved by rewriting S and s with T and t throughout the specification.)

```
domain discrete-time(f::facet) is
  T::natural;
  t::T;
  delta::natural--0;
  next(t::T)::T is t+delta;
  __>__(s1::T,s2::T)::boolean;
```

```
    next(s::S)::S;
    __>__(s1::S,s2::S)::boolean;
    M__parse(l::M__labels(f))::universal is M__value(M__item(l,s));
    __@__(l::label; s1::S)::universal is M__value(M__item(l,s1)));
    __'(x::label)::item is x@next(s);
  begin logic
    a1: forall(<*(s::S)::boolean is
          forall(<*(l::M__labels(f))::boolean is
            M__item(l,f) = M__item(l,s)*>)
    l1: forall(<*(s1::S)::boolean is next(s1) > s *>);
    l2: forall(<*(s1::S)::boolean is -(s1 < s1) *>);
    l3: forall(<*(s1::S,s2::S)::boolean is s1 < s2 => -(s2 < s1) *>);
    l4: forall(<*(s1::S,s2::S,s3::S)::boolean is
          s1 < s2 and s2 < s3 => (s1 < s1) *>);
  end infinite-state;
```

Summary

The `infinite-state` domain is another extension of the `state-based` domain where the set of possible states is known to be infinite and ordered. Using the `infinite-state` domain is exactly the same as using the `state-based` domain with the additional restriction of assuring the existence of a state ordering and that `next` generates new states in order. Note that all `infinite-state` specifications can be expressed as `state-based` definitions with appropriate added restrictions on state ordering.

The `infinite-state` domain is useful when defining systems where states are ordered and potentially infinite numbers exist. For example, representing a discrete event simulation system is appropriate for the `infinite-state` domain.

The `infinite-state` domain should not be used when the state type is know to be finite or if no state sequencing is known. Nor should the `infinite-state` domain be used when timing models are known. As such, most specifiers will choose to use the `discrete-time` or `continuous-time` domain over the `infinite-state` domain in most modeling situations.

6.1.6 Discrete Time

The `discrete-time` domain is a special case of the `infinite-state` domain where: (i) each state has an associated time value; and (ii) time values increased by a fixed amount. Specifically, in the `discrete-time` domain, time is a natural number denoted by `t` and discrete time quanta is a non-zero natural number denoted by `delta`. The next state function is defined as `next(t)=t+delta` and following from previous domain definitions, `x@t` is the value of `x` and time `t` and `x'` is equivalent to `x@next(t)`.

Specifications are written in the discrete time domain in the same fashion as the infinite and finite state domains. The additional semantic information is the association of each state with a specific time value. Thus, the term:

```
  t1: x' = f(x)
```

constrains the value of `x` at time `t+delta` to be the value of `f(t)` in the current state. This specification style is common and reflects the general syntax and semantics of a VHDL signal assignment.

The `finite-state` domain is useful when defining systems known to have finite states. Whenever a sequential machine is the appropriate specification model, the `finite-state` domain is the appropriate specification model. Typically, the elements of the state type are specified by extension or comprehension over another bunch to assure the state type is finite. Most RTL specifications can be expressed using the `finite-state` domain if desired.

The `finite-state` domain should not be used when the state type is not known to be finite. If the additional finite state property does not add to the specification, then the `state-based` domain should be used. Of particular note is that `finite-state` specifications should not typically be used when timing information is specified as a part of function. In such circumstances, the set of possible states is almost always know to be infinite.

6.1.5   Infinite State

Like the `finite-state` domain, the `infinite-state` domain extends the `state-based` domain by restricting the state definition. Instead of making the state type finite, the `infinite-state` domain explicitly makes the state type infinite by adding the term `next(s) > s`. With this definition, two concepts are defined: (i) an ordering on states; and (ii) the next state is always greater than the current state. Strictly, the latter definition is somewhat too restrictive as it implies no loops in the state transition diagram. However, the restricted model does lend itself to systems level specification.

Examples

Semantics

The `infinite-state` domain extends the `state-based` domain by adding the assertion that the next state is always greater than the current state. This is expressed in term l1 in the following domain definition:

```
domain infinite-state(f::facet) is
  __>__(s1::S,s2::S)::boolean;
begin state-based
  l1: forall(<*(s1::S)::boolean is next(s1) > s1 *>);
  l2: forall(<*(s1::S)::boolean is -(s1 < s1) *>);
  l3: forall(<*(s1::S,s2::S)::boolean is s1 < s2 => -(s2 < s1) *>);
  l3: forall(<*(s1::S,s2::S,s3::S)::boolean is
        s1 < s2 and s2 < s3 => (s1 < s1) *>);
end infinite-state;
```

The addition of l1 to the definition implies a potentially infinite number of states as a side effect of state ordering. Specifically, a total order, ">", must be defined over any bunch of items used to define state. To assure the total order property, terms constrain the "<" to be a total order. Note that although this declaration is local to the `infinite-state` domain, the specifier is responsible for assuring the definition of the operation. Specifically, the specifier must define what "<" means for the state bunch `S`. When specifying time models in the `discrete-time` and `continuous-time` domains later, the elements of `S` will be fixed in such a way that the ordering property is assured.

Like the `finite-state` domain, because the `infinite-state` domain extends `state-based`, all definitions remain true. Expanding the `state-based` domain gives the following definition of `infinite-state`:

```
domain infinite-state(f::facet) is
  S::bunch(items);
  s::S;
```

6.1.4 Finite State

The finite-state domain provides a mechanism for defining systems whose state space is known to be finite. As finite-state extends state_based, all definitions from state_based remain valid in the new definition. Specifically, next, "@", and tick retain their original definitions. The only addition is the constraint that S must be a finite bunch. Consider the simple definition of a counter:

```
facet counter(clk::in bit; c::out real) is
S :: bunch(real);
begin finite-state
   state-space: S=0++1++2++3;
   next-state: next(S)=if S=<2 then S+1 else 0 endif;
   output: c'=next(S);
end counter;
```

In this counter, the state space is explicitly defined as the bunch containing 0 through 3. Because #S=4, the state space is clearly finite causing no inconsistency with the axiom added by the finite-state domain. Here, instead of defining the next state function in terms of properties of the next state, it is explicitly defined as modulo 4 addition on the current state. This is quite different than the previous state-based specifications where the actual value of the next state was not defined.

The final term in the counter asserts that the output in the next state is the value of the next state. Remember that here state is defined as a collection of numeric values from 0 through 4. Expanding the final term reveals the following:

```
c@next(s)=next(s)
```

The output in the next state is equal to the value of the next state.

Examples

Semantics

The domain finite-state is an extension of the state-based. Specifically, in the finite-state domain the number of defined states is finite. The domain definition extends the finite-state facet by simply adding the assertion that the size of the state type, S, is finite:

```
domain finite-state(f::facet) is
begin state-based
   l1: #S < TRUE;
end finite-state;
```

Because finite-state is an extension of state-based, all other definitions remain true.

Summary

The finite-state domain is simply an extension of the state-based domain where the set of possible states is known to be finite. Using the finite-state domain is exactly the same as using the state-based domain with the additional restriction assuring that the state bunch is finite. Note that all finite-state specifications can be expressed as state-based definitions with the added restriction on the state space size.

```
    M__parse(l::M__labels(f))::universal is M__value(M__item(l,s));
    __@__(l::label; s1::S)::universal is M__value(M__item(l,s1)));
    __'(x::label)::item is x@next(s);
  begin logic
    a1: forall(<*(s::S)::boolean is
          forall(<*(l::M__labels(f))::boolean is
            M__item(l,f) = M__item(l,s)*>)
  end state_based;
```

The Rosetta Semantics Guide defines the function `M__parse` as a function that assigns semantics to Rosetta structures. When applied to an arbitrary Rosetta structure, `M_parse` is equal to the semantic definition of that structure. In the `state-based` domain, the definition of `M_parse` is specialized to provide a mapping from labels to their associated values *in a particular state*. The function is specialized only for atomic items forming the leaves of the parse tree. In other words, items associated with lexical tokens.

`M__parse` for a given label is defined to return the value associated with that label in the current state. The meta-function `M__item(l,c)` refers to the item associated with `l` in the state (or context) `c`. In `M__parse`, `M__item` is instantiated to reference specifically the state `s` defined in the `state_based` domain to be the current state. `M_value` is used to access the value of the item. Thus, any label appearing in a `state_based` facet refers to the value of the item named by the label in the current state.

The infix function `@` is used to reference a label's value in an arbitrary state. The notation `x@s5` refers to the value of `x` in a state referred to by `s5`. The definition of `@` is nearly identical to `M_parse` using `M__item` and `M__value` to obtain the value of a label in an arbitrary state. As anticipated, the theorem that `x==x@s` easily proven by instantiating the state variable in `@` with the current state `s`. Of special note is that the first argument to `@` is a label. Because the label rather than the value is desired in this situation, the `M__parse` function is not applied to the first argument of `@`.

In the `state_based` domain, the dominant specification methodology is axiomatic specification. Thus, the primary use of `@` is to refer to label values in the next state. Specifically, statements such as `R(x,x@next(s))` are used to constraint the value of `x` in the next state based on the value of `x` in the current state. In axiomatic specification, the standard notation `x'` is used as a shorthand for `x@next(s)`. Thus, Rosetta provides a shorthand notation for any symbol, `x`, in the next state as `x'`. The previous example specification can thus be rewritten in a more compact notation as `R(x,x')`.

Summary

The `state_based` domain provides a mechanism for specifying how a component or system changes state. It provides a basic type, `S`, for states and a state variable, `s::S`, that represents the current state. In addition, the STATE_BASED domain provides a definition for `next`, a function that generates a new state from a given state. Thus, if `s` is the current state, then `next(s)` is the next state.

To refer to the value of a variable in a state, the "`@`" operation dereferences a label in a bunch of items. Specifically, given an item bunch `c` and a label `l`, the notation `l@c` refers to the value of the item associated with `l` in the context `c`. As states are defined as bunches of items, the "`@`" operation is easily used to obtain values of items in a state.

Because relationships between the current and next state are frequently the objective of state based specification, the `state_based` domain provides a shorthand for referencing items in the next state. Specifically, the notation `x'` is equivalent to `x@next(s)` providing a convenient shorthand for referencing next state values.

The `state-based` domain should be used whenever a system level description of component or system is needed. At early design stages when working at high levels of abstraction, the `state-based` domain provides a mechanism for describing state transformations without unnecessary details.

The `state-based` domain should not be used when details such as timing are involved in the specification. Furthermore, the `state-based` domain provides no automatic mechanism for composing component states when developing structural models.

```
x == x@s
```

and

```
x' == x@next(s)
```

The previously defined counter specification is equivalent to the following expansion:

```
facet counter(v::out natural) is
  n::natural;
begin state_based
  next: if n@s < 7 then n@next(s)=(n@s)+1 else n@next(s)=0 endif;
  output: v@next(s) = n@s;
end counter;
```

where the tick notation is replaced by it's definition and references to labels in the current state are expanded to explicitly reference the state. Readers curious about the actual definition of the state_based domain should refer to Section 6.1.3 defining the semantics of state_based. Readers needing only to understand use of the state_based domain may safely skip Section 6.1.3.

Examples

```
%% Steal the examples from the tutorial and add a few more.
```

Semantics

This and subsequent semantics sections may be skipped by readers who do not wish to see the internals of a domain definition.

The state_based domain provides a basic definition of state and change. Two basic mechanisms are provided: (i) a definition of state; and (ii) a definition of what next state means. The definition of state provides a state type that can be referenced in definitions. In the state_based domain, relatively few restrictions are placed on the state definition. The next state function provides the concept of change by sequencing states. Like the state type, the state_based domain places few restrictions on the next state function.

The state type, S, is defined as an uninterpreted type representing possible states. The function M_item defined in the Semantics Guide is a function mapping a label and state to the item associated with that label. Within the body of the state_based domain definition, M_item(l,s) is defined to refer to the same object in every state as is defined in the facet. In other words, each state is an extension of the facet definition. Information can be added, but base definitions cannot be change. A specific instance, s::S, is defined as the *current state*. Effectively, s provides a name that can be referenced in definitions rather than quantifying over all states.

The next state function, next, is defined as a function that maps one state to another. No other constraints are defined for the next function. In other domains related to state_based, next will be restricted and specialized to model varying definitions of time and state change. In the basic domain, only the existence of a current and next state are defined.

```
domain state_based(f::facet) is
  S::bunch(items);
  s::S;
  next(s::S)::S;
```

Summary

6.1.3 State Based

The domain `state_based` is used to define systems that change state. The `state_based` domain extends the mathematical capabilities provided by the `logic` domain to include the concept of state and change. Recall that in the `logic` domain, the values associated with items could not change. Doing so created inconsistencies with the original definitions. The `state_based` domain provides the basis for modeling the concepts of state and change by defining: (i) the state of a facet; (ii) the current state; and (iii) a next state function that derives the next state from the current state.

Consider the following trivial definition of a counter that counts from 0 to 7 and repeats:

```
facet counter(v::out natural) is
  n::natural;
begin state_based
  next: if n < 7 then n'=n+1 else n'=0 endif;
  output: v' = n;
end counter;
```

This definition uses a natural number, `n`, to maintain the current counter value and uses two terms to define the next state and output respectively. The first term, labeled `next`, defines the next state given the current state:

```
next: if n < 7 then n'=n+1 else n'=0 endif;
```

In this term, `n` refers to the value of `n` right now in the current state. The notation `n'` refers to the value of `n` in the next state after the component or system represented by the facet has completed its computation. Understanding this convention, the term can now be interpreted as a conditional statement stating 'if `n` is less than seven in the current state, then `n` in the next state is `n+1`, else `n` in the next state is 0.' This is precisely how a counter calculates its next value.

Similarly, the second term defines the next output:

```
output: v' = n;
```

Using the same interpretation mechanism, the next value of `v` will be the current value of `n`. This is somewhat interesting as the output lags the current state by one value. If such behavior is not desired, then this term can be modified to state `v'=n'`.

It is exceptionally important to recognize that the following term similar to a C-like programming statement is not correct:

```
next: if n < 7 then n=n+1 else n=0 endif;
```

Remember that terms state things that are true. These are not executed and there is no notion of assignment. Although legal in C where = is an assignment operator, in Rosetta this statement asserts that if $n < 7$, then $n = n + 1$ is also true. Looking at = as equality rather than assignment makes the second statement inconsistent as there is no natural number that is equal to itself plus one. The key to using `state_based` domains is recognizing that no label tick indicates the current state and label tick indicates the next state.

The `state_based` tick notation is defined based on the more fundamental `state_based` domain definitions of current state and the next state function. In reality, the notation `x'` is shorthand for the notation `x@next(s)` where: (i) `@` refers to the value of a label in a state; (ii) `next` defines the state following a given state, and `s` is the current state. Specifically:

the parameter `f` in `state_based` refers to the including facet `register`. Thus, the domain definition can generically reference elements of the including facet in its definition. For example, it is possible to reference `M_labels(f)` or `M_items(f)` to reference the labels and items defined in `f` respectively.

As with traditional facet definition, a domain definition extends the theory provided by its referenced domain. It is therefore possible to define a lattice of domains that inherit and specialize each other. Figure 6.1 shows one such specification lattice including pre-defined domain definitions. Solid arrows represent extension between domains.


Figure 6.1: Lattice of pre-defined specification domains.

The following sections provide basic definitions and usage examples for each pre-defined domain.

6.1.1 Null

The `null` domain refers to the empty domain. It is included to provide a basis for defining domains that inherit nothing from other domains. The `logic` domain that provides basic mathematics will use `null` as a domain to indicate that it is self contained. There is no constructive definition of `null` because it has no domain definition.

6.1.2 Logic

```
domain logic(f::facet) is
begin null
end logic;
```

Chapter 6

Domains and Interactions

Domains and interactions are special facets that define domain theories and interactions between domain theories respectively.

6.1 Domains

```
%% Add reference to the fact that facets extend domains.  This
%% defines what domain inclusion means.
```

A *domain* is a special purpose facet that defines a domain theory for facets. The syntax for a domain is defined as:

```
domain <name>(f::facet) is
  <declarations>;
begin <domain>
  <terms>
end <name>;
```

where ¡name¿ is the label naming the domain, ¡declarations¿ are items defined in the facet, ¡domain¿ is the domain facet extended by the new definition, and ¡terms¿ define the new domain facet. All domains are parameterized over a single facet variable that represents a place holder for the facet including the domain theory. Given a domain called state_based used in the following facet:

```
facet register(i::in bitvector; o::out bitvector; s0::in bit, s1::in bit) is
  state::bitvector;
begin state_based
  l1: if s0=0 then
          if s1=0 then state'=state
              else state'=lshr(state) endif
          else if s1=0 then state'= lshl(state)
              else state'=i endif
      endif;
  l2: o'=state';
end find;
```

5.5.2 Parameter Ordering

The pragmatics of using parameter list union insist that some ordering be placed on the results. Typically, specifiers use the order of parameters in parameter list to associated actual parameters with formal parameters. Rosetta provides two mechanisms for handling this situation. The first is for the user to define an ordering and the second is to use explicit parameter assignment.

To explicitly define parameter ordering in the facet resulting from a conjunction the user specifies parameters in the facet declaration. For our `and_gate` gate example previously, the following definition specifies an ordering for resulting parameters:

```
and_gate(z,y,x::null)::facet is and_discrete and and_cont;
```

In this definition, the parameter ordering in the definition of `and_gate` defines the parameter ordering. The `null` type is used to specify the parameter types as for any type `T`, `T and null == T`. Thus, the parameter definitions add ordering information, but add no additional type information to the definition.

Users may also allow Rosetta to order the types for them.

```
%% Ordering definition here
```

Examples:

Summary:

5.5.1 Type Composition

The more interesting case occurs when a parameter is shared betweeen facets, but the declarations specify different types. Consider the following two facet declarations:

```
facet F1(x::R, y::S, t::T) is        facet F2(w::Q,x::P) is
...                                  ...
end F1;                              end F2l;
```

In this case, the parameter list of F1 and F2 is (x::R, x::P, y::S, t::T, w::Q). Note that two declarations of x exist in the parameter list definition. Recall that parameter declarations are simply terms appearing in the parameter list. Specifically, a variable or parameter declaration is shorthand for:

```
x:e::R
```

Viewing the parameter definition in this way allows application of label distribution laws. This application yeids the parameter list (x::R and P, y::S, t::T, w::Q). Note that in this parameter list x is of type R and P implying that x can be viewed both as type R and type P.

When conjuncting and disjunction facets, care must be taken to assure that parameters having the same name represent the same physical quantity. The type declaration R and P results in a type that, in principle, behaves like the result of facet conjunction. Specifically, an item of this type is simultaneously viewed as being of both types. It is also important to understand that type composition is not type union. Specifically R and P is not equal to R ++ P. In the latter case, elements of R ++ P can take values from either R or P.

An excellent example of type composition occurs when looking at a circuit component such as a simple gate from multiple perspectives. Consider a simple and gate viewed in both the analog and digital domains:

```
facet and_discrete(x,y::in bit;      facet and_cont(x,y::in real;
                z::out bit) is                   z::out real) is
begin state_based                    begin continuous
  l1: z' = x*y;                        <and gate definition here>
end and_discrete;                    end and_cont;
```

The definition of a completely modeled and_gate gate becomes:

```
and_gate :: facet is and_discrete and and_cont;
```

The parameter list resulting from this definition is:

```
(x,y::in bit and real, z::out bit and real)
```

Thus, each parameter can be viewed as either a real or discrete value.

```
%% Need discussion of parameter interaction here.  Defer semantics
%% to the semantics guide, but some discussion must occur.
```

5.3   Facet Implication

Facet implication, $F_1 \Rightarrow F_2$, states that properties specified by term T_1 must imply properties specified by term T_2. Note that $F_1 \Rightarrow F_2 \equiv \neg F_1 \vee F_2$. The most obvious use of refinement is showing that one facet "implements" the properties of another. Specifically, if $F_1 \Rightarrow F_2$, then the theory of F_2 is a subset of the theory of F_1.

Example 21 (Implementation) *Given the requirements defined for sort in* `sort_req`, *any legal implementation of a sorting algorithm must implement these properties. We say that* `sort_req` *can be* refined *into* `bubble_sort` *and state this as:*

```
sort_ref::facet is bubble_sort => sort_req;
```

Additional constraints may be added by conjuncting facets in the consequent of the implication. The following is an example of adding a low power constraint to the specification:

```
constrained_sort_ref::facet is bubble_sort => low_power and sort_req;
```

This is an interesting result because it insists that `bubble_sort` *be a low power solution to the sorting problem. As an aside, the definition of conjunction requires that* $F_1 and F_2 => F_1$.

Summary:

5.4   Facet Equivalence

Facet equivalence, $F_1 \Leftrightarrow F_2$, states that properties specified by terms T_1 and T_2 in domains D_1 and D_2 must be equivalent. The formaly definition of equivalence can be expressed in terms of implication. Formally:

$$F_1 \Leftrightarrow F_2 = F_1 \Rightarrow F_2 \wedge F_2 \Rightarrow F_1$$

Summary:

5.5   Parameter List Union

Throughout the defintion of the facet algebra, reference is made to the union of parameter lists. Specifically, when facets are combined the parameter list of the new facet is defined as $I_1 + + I_2$. Viewed as bunches, this definition is literally true where all parameters from both facets become parameters in the new facet.

Given the facet declarations:

```
facet F1(x::R, y::S, t::T) is        facet F2(w::Q,x::R) is
  ...                                   ...
end F1;                               end F2;
```

The parameter list of `F1` and `F2` is `(x::R, y::S, t::T, w::Q)`. Note that the declaration of parameter x is shared in both facet declarations. Bunch union implies that a single x appears in the result of parameter list union.

Example 20 (Component Version) *Another excellent example of disjunction use is representing a family of components. Consider the following definitions using sort facets defined previously:*

```
multisort::facet is sort_req and (bubble_sort or quicksort);
```

The new facet multisort describes a component that must sort, but may do so using either a bubble sort or quicksort algorithm. [1]

A more interesting definition configures a component to represent both low and high power configurations of a device:

```
facet low_power is                      facet power is
  power::real;                            power::real;
begin constraints                       begin constraints
  power =< 1mW;                           power =< 5mW;
end low_power;                          end power;

facet tx_req(d::in data;               low_power_tx::facet is
             s::out signal) is            tx_req and low_power;
begin continuous
  <transmitter definition here>        high_power_tx::facet is
end tx_req;                               tx_req and power;
```

In this example one specification for a transmitter function is provided along with two definitions of low and high power versions. Facet conjunction is used to combine power constraints with functional transmitter properties.

Consider the following specification:

```
tx(select::boolean)::facet = if select then
                                  low_power_tx
                                else high_power_tx
                             endif;
```

Here a generic parameter is introduced into the definition to select one version over another. When `select` is instantiated, then `tx` resolves to the appropriate model. A more interesting case occurs when `select` is skolemized to an arbitrary boolean constant a:

```
tx(a) ==  if a then low_power_tx else high_power_tx endif;
```

Whenever facet `tx` is used in this manner, both specifications must be considered. Effectively, `tx` defines two transmitter models. When instantiated in a structural facet, both models must be considered in the analysis activity. It must be noted that the parameter `select` is a boolean valued parameter and not a facet. It is tempting to attempt a definition of `if-then-else` that uses facets as all its parameters. However, such a definition has been shown to have little utility.

```
%% These must be dealt with separately because they do not result in
%% facets.  Although implication should if it's defined in terms of
%% disjunction.
```

Summary:

---

[1] Assume the facet quicksort has been defined in the canonical fashion.

Alternatively, the following definition can be used to define sort*:*

```
    sort::facet;
  begin domain
    l1: sort = sort_req and sort_const;
    ...
  end domain;
```

Another alternative is using relabeling to define a single sort component in a structural Rosetta description:

```
  begin domain
    sort: sort_req and sort_const;
    ...
  end domain;
```

In each case, the resulting sort *definition is the conjunction of the* sort_req *and* sort_const *definitions.*

Summary:

5.2 Facet Disjunction

Facet disjunction, $F_1 \vee F_2$, states that properties specified by either terms T_1 in domain D_1 or T_2 in domain D_2 must be exhibited by the resulting facet. Like conunction, the interface of the resulting facet is $I_1 + +I_2$, the union of the facet interfaces.

The most obvious use of facet disjunction is the definition of cases. Two situations are of particular interest: (i) using predicatative semantics to define component behavior; and (ii) defining families of components.

Example 19 (Case Specification) *Given a container C defined as a collection of key (K), element (E) pairs, naive requirements for a simple search algorithm are defined as:*

```
  facet search(c::in C; k::in K; o::out E) is
  begin state_based
    member((k,o'),c);
  end search;
```

Clearly, this specification will be inconsistent if there is no element in c corresponding to k. Thus, it is traditional to break the requirements into two cases: (i) the element is present and is returned; and (ii) the element is not present. Such a situation is modeled by the following two specifications:

```
facet searchOK(c::in C; k::in K;            facet searchErr(c::in C; k::in K;
              o::out E) is                                 o::out E) is
begin state_based                           begin state_based
  l1: exists(x::E | member((k,x), c));        l1: -exists(x::E | member((k,x), c));
  l2: member((k,o'),c);                     end searchErr;
end searchOK;
```

Facet search *is now defined:*

```
  search::facet is searchOK or searchErr;
```

Chapter 5

The Facet Algebra

```
%% Needs a good introductory paragraph

%% Still need to define syntax for the facet composition operators.
%% Also need some reference to the formal semantics.  Put a small
%% section in on the theory calculus from the interactions white
%% paper.
```

The following sections describe several prototypical uses of facet composition. Please note that domains use in these examples are defined in Chapter 6. In the following definitions, assume that all F_n are facets where T_n, D_n and I_n are the term set, domain and interface associated with F_n respectively.

5.1 Facet Conjunction

Facet conjunction, $F_1 \wedge F_2$, states that properties specified by terms T_1 and T_2 must by exhibited by the composition and must be mutually consistent. Further, the interface is $I_1 + + I_2$ implying that all symbols in the parameter lists of F_1 and F_2 are also visible in the parameter list of the composition.

The most obvious use of facet conjunction is to form descriptions through composition. Of particular interest is specifying components using heterogeneous models where terms do not share common semantics. A complete description might be formed by defining requirements, implementation, and constraint facets independently. The composition forms the complete component description where all models apply simultaneously.

Example 18 (Requirements and Constraints) *Consider the following facets describing a sorting component:*

```
facet sort_req(i::in T; o::out T) is      facet sort_const is
begin state_based                           power::real;
  l2: permutation(o',i);                  begin constraints
  l1: ordered(o');                          p1: power =< 5mW;
end sort_req;                             end sort_const;
```

A sorting component can now be defined to satisfy both facets:

```
sort::facet is sort_req and sort_const;
```

Example 16 (Low Power Transmitter) *Define a new facet for transmitters and receivers that constrains power consumption:*

```
facet low_power is
  power::real;
begin constraints
  p0: power =< 10MW;
end low_power;


%% Cindy, look at this syntax...
```

One can now define a low power transmitter as:

```
tx_low_power(data::in T; output::out T)::facet is tx and low_power;
```

In this definition, a new facet called ***tx_low_power*** *is defined that is the composition of the transmitter functional facet and the low power constraints.*

```
%% The definition below should be contained in a facet definition as
%% it's really a function.
```

Example 17 (Transmitter Configuration) *Define a new facet for high power transmission:*

```
facet high_power is
  power::real;
begin constraints
  l1: power =< 100Mw;
end high_power;
```

Now define a configurable device that represents either the high or low power version:

```
tx_power_select(select::boolean)::facet is
   tx(data,output) and
   if select
      then low_power
      else high_power
   endif;
```

When the ***select*** *parameter is true, then the* ***tx*** *facet is composed with the* ***low_power*** *constraints facet. Otherwise, the* ***tx*** *facet is composed with the high power constraint facet.*

Recall that in the presense of an `export` statement, only specified labels are visible outside the facet. Here, a facet is used rather than a facet to allow specification of the `encode_decode` axiom that states the inverse relationship between the encode and decode functions. The `use` clause makes the functions visible and available to the transmitter and receiver specifications. The axiom is not visible, but does remain present in the definition.

Example 12 (Transmit/Receive Pair) *The following defines the simplest possible communications channel transmitting and receiving encoded, baseband signals:*

```
facet tx_rx_pair (data_in::in T; data_out::out T) is
    channel::T;
begin logic
    txb: tx(data_in,channel);
    rxb: rx(data_out,channel);
end tx_rx_pair;
```

The resulting component represents a perfect transmitter receiver pair where input data is perfectly transmitted to an output data stream.

Example 13 (Transceiver) *The following defines a simple transceiver combining the transmitter and receiver functions into a single component:*

```
facet transceiver (data_in::in T; data_out::out T;
                    out_chan::out T; in_chan::in T) is
being structural
    txb: tx(data_in, out_chan);
    rxb: tx(data_out, in_chan);
end transceiver;
```

Note that in this specification, the transmitter and receiver do not interact. They simply operate in parallel on indepednent data streams.

Example 14 (Transceiver Pair) *Now consider a tranceiver pair constructed from two transceivers:*

```
facet trx_pair (data_in1, data_in2::in T;
                data_out1, data_out2::out T) is
begin logic
    chan1,chan2::T;
    trx1: transceiver(data_in1,data_out1,chan1,chan2);
    trx2: transceiver(data_in2,data_out2,chan2,chan1);
end trx_pair;
```

Example 15 (Transceiver Pair - Common Channel) *An adaptation of a tranceiver pair is one where transmission from both devices occurs on the same channel. Here, only one channel parameter is defined:*

```
facet trx_pair (data_in1, data_in2::in T;
                data_out1, data_out2::out T) is
    chan::T;
begin logic
    trx1: transceiver(data_in1,data_out1,chan,chan);
    trx2: transceiver(data_in2,data_out2,chan,chan);
end trx_pair;
```

Here, conjunction distributes across the structural definition. Proper label selection allowed power constraints to be associated with each component. The result can be viewed as either the conjunction of a power and functional model or the composition of two component models both having constraint and functional models.

```
%% Working Here %%
```

Example 11 (Structural Example) *Consider the following facets:*

```
facet sort(x::in array(T); y::out array(T)) is
begin state_based
  l1: permutation(x,y');
  l2: ordered(y');
end sort;


facet binsearch(k::in keytype; x::in array(T); y:out T) is
begin state_based
  l1: ordered(y);
  l2: member(k, dom(x)) => member(y', dom(x)) AND key(y')=k;
end binsearch;


facet find_structure(k::in keytype; x::in array(T); y:out T) is
  buff::array(T);
begin logic
  b1: sort(x,buff);
  b2: binsearch(k,buff,t);
end find_structure;
```

The sort *and* binsearch *facets define requirements for sorting and binary search components. The* find-structure *facet defines a* find *architecture by connecting the two components. The state variable* buff *is shared by the binary search and sorting components and facilitates sharing information. Note that* new *does not generate a new copy of* buff *because new is called on both* sort *and* binsearch *before parameters are instantiated. Thus, the same object* buff *is references in the terms of both components and constrained by those terms.*

The following collection of examples are designed to demonstrate several configurations of a simple transceiver system. The following represent simple specifications of a transmitter and receiver used throughout the examples:

```
use signal_processing_requirements(T);    use signal_processing_requirements(T);
facet tx (data::in T; output::out T) is    facet rx (data::out T; input::in T) is
begin state_based                          begin state_based
  l1: output'=encode(data);                  l1: data'=decode(input);
end tx;                                    end rx;
```

These specifications assume the following domain facet for signal processing:

```
facet signal_processing_requirements(T::subtype(univ)) is
  encode(t::T)::T;
  decode(t::T)::T;
  export encode,decode;
begin logic
  encode_decode: forall(t::T | decode(encode(t))=t);
end signal_processing_requirements;
```

the renaming is that the `delay` physical variable associated with the adder definition is duplicated. *I.e.* the values `b0.delay` and `b1.delay` are available for reference and represent distinct objects. Without renaming using labels, both `one_bit_adder` instances would refer to the same physical variable, `one_bit_adder.delay`. This is not appropriate as the adders should be distinct. The same result can be achieved using parameter for delay. In large specifications including parameters for physical variables representing constraint specifications becomes cumbersome. Further, delay is not a parameter but a characteristic of the component.

After including the two adder instances, the value of `delay` in the `two_bit_adder` specification is constrained to be equivalent to the sum of the `one_bit_adder` delays. In this way, it is possible to specify composition of non-behavioral characteristics across architectures.

Logical operators are defined to distribute across structure components. Assume the following facets defining power constraints on a one bit adder and an architecture defining constraints on a two bit adder composed of two one bit adders:

```
facet one_bit_adder_const is          facet two_bit_adder_const is
   power::posreal;                        power::posreal;
begin constraints                      begin constraints
   p0: power <= 5mW;                      b0: one_bit_adder_const;
end one_bit_adder_const;                  b1: one_bit_adder_const;
                                          p0: power = b0.power + b1.power;
                                       end two_bit_adder_const;
```

The facet conjunction `two_bit_adder = two_bit_adder and two_bit_adder_const` is equivalent to:

```
facet two_bit_adder(x0,x1,y0,y1::in bit; z0,z1,c::out bit) is
  delay::real;
  power::posreal;
  cx::bit;
  export delay,power;
begin logic
  b0: one_bit_adder(x0,y0,0,z0,cx);
  b0: one_bit_adder_const;
  b1: one_bit_adder(x1,y1,cx,z1,c);
  b1: one_bit_adder_const;
  d0: delay = b0.delay+b1.delay;
  p0: power = b0.power+b1.power;
end two_bit_adder;
```

This definition results from the definition of facet conjunction. The term set is simply the set of all defined terms in the two facets.

This definition results from the distributivity of labeling. The same result holds for disjunction, implication and logical equivalence. Application of label distribution results in:

```
facet two_bit_adder(x0,x1,y0,y1::in bit; z0,z1,c::out bit) is
  delay::real;
  power::posreal;
  cx::bit;
  export delay,power;
begin logic
  b0: one_bit_adder(x0,y0,0,z0,cx) and one_bit_adder_const;
  b1: one_bit_adder(x1,y1,cx,z1,c) and one_bit_adder_const;
  d0: delay = b0.delay+b1.delay;
  p0: power = b0.power+b1.power;
end two_bit_adder;
```

%% Why does 'power' appear in facet find below? It is not used there.

```
use find_primitives(T,keytype);
facet find(k::in keytype; i::in array(T); o::out T) is
  power::real;
begin state_based
  postcond1: key(o') = k;
  postcond2: elem(o',i);
facet find;
```

Here the `use` clause makes all labels defined in `find_primitives` visible in the current scope. When using this approach, the "." notation is no longer necessary as the functions `key` and `elem` are now visible. In most situations, this is the desired mechanism for packaging and using definitions. The special `package` definition provides a facet construct specifically for this purpose. Please see Chapter 2 and Chapter 5 for more details on the semantics and use of packages.

4.5.2 Structural Definition

System structure is defined using facet inclusion and labeling in the sam manner as defined previously. Facets representing components are included and interconnected by instantiating parameters with common objects. Labeling provides name spaced control and supports defining multiple instances of the same component. Consider the following specification of a two bit adder using two one bit adders:

```
facet one_bit_adder(x,y,cin::in bit; z,cout::out bit) is
  delay::real;
  export delay;
begin state_based
  l1: z'= x xor y xor cin;
  l2: cout' = x and y;
end one_bit_adder;

facet two_bit_adder(x0,x1,y0,y2::in bit; z0,z1,c::out bit) is
  delay::real;
  cx::bit;
  export delay;
begin logic
  b0: one_bit_adder(x0,y0,0,z0,cx);
  b1: one_bit_adder(x1,y1,cx,z1,c);
  l1: delay = b0.delay+b1.delay;
end two_bit_adder;
```

Facet interconnection is achieved by sharing symbols between component instances. When a facet is included in the structural facet, formal parameters are instantiated with objects. When objects are shared in the parameter list of components in a structural facet, those components share the object. Thus, information associated with the object are shared between components. The *two_bit_adder* specification includes two copies of `one_bit_adder`. Parameters of the two adders are instantiated with parameters from `two_bit_adder` to associated signals with those at the interface. The internal variable `cx` is used to share the carry out value from the least significant bit adder with the carry in value from the most significant bit adder.

When the two `one_bit_adder` instances are included in the `two_bit_adder` definition, they are labeled with b0 and b1. The result is that the first `one_bit_adder` is renamed b0 and the second b1. The implication of

4.5 Relabeling and Inclusion

The ability to rename object in conjunction with label distribution laws allows definition of: (i) facet inclusion and instances; (ii) system structures; and (iii) type combination. Facet inclusion supports use of facets as units of specification modularity. If renamed when included, the new facet represents a renamed instance of the original. With inclusion, describing structural definitions becomes possible. Finally, using variable labels allows definition of type combination and interface union.

4.5.1 Facet Instances and Inclusion

Facet inclusion allows compositional definition in a manner similar to packages or modules in programming languages and theory inclusion in formal specification language. Whenever a facet label is referenced in a term, that facet is included in the facet being defined. Consider the following extended `find` specification:

```
facet find_primitives(T,K::subtype(univ)) is
  key(t::T)::K;
  elem(t::T,a::array(T))::boolean;
  export all;
begin logic
end find_primitives;
```

`%% Why does 'power' appear in facet find below? It is not used there.`

```
facet find(k::in keytype; i::in array(T); o::out T) is
  power::real;
begin state_based
  findpkg: find_primitives(T,keytype);
  postcond1: findpkg.key(o') = k;
  postcond2: findpkg.elem(o',i);
facet find;
```

In previous `find` specifications, definitions for `key` and `elem` remain unspecified. In this example, the facet `find_primitives` defines those operations. The `find` facet includes a copy of `find_primitives` in the term labled `findpkg`. Semantically, this term includes a copy of `find_primitives` and relables the facet with `findpkg`.

In the resulting definition, elements of the newly renamed facet are accessed using `findpkg` as their associated facet name. Specifically, the `elem` and `key` functions defined in `find_primitives` are referenced using the `findpkg.elem` and `findpkg.key` notations respectively. `findpkg` is said to be an instance of the original facet because each newly named copy is distinct from the original. This includes physical variables as well as terms. Thus, two renamed copies of the same facet will not inadvertently interact. This is exceptionally important when defining structural definitions where many instances of the same component may be required.

Alternatively, a facet or package may be referenced in a `use` clause to make their definitions visible in the current scope. Consider the following definition:

```
facet find_primitives(T,K::subtype(univ)) is
begin requirements
  key(t::T)::K;
  elem(t::T,a::array(T))::boolean;
begin logic
  ...
end find_primitives;
```

4.4.2 Distributing Declarations and Terms

Label distribution over semicolons occurs when two declarations or terms share the same label. Specifically, an example in the case of declarations:

```
x::integer;
x::character;
```

and in the case of terms:

```
b1: and_gate(x,y,z);
b1: constraint(p);
```

In both cases, the terms or declarations share the same label. In such circumstances, the semantics of distribution is the conjunction of the definitions. In the case of declarations:

```
v::S is c; v::T is d;
```

is equivalent to:

```
   v::S and T;
begin <domain>
   t: v=c and v=d;
```

The semantics of the declaration are such that v is the coproduct of S and T.[1] Specifically, any value associated with v has both the properties of S and the properties of T. This is not type intersection, but product in the category theoretic sense. The definition does not say that v is in the intersection of the original types, but says that it has a projection into both types.

The semantics of distribution over term declarations is similar. The definition:

```
b1: and_gate(x,y,z);
b1: constraint(p);
```

is equivalent to:

```
b1: and_gate(x,y,z) and constraint(p);
```

If the two conjuncts are boolean expressions, the definition of conjunction applies. If the two conjuncts are facets, then the new facet b1 has the properties of both an and_gate and constraint simultaneously.

Summary: Label distribution is defined across all boolean operators as well as semicolons as used in declarations and term definitions. In all cases for boolean operations, identical labels distribute across operations. In all cases for semicolons, declarations and terms sharing labels can be combined into a single declaration or term resulting from the product (conjunction) of their definitions.

---

[1] See Chapter 5 for details of conjunction usage.

4.3.4 Explicit Exporting

Visibility of labels is controlled using the `export` clause that appears in the declaration part of a facet. The convention for label exporting is that any label listed in the `export` clause is visible outside the enclosing facet using the standard *facet.label* notation. Labels not listed in the export clause are not visible and cannot be referenced. All labels within a facet can be exported using the special shorthand `export all` notation. If the export clause is omitted, then no labels from the facet are visible.

Summary: All Rosetta items are labeled and can be referenced in a specification by their associated label. Three major labeling operations are the definition of facets, the declaration of variables and constants, and the definition of terms within a facet. Any label may be referenced outside its enclosing facet using the canonical notation *facet.label* where *facet* is the containing facet's name and *label* is the label being access. Controlling access is achieved using the `export` clause. If an `export` clause is present, all listed labels are visible and all unlisted labels are not. If `all` appears in the `export` clause, then all labels are exported. If no export clause is present, then no labels are visible outside the facet.

4.4 Label Distribution Laws

Given two labeled Rosetta items, distributive properties of labels over logical operations can be defined as follows:

Equivalence	Name
`l1:`j `and` i `= l1:`j `and l1:`i	and Distribution
`l1:(`j `or` i`) = l1:j or l1:i`	or Distribution
`l1:(not` i`) = not l1:`i	not Distribution
`l1:` i`; l1:` j `= l1:` i `and` j`;`	term Distribution
`l1::`S`; l1::`T `= l1::`S `and` T`;`	Declaration Distribution

Label distribution works consistently for any Rosetta definition. An identical label can be distributed into or factored out of any logical or collection operation regardless of the types of it's arguments. For example, labels distribute over `and` in exactly the same manner whether the arguments are expressions, facets, or terms. Let's examine distribution law in two general classes: (i) boolean operations; and (ii) term and declaration distribution.

4.4.1 Distribution Over Logical Operators

Label distribution over logical operations follows the same process regardless of the specific operation. Namely:

`l1:` A `∘` B `== l1:` A `∘ l1:` B

for any logical operator `and`, `or` or `not`. The definition easily extends to cover cases for `=>`, `=` and other logical connectives. For example, the definition:

`l1: P(x) or Q(y)`

is equivalent to the definition:

`l1: P(x) or l1: Q(y)`

The semantics of each term depends on the specifics of the term value. In this case, if `P` and `Q` are both boolean valued operations, then the terms assert that the disjunction of the two properties holds. If the term types are facets, then the resulting definition defines and labels a new facet.

```
l : term;
```

where *l* is the term label and *term* is the term body. Any reference to the label *l* in the scope of this definition
refers to the term specified. Consider again the term from the earlier specification for EMI:

```
emiConst: emi power;
```

This simple definition defines a term `emiConst` that asserts `emi power`. Thus, the item referred to by `emi`
instantiated with the item `power` is asserted as a true statement.

Consider the following term involving a `let` expression:

```
l1: let (x::natural = 1) in inc x;
```

The label `l1` refers to the term defined by the `let` expression. Simplifying this definition based on the
definition of `let` results in the term `l1:  inc 1`. Given the classic definition of `inc`, this term is not legal
as it asserts the value associated with `inc 1`. The only condition where this could be legal is if `inc` returns
a boolean value or a facet.

4.3.3   Variable and Constant Labels

Labels for variable and constant items are labels for the objects they represent. Like term and facets,
variables and constants are also made visible using their label. Like all other items, variables are referenced
using the *name.label* notation where *name* is the facet label and *label* is the physical variable name. Consider
the definition of `power` from the earlier constraint facet:

```
power::real;
```

This declaration defines a variable item referenced by the label `power`. Outside the facet definition, this
variable is accessed using the notation `find.power`.

Constants work similarly. Consider the following constant definition:

```
pi :: real = 3.14159;
```

Within the scope of this definition, the label `pi` refers to the defined item whose value is the constant
`3.14159`.

It is important to remember that functions, types and facets are all items that can be declared within a facet.
Thus, they may all be referenced using their associated labels. Recall the definition of `increment_minutes`
from the alarm clock specification:

```
increment_minutes(t::time)::minutes is
    if m(t) =< 59 then m(t) + 1 else 0;
```

This definition is interpreted exactly like the previous constant definition. The label `increment_minutes`
refers to the item of type `time->minutes` whose value is specified by the constant function definition. Thus,
`timeTypes.increment_minutes` is used in the body of including specifications to reference the functions.
This practice of collecting declarations within facets will form the basis of the Rosetta `package` construct
defined later.

This definition produces a facet item labeled `find` whose type is `facet` and whose value results from parsing all declarations and terms within the facet.

Items declared in and exported from a facet visible outside the facet. Such items are referenced using the standard notation *name.label* where *name* is the facet label and *label* is the item label. For example, `key(o')` = k is accessed using the name `find.postcond1`. Consider the following facet definitions:

```
facet find_power is                   facet find_emi is
  power::real;                          power::real;
begin constraint_requirements        begin constraint_requirements
  heatConst: heatDiss power <= 10mW;
end find_power;                          emiConst: emi power;
                                     end find_emi;
```

These facets describe electromagnetic interference (EMI) constraints and heat dissipation constraints in facets labeled `find_emi` and `find_power`. Both facets are defined over a physical variable representing power consumption. Consider the composition of these facets into a single electrical constraints facet. The new facet is defined by conjuncting the `find_power` and `find_emi` facets and providing a new label, `find_electrical`:

```
    find_electrical :: facet is find_power and find_emi;
```

Note that this declaration is identical to all other Rosetta declarations. An item of type `facet` is declared and named `find_electrical`. Then, the value of `find_electrical` is constrained to be the product (conjunction) of `find_emi` and `find_power`. The declaration does not assert the facet in the current scope, but asserts that `find_electrical` references the new facet. This definition is the equivalent of saying:

```
    find_electrical :: facet;
  begin domain
    l1: find_electrical = find_emi and find_power;
```

Alternatively, a facet can be defined and referenced in a definition using the following form:

```
    find_electrical:: find_power and find_emi;
```

```
  %% Work on this.  It's a bit old and I think there's a much better
  %% way to assert facets as terms.
```

As stated earlier, all terms are boolean valued expressions. However, in the earlier definition `and` is used to define a new facet. The distinction is this form defines a new facet and asserts it to be true in the current context. Again it is named `find_electrical` and all labels and variables defined in it are accessed using `find_electrical`, not their original names. Facet labels do not nest, but instead the new label always replaces the old. Because no export clause is specified, `power`, `heatConst` and `emiConst` are all visible using the `find_electrical.label` notation. Further discussion of facet inclusion and assertion is presented in Section 4.5. It suffices here to understand that a new facet is being defined and it's resulting label is the assigned term label.

4.3.2 Term Labels

Each term defined in a facet must have a label. The Rosetta syntax allows labels to be omitted, however the resulting term's label is simply undefined and may be constrained to particular value by language tools. The label identifies the term and is effectively equal to the term throughout the facet definition. All term definitions have the form:

```
   inc(x::integer)::integer;
begin logic
   incdef: inc = <* (x::integer)::integer is x + 1 *>;
```

Semantically, this is identical to the standard definition. Like the previous definition, it is not as easy for the compiler to determine the value of `inc`.

The `let` form is also used to form terms. Consider the following definition:

```
l1: let (x::integer is a+1) in f(x,5);
```

When the `let` form is evaluated, the following term results:

```
l1: f((a+1),5);
```

Note that `let` currently supports defining variables over a single expression. Let forms cannot define variables over multiple terms.

Summary: A term is a labeled, boolean expression defined within the body of a facet. Each term is separated by a semicolon and is simultaneously true within the facet scope. Terms must be evaluated with respect to the domain associated with their enclosing facet to be fully interpreted.

4.3 Labeling

Labeling is the process of assigning a name to a Rosetta item. Facet definitions, item declarations, and terms all define items and provide labels. Recall that all Rosetta items consist of a label, value and type. Where the value and type define current and possible values associated with the item, the label provides a name used to reference the item. Specifically, labels serve as names for terms, variables, constants, and facets. Any item may be referenced using it's label. This provides the basis of reflection in Rosetta allowing Rosetta specifications to reference elements of themselves.

4.3.1 Facet Labels

Facet labels name facets and provide a mechanism for controlling visibility within a facet. Facets are labeled when they are defined directly. Further, they are defined when labeled terms define new facets from existing definitions using the facet algebra described in Chapter 3 and later in Chapter 5. When that label appears within a definition, it references the defined facet.

In a traditional facet definition, the facet name following the `facet` keyword becomes the defined facet's label. Consider the following definition of `find`:

```
%% Question: Why is a local var called 'power' declared in facet 'find'
%%  below? It is not used so it seems superfluous.

facet find(k::in keytype; i::in array[T]; o::out T) is
   power::real;
begin state_based
   postcond1: key(o') = k;
   postcond2: elem(o',i);
end find;
```

67

```
begin logic
  l1: x = f(x);
  ...
```

The domain for this term is `logic`, referring to Rosetta's basic mathematical system. There is no concept of state, time or change in this domain. Thus, `x = f(x)` is an assertion about `x` that must always hold. This domain is frequently termed the *monotonic* domain because change is not defined. If `f(x)` is not equal to `x`, then this term is inconsistent and the specification is in error.

The following term asserts that `x` in the next state is equal to `f` of `x` in the current state:

```
begin state_based
  l1: x' = f(x);
  ...
```

The `state-based` domain provides the basics of axiomatic specification. Specifically, the notion of current and next state. `x'` refers to the value of `x` in the state resulting from evaluating the facet's function. `x` refers to the value in the current state. This specification fragment has roughly the same semantics as an assignment statement as it specifies that `x` in the next state is equal to `f x`. Thus, if `x /= f(x)`, no inconsistency results. It is interesting to note that this statement is quite similar in nature to a basic signal assignment in VHDL. Specifically in VHDL:

```
  x <= f(x);
```

The following term asserts that `x` at current time plus `5ms` is equal to `f` of `x` in the current state:

```
begin continuous
  l1: x@(t+5ms) = f(x);
  ...
```

This specification is quite similar to the previous specification in that the value of `x` in some future state is equal to `f(x)`. It differs in that the specific state is defined temporally. Specifically, in the state associated with `5ms` in the future, `x` will have the value associated with `f(x)` where the argument to `f` is the value of `x` in the current state. Again, this definition bears some resemblance to VHDL signal assignments. This time, a wait statement is specified in conjunction with the signal assignment:

```
  x <= f(x) after 5ms;
```

Other domains and semantics are available for discrete time, constraints and mechanical specifications. The intent here is to demonstrate only the relationship between a term and its associated domain.

Another example uses classical axiomatic specification to define a function. The function `inc` has been used repeatedly as an example of constant function definition. Here, the function is defined as an abstract function and constrained using a term in the specification body:

```
    inc(x::integer)::integer;
  begin logic
    incdef: forall(x::integer | inc(x) = x + 1);
    ...
```

The definition states that for every integer, `x`, calling the function `inc` on `x` is equal to adding `1` to `x`. This is semantically equivalent to previous the previous definition, however it is more difficult for an interpreter to evaluate.

An alternate definition assigns a specific function to the function variable defined:

The type of an expression is the bunch of items resulting from all possible instantiations of the expression. For example, given the declaration `x::natural` and the expression `x+1`, the type of `x+1` is the bunch `ran(x::natural | x+1)` equl to the whole or counting numbers. Any function or facet parameter or variable may be legally instantiated with any expression of the same type.

4.2 Terms

A Rosetta term is a labeled, expression that appears within the scope of a `begin-end` pair within a facet. All terms are asserted as true within the scope of the facet. Note that simply because a term is boolean valued does not imply the term cannot represent an operational specification. It simply says that the statements made within the term are declared to be true.

The general format for a Rosetta term is a label, followed by an expression, terminated by a semicolon. Specifically:

```
label : term;
```

For example, the following term states that `inc 3` is equal to 4:

```
l1: inc(3) = 4;
```

Term label is `l1`, the term is `inc(3) = 4` and the semicolon terminates the term definition. Effectively, the semicolon terminates the scope of the label's assigned term. All terms defined in this fashion must be labeled.

The function of the semicolon is to terminate a labeled expression. Thus, the specification fragment:

```
l1: inc(3) = 4;
l2: forall(x::1++2 | x<4);
```

defines two terms with labels `l1` and `l2` and term expressions `inc(3) = 4` and `forall(<*x::1,2 -> x<4 *>)` respectively. In contrast:

```
l1: inc(3) = 4
l2: forall(x::1++2 | x<4);
```

is illegal as `l2:` is not an operation in the specification grammar.

```
%% Operators do in fact distribute over semicolons.  This is,
%% unfortunately wrong and needs to be corrected.
```

Semantically, the semicolon behaves as a conjunction. Terms delineated by semicolons in the body of a specification are simultaneously true and form a set of terms associated with the facet. This set of terms must be consistent with respect to the facet domain. A facet is consistent if and only if its domain, bunch of terms, and declarations are mutually consistent.

No term's semantic meaning can be inferred without reference to the including facet's domain. For example, the following definitions seem quite similar, but with proper interpretation mean quite different things. The following examples demonstrate this fact by showing how similarly defined terms have different semantics based on the definition domain. In each case, reference to the VHDL signal assignment semantics is mentioned to aide in understanding what is being specified. The various domains are explained in Chapter 6.

The following term asserts that `x` is equal to `f` of `x`:

Chapter 4

Expressions, Terms, Labeling and Facet Inclusion

4.1 Expressions

A Rosetta expression is constructed using operators and variables as defined in the current scope. Predefined operators and types are defined in Chapter 3 and form the basis of the Rosetta expression syntax. All rosetta expressions are recursively defined in terms of unary and binary operations. Parenthesized operations have the highest priority followed by unary operations and binary operations in traditional fashion.

Expressions are formed using unary operations, binary operations, grouping operations, and with function calls. The following general rules are used to define Rosetta expressions:

- Any constant a is an expression.
- Given any unary operation, $\circ$, and expression e, then $\circ e$ is also an expression.
- Given any binary operation, $\circ$, and expressions e_1 and e_2, then $e_1 \circ e_2$ is also an expression.
- Given any function, f, and expressions $e_1...e_n$, then $f(e_0, ..., e_{n-1})$ is an expression where n is the arity of f.
- Given any expression, e_1, then (e_1), $\{e_1\}$ and $[e_1]$ are also expressions.

Precedence for Rosetta unary and binary operations follow the canonical style. The following table lists Rosetta operators in tabular form:

Operator	Type
`(),{},[]`	Grouping
`-, not, %, $, #, ~`	Unary
`^ , in, ::`	Power and membership
`*, /, **`	Product
`+, -, ++, --, ;;`	Sum
`<, =<, >=, >, <<, >>`	Relational
`=, /=`	Equality
`min, and, nand`	Boolean product
`max, nmax, nmin, or, nor, xor, xnor, <=, =>`	Boolean Sum
`==`	Equivalence

Table 4.1: Precedence table for pre-defined Rosetta operations.

3.9.3 Composite Types

The following composite types are defined in Rosetta:

- `bunch` — The basic collection representation in Rosetta. A `bunch` is an unpackaged collection of items that resembles a `set` without the power-set concept. The notation `bunch(B)` (or `type(B)`) refers to any bunch of Rosetta items formed from the elements of bunch `B`.

- `set` — All possible sets that may be defined in Rosetta. A `set` is a packaged collection of items and is formed by packaging a bunch. The notation `set` refers to any set of Rosetta items. The notation `set(B)` refers to any set formed from the elements of bunch `B`.

- `sequence` — The basic ordered structure representation in Rosetta. A `sequence` is an unpackaged, indexed collection of times. The keyword `sequence` refers to any indexed collection of Rosetta items. The notation `sequence(B)` refers to any sequence of Rosetta items formed from the elements of bunch `B`.

- `bitvector` — A special sequence defined as `bitvector=sequence(bit)`.

- `string` — A special sequence defined as `string=sequence(character)`.

- `array` — All possible arrays that may be defined in Rosetta. An array is a packeted collection of indexed items and is formed by packaging a sequence. The keyword `array` refers to any indexed collection of Rosetta items. Then notation `array(B)` refers to any array of Rosetta items formed from the elements of bunch `B`.

- `univ` — The bunch containing `element` and all composite and constructed types.

3.9.4 User Defined Types

As types are first class items in Rosetta, user defined types are declared exactly as any other constant or variable. The following classes of type definitions are defined:

- `T::subtype(univ)` — Unconstrained type

- `T::type` — Unconstrained type (shorthand for previous)

- `T::subtype(`B`)` — Unconstrained subtype of `B`

- `T::subtype(univ) = `b — Defined type whose value is bunch b

- `T::type = `b — Defined type whose value is bunch b (shorthand for previous)

- `T::subtype(`B`) = `b — Defined subtype of B whose value is bunch b. $b::B$ must hold

3.9 Summary

The Rosetta type system provides a rich environment for specifying types and functions independent of any specific specification domain. Rosetta specifies three basic types of data types: (i) elements; (ii) composite types; and (iii) functions. Elements include basic data items that cannot be decomposed. Composite types include data structures containing other structures such as bunches, sets, sequences and arrays. Functions and function types include mappings from one bunch to another.

3.9.1 Declaring Types, Variables and Constants

Every Rosetta type is defined as a bunch. Thus, the operation for bunch membership is used to define variables and constants of a particular type. Specifically, defining a Rosetta variable of type `T` is achieved using the following notation:

```
v::T;
```

Defining a similarly typed Rosetta item with a constant value `c` is achieved using the following notation:

```
v::T is c;
```

Because Rosetta types are bunches, they are first class values and can be manipulated just as any other value would. Any operation defined on bunches can be used on a type. Additionally, the terms `bunch` and `type` are used interchangeably and are synonyms of each other.

3.9.2 Elements

The following element types are pre-defined in all Rosetta specifications:

- `null` — The empty bunch containing no elements.

- `boolean` — The named values `true` and `false`. `Boolean` is a subtype of `integer`. `True` is the greatest integer value while `false` is the smallest.

- `integer` — Integer numbers, from `false` to `true`. `Integer` is a subtype of `real`.

- `natural` — Natural numbers, from zero to `true`. `Natural` is a subtype of `integer`.

- `rational` — Rational numbers, consisting of two integers, a numerator and a denominator. `Rational` is a subtype of `real`.

- `real` — Real numbers consisting of all real valued numbers expressed as strings of decimals in traditional decimal or exponential form.

- `number` — Any legally defined Rosetta number.

- `character` All traditional character values.

- `element` — All elementary values, including all elements of all types named above and all values.

- `null` – The empty type.

```
sort(qs::boolean)::facet is
   sort_const and (if qs then quick_sort_req else sort_req);
```

In this definition, the parameter `qs` selects whether requirements for a quicksort or more general sorting requirements are included in the conjunction.

The following operators are defined over facets:

Operation	Format	Definition
and	F and G	co-product of F and G
or	F or G	product of F and G
not	- G, not G	Inverse of G
implies	F => G, F implies G	-F or G
=, /=	F = G, F /= G	Equivalence operations

Facet algebra operations are not logical operations. For example, `F and G` produces a new facet that is the product of `F` and `G`. The properties of each operation are defined by the category theoretic notations of co-product and product. When the co-product of two items is formed, the new item must have the properties of both the original items. Specifically, the facet `F and G` must have all properties of both `F` and `G`. When the product of two items is formed, the new item must have the properties of one or the other of the original items. Specifically, the facet `F or G` must have either properties of `F` or `G`. Negation is similarly defined. `not F` has none of the properties of `F`.

As facets are a Rosetta type, all operations defined over Rosetta bunches are also defined over facet types. Some specific examples include defining functions over facets and using quantification on functions with facet domains. This special property will be used heavily when defining reflection and defining meta-functions operations over facets. The facet algebra is described fully in Chapter 5. It is necessary here only to understand the mechanisms used to declare facets and the format of facet algebra expressions.

3.8.2 Facet Subtypes

Facet subtypes are defined by the use of domains in their definitions. Specifically, facet domains can be used as type and subtype qualifiers in their definition. For example, a facet `f` defined as follows:

```
facet f(x::in integer, z::out integer) is
begin finite-state
   t1: ... ;
   t2: ... ;
end f;
```

is considered to be of subtype `finite-state`. Thus, the declaration:

```
f::finite-state;
```

could be used to declare the same facet. Note that the details defined in terms and declarations from the previous facet are not included in this declaration.

Facet subtypes also provide a *domain polymporphism* capability. In the same way that `integer` is a subtype of `real` because integers are defined by restricting reals, the `finite-state` domain is a sub-domain of the `state-based` domain. This is true because the `finite-state` domain is obtained from the `state-based` domain by *extension* by adding new definitions to constrain the `state-based` domain. Thus, a homomorphism exists between the `state-based` and `finite-state` domains.

```
%% Summary section must be updated to reflect the chapter or
%% deleted.
```

```
lTree(t::NODE(lt,v,rt))::intTree is lt;

rTree(t::NODE(lt,v,rt))::intTree is rt;
```

The utility of pattern matching is more obvious here. The two functions return actual parameters associated with the constructor function `NODE`. Furthermore, both functions are defined only over trees constructed with `NODE` and are not defined over trees constructed with `NULL`. This is the desired result for high level specification.

In the definitions of `lTree`, `rTree` and `is_empty`, some or all of the constructor parameters are not used in the internal function. Thus, they need not be named in the definition. We use "_" to designate such a parameter as in the following:

```
is_empty(t::intTree)::boolean is
  (<*(t::NULL):boolean is true *> |
   <*(t::NODE(_,_,_))::boolean is false*>);

lTree(t::NODE(lt,_,_))::intTree is lt;

rTree(t::NODE(_,_,rt))::intTree is rt;
```

In both cases, parameters that are not used are not named or available in the function definition.

3.8 Facet Types

Like other items in Rosetta, facets are also a type defined by a bunch of items. The specifics of the facet type are defined in the following chapters and in the Rosetta Semantics Guide. Here, the declaration and an alternate definition mechanisms for facets are presented.

3.8.1 Facet Operations

In Chapter 2 a format for defining facets directly is provided. Specifically, the following defines a simple facet that increments an input value and outputs it:

```
facet inc(i::in integer; o::out integer) is
begin state_based
  l1: o'=i+1;
end inc;
```

Most basic facets will be described using this method.

In contrast, facets may be defined by composing other facets using the *facet algebra*. To achieve this, a facet is declared and assigned to the composition of other facets. An example from Chapter 2 describes the composition of requirements and constraints for a sorting component. Specifically:

```
sort :: facet is sort_req and sort_const;
```

This declaration follows the definitional style used for all Rosetta declarations. The label `sort` names the facet while the built-in type `facet` defines the collection of facets. In this case `and` forms a new facet from `sort_req` and `sort_const`. Note that `and` is a facet forming operation and not a boolean operation in this case.

Like types, parameterized facets may be defined using the function notation. The `facet` type is a type like any other and can be returned by functions. Thus, the signature of a parameterized `sort` facet definition is:

where v1 through vn name the specific values for fields f1 through fn. Defining a coordinate in Cartesian space using the definition above is achieved by:

```
cartFormer(1,0,0);
```

Accessing the result is achieved using the observer functions:

```
x(cartFormer(1,0,0))==1;
y(cartFormer(1,0,0))==0;
z(cartFormer(1,0,0))==0;
```

Using the "_" notation, it is possible to create records whose specific field values are not known. The following creates a cartesian coordinate whose x and y values are known, but whose z value is not specified:

```
cartFormer(1,0,_);
```

Should the function z be instantiated with this record, the return value is undefined.

3.7.3 Pattern Matching

Pattern matching in parameter lists dramatically simplifies defining observer functions over type constructors. Parameter matching takes advantage of the mechanism used to create its input parameters. Consider the integer tree definition presented above. Two constructor functions, NULL and NODE are defined to construct two different types of trees. Viewed differently, they also partition trees into the subclasses constructed by those individual functions. Specifically, the empty and nonempty trees. Viewed in this manner, it follows the the constructor functions can be used to generate types like any other types. For example:

```
nonemptyIntTree :: type is ran(NODE)
emptyTree :: type is ran(NULL)
```

Due to the nature of constructed types, the constructor for a particular instance of the type is always known. This fact can be utilized to perform pattern matching when instantiating function parameters. Consider the following definition of is_empty using selective union:

```
 is_empty(t::intTree)::boolean is
    (<*(t::NULL):boolean is true *> |
     <*(t::NODE(lt,v,rt))::boolean is false*>);
```

The first function accepts a single parameter of type NULL. This shorthand is equivalent to saying that t is contained in the bunch of all trees generated by NULL. Of course, this contains the single NULL tree. In the second definition, the type NODE(lt,v,rt) refers to all trees that can be constructed with NODE. Furthermore, lt, v and rt become parameters in the function that are bound to the actual parameters of any invocation of NODE. Specifically, in the following function call:

```
<*(t::NODE(lt,v,rt))::boolean is false*>(NODE(NULL,5,NODE(NULL,6,NULL)))
```

lt = NULL, v=5, and rt=NODE(NULL,6,NULL) within the scope of the function. These values are determined by matching the constructor function NODE with the parameter specification for t. The parameters are implicitly defined and their associated types determined from the constructor specification. Specifically, lt and rt are of type intTree while v is of type integer.

A more interesting case is defining accessor functions for the left and right subtrees of a nonempty tree. This is accomplished using the following definitions:

```
   b11(t::f1(x,_,_ ... _))::T11 is x;
   b12(t::f1(_,x,_ ... _))::T12 is x;
   ...

begin logic
   t1: forall(x::T | exists(x1::b11,...,xi::b1i | f1(x1,x2, ... xi) = x) or
                     exists(x1::b21,...,xj::b2j | f2(x1 x2 ... xj) = x) or
                     ...
                     exists(x1::bn1,...,xk::bnk | fn(x1 x2 ... xk) = x))
```

Constructed types, and constructor functions, are well known in the areas of functional programming and type theory. In addition to making a type theory decidable (which is not a factor in the design of Rosetta), they allow the user to distinguish otherwise identical types from each other, assuring that type checking and proofs will catch accidental substitution of one type of values for another. It also allows definitions of standard operations (such as addition and multiplication) to use different algorithms when necessary.

3.7.2   Records

In Rosetta, no special syntax for defining records is defined as record structures follow directly from constructed types. A record type is a constructed type with a single constructor function that associates values with parameters used as field names. A typical record type will be defined with the following constructive technique:

```
record::type is data
   recordFormer(f0::T0 | f1::T1 | ... fn::Tn)::recordp;
```

where `recordFormer` is the single constructor, `f1`through `fn` are the names of the various fields and `T1` through `Tn` are the types associated with those fields. The recognizer `recordp` is also defined, but is largely unused. To define a specific record type that represents Cartesian coordinates, the following notation is used:

```
cartesian::type is data
   cartFormer(x::real, y::real, z::real)::cartp;
```

To define an item of this type, the standard Rosetta declaration syntax is used:

```
c :: cartesian;
```

Values an be assocated with record items using the canonical `is` form:

```
origin :: cartesian is cartFormer(0,0,0);
```

Accessing individual fields of the record is achieved by applying one of the observer functions associated with a field name. To access field `y` in the record `c`, the following notation is used:

```
y(c)
```

Forming a record is achieved by calling the constructor function:

```
recordFormer(v1,v2,...,vn)
```

```
nullp(x::intTree)::boolean is x=NULL;

nodep(x::intTree)::boolean is
   exists(lt::intTree, v::integer, rt::intTree | NODE(lt,v,rt)=x);
```

Finally, parameter names are used to generate observer functions that return actual parameters from constructor functions. Specifically, the following functions are generated from the integer tree definition:

```
lt(t::sel(x::intTree | nodep(x)))::intTree
rt(t::sel(x::intTree | nodep(x)))::intTree
v(t::sel(x::intTree | nodep(x)))::integer
```

These functions return the actual parameter instantiation of their associated formal parameter. For example:

```
lt(NODE(NULL,1,NODE(NULL,2,NULL))) == NULL
v(NODE(NULL,1,NODE(NULL,2,NULL))) == 1
v(rt(NODE(NULL,1,NODE(NULL,2,NULL)))) = 2
```

In general, all constructed type definitions have the following general form:

```
T :: subtype(S) is data
   f1(b11::T11, b12::T12 ... b1i::T1i)::r1 |
   f2(b21::T21, b22::T22 ... b2j::T2j)::r2 |
   ...
   fn(bn1::Tn1, bn2::Tn2 ... bnk::Tnk)::rn ;
```

This data type definition defines n functions that create all elements of type T. Instantiating any of the fk functions creates an element of type T. Note that these functions have no body and are not evaluated. Associated with each fk is a recognizer function rk that is true when its argument was created with the constructor function. Specifically, rk will return true when its argument was created using fk. Finally, associated with each constructor function parameter is a function that observes the parameter. Given an instantiated constructor function, the observer associated with a parameter will return the actual parameter instantiating it. Like any other function, constructor functions can be partially evaluated. If this is the case, then the results of applying observer functions associated with uninstantiated parameters are not defined.

The general expression above is equivalent to the following definitions and laws (where the definitions are in the definition section and the laws in the predicate section):

```
   T :: subtype(S);
   f1(b11::T11, b12::T12 ... b1i::T1i)::T;
   f2(b21::T21, b22::T22 ... b2j::T2j)::T;
   ...
   fn(bn1::Tn1, bn2::Tn2 ... bnk::Tnk)::T;

   r1(t::T)::boolean is exists(b11::T11, b12::T12 ... b1i::T1i |
                                 f1(b11, b12 ... b1i)=t)
   r2(t::T)::boolean is exists(b21::T21, b22::T22 ... b2i::T2j |
                                 f2(b21, b22 ... b2i)=t)
   ...
   rn(t::T)::boolean is exists(bn1::Tn1, bn2::Tn2 ... bni::Tnk |
                                 fn(bn1, bn2 ... bni)=t)
```

57

Remembering that `subtype` is a synonym for bunch, the function signature defines a mapping from natural numbers to a bunch of bitvectors. That bunch of bitvectors is defined by the `sel` operations to be those whose lengths are equal to the parameter n. Thus, `boundedBitvector` will return the bunch of bitvectors of length equal to its parameter. We can now use `boundedBitvector` as as a type definition construct.

The notation:

```
reg::boundedBitvector(8);
```

defines `reg` to be a bitvector of length 8.

The notation:

```
bv8::subtype(bitvector) is boundedBitvector(8);
```

defines `bv8` to be the bunch of all bitvectors of length 8.

Summary: User defined types are declared exactly as are other Rosetta variables and constants. While the notation `x::T` forces `x` to be a singleton element of `T`, the notation `x::subtype(T)` allows `x` to be a bunch selected from `T`. Types can be formed from any element or composite type.

Uninterpreted types are defined as subtypes of the `univ` type.

Parameterized types are defined by using functions to return bunches as types.

3.7   Constructed Types

3.7.1   Defining Constructed Types

Rosetta provides a general shorthand for defining types in a constructive fashion. Constructor, observer and recognizer functions are defined for the type and encapsulated in a single notation. These types are called *constructed types* and are created with the special `data` keyword and notation. As an example, consider a definition for a binary tree of integers:

```
intTree :: type is data
  NULL::nullp |
  NODE(L::intTree,v:integer,R::intTree)::nodep;
```

This definition provides two constructors for `intTree`: (i) the nullary function `NULL`; and (ii) the ternary function `NODE`. The `NULL` function creates an empty tree while the `NODE` function creates a node from a value and a left and right subtree. A tree with one node whose value is 0 can be generated with the following function instantiation:

```
NODE(NULL,0,NULL);
```

A balanced tree with 0 as the root and 1 and 2 as the left and right nodes respectively can be generated:

```
NODE(NODE(NULL,1,NULL),0,NODE(NULL,2,NULL));
```

The recognizers `nullp` and `nodep` indicate the constructor used to generate a tree. Specifically, `nullp` is true if its argument is `NULL` and `nodep` is true if its argument is an instantiation of the node function. Semantically, these functions are defined as follows:

Here, the value of `natural` is known to be the bunch of integers greater than or equal to zero. Like T, `natural` is a sup type of `integer`, but its value is given by the expression.

In general, the notation:

```
T::subtype(supertype);
```

defines a new type T whose value is not specified, but constrained to be any bunch contained *supertype*. Such variable types are again referred to as uninterpreted subtypes. Their specific values are unknown, but they are restricted to be a sub-bunch of their associated supertype.

Finally, the notation:

```
T::subtype(supertype) is b;
```

Declares a new type whose supertype is a sub-bunch of *supertype* and whose value is b. It is implied that b::*supertype*. If the supertype is left unconstrained, the new type will have no explicit supertypes.

Consider the definition of the type `bit` as a subtype of `natural`. The specific definition of `bit` has the following form:

```
bit::subtype(natural) is 0++1;
```

Because `0++1 :: natural`, this represents a perfectly legal type definition.

Consider the definition of the type `natural` as a subtype of `integer`. The specific definition of `natural` has the following form:

```
natural::subtype(integer) is sel(x::integer | x >= 0);
```

As Rosetta is declarative, there is no reason why the declaration cannot include an expression. Further, it is quite possible that that expression may not be evaluated until analysis time.

In addition to constructing new types comprised of elements, the `subtype` construct can be used to define types comprised of composite values. The following definition:

```
time::subtype(record[h::hours | m::minutes | s::seconds]);
```

defines a new type named `time` that is comprised of records. Similarly, it is possible to define types containing sets, sequences, and arrays.

One final note about the distinction between the notations `x::T` and `x::subtype(T)` as declarations. The first says that the value of x is a *single* element of type T. The second says that the value of x is a *bunch* of values selected from T. If the first definition is used as a type, then only single element types are allowed. The second explicitly allows bunches. In contrast, if these statements are used as terms in a specification, they are identical. This distinction will become clearer when terms are presented.

3.6.1 Parameterized Types

Any function returning a bunch can be used to define a Rosetta parameterized type. Consider the following function definition:

```
boundedBitvector(n::natural)::subtype(bitvector) is
  sel(b::bitvector | $b = n);
```

Antiderivative, or indefinite integral, is the inverse of derivative. The antiderivative of `f` with respect to `x` is expressed as:

```
antideriv(f,x,c) =  ∫  f(x)dx  +  c
```

`f` being the function in question, `x` being the variable integrated over, and `c` being the constant of integration.

As antiderivative is the dual of derivative, the following axiom is defined for all real valued functions:

```
antideriv(deriv(f,x),x,0) == deriv(antideriv(f,x,0),x) == f
```

The definite integral of `f` with respect to `x` over the range `u` to `l` is expressed as:

$$\texttt{integ(f,x,l,u)} == \int_{l}^{u} f(x)dx$$

The definite integral is defined as the difference of the indefinite integral applied at the upper and lower bounds:

```
integ(f,x,l,u) == antideriv(f,x,0)(u) - antideriv(f,x,0)(l)
```

It is possible to express a definite integral over an infinite range using the notation:

$$\texttt{integ(f,x,false,true)} = \int_{-\infty}^{\infty} f(x)dx$$

It should be noted that limit, derivative, antiderivative and integral functions are defined over real valued functions only. Further, the functions provide a mechanism for expressing these operations and some semantic basis for them. Solution mechanisms are not provided.

3.5.5 Univ Types

```
%% Make sure this definition is correct
```

The type `univ` is now introduced to contain all `element`, `bunch`, `set`, `sequence`, `array`, and constructed types. This type contains all basic data types provided in the Rosetta type system. It differs from `universal` in that it does not contain function types, facet types or other types traditionally used to represent non-data constructs.

3.6 User Defined Types

User defined types are declared in the same manner as any constant or variable. The notation:

```
T::subtype(integer);
```

defines a variable type `T` whose value is unspecified, but whose elements must come from the integers. The expression specified in the `subtype` declaration is the supertype of the type being defined. Recall that the expression `subtype` is synonymous with bunch, thus `T` is a subtype of the integers. Such variable types are referred to as uninterpreted as their values cannot be determined at compile time. The natural numbers can be defined from the integers using the following definition:

```
natural::subtype(integer) is sel(x::integer | x >= 0);
```

We wish to determine:

```
    succ(n::natural)::natural is n+1 :: succ(n::natural)::natural;
==  (natural::natural) and forall(<*n::natural -> succ(n)::natural*>)
==  true               and forall(<*n::natural -> (n+1)::natural*>)
==  true               and true
==  true
```

Assume that `f(x::df)::rf` and `g(x::dg)::rg`. The following operations are defined over two functions:

Relation	Format	Definition
=,/=	f = g, f /= g	f::g and g::f
		-(f::g) or -(g::f)
>=, =<, >, <	f >= g, f > g	g::f
		f > g = g::f and f /= g

Functional equivalence checks to determine if every application of `f` and `g` to elements from the union of their domains results in the same value. Specifically, `f(x) = g(x)` for every `x` in either domain. Function inequality is defined as the negation of function equality.

Function inclusion comes in several forms and is specified using the same relational operators used for numerical relations, subset and sub-bunch relationships. One function is included in another (`f=<g`) when applying the including function to every element of the included function's range is the same as applying the included function. The distinction here is that the range of the included function, not both functions, is evaluated. Proper inclusion `f<g` occurs when `f` is included in, but not equal to `g`. It is easy to show that `f =< g and g =< f == f = g`. Further, the operations define a partial order over functions.

3.5.4   Limits, Derivatives and Integrals

A special class of functions for defining limits, derivatives and integrals are provided for use with real valued functions. These functions exist primarily to allow specification of differential equations (both ordinary and partial) over real valued functions. Given a real valued function `f(x::real)::real`, the following definitions are provided:

Function	Format	Definition
Limit	lim(f,x,n)	$\lim_{x \to n} f(x)$
Derivative	deriv(f,x)	$\frac{df}{dx}$
Indefinite Integral	antideriv(f,x,c)	$\int f(x)dx + c$
Definite Integral	integ(f,x,u,l)	$\int_{l}^{u} f(x)dx$

The derivative of a function is defined with using limit in the canonical fashion. The following axiom is defined for all real valued functions and real valued nonzero `delta`:

```
deriv(f,x) = lim((f(x+delta)-f(x))/(x+delta)-x,delta,0)
```

In the derivative function, `f` is the object function and `x` is the label of the parameter subject to the derivative. In the above function, the following holds:

```
deriv(f,x) = 
```
$\frac{df}{dx}$

The derivative function is generalizable to expressing partial derivatives. Assuming that `g` is defined over multiple parameters, such as `g(x::real;y::real;z::real)::real`, then:

```
deriv(g,x) = 
```
$\frac{\partial g}{\partial x}$

Thus, increment is of type `<*(x::natural)::  natural*>`. It is interesting to note that this is exactly the relationship that must be checked for every definition of the form:

```
f(x::R)::S is T;
```

as it indicates that the actual function is of the same type as the specified signature.

The final case defines when one constant function is included in another. In this case, both `d1` and `d2` are expressions and the most general expression of the function inclusion law must be applied.

First, consider determining if the increment function is included in itself. Clearly, this should be the case and the function inclusion law supports the assertion:

```
inc :: inc
   == dom(inc) :: dom(inc) and forall(x::natural | inc(x) :: inc(x))
   == natural :: natural and forall(x::natural | x+1 :: x+1)
   == true                and true
```

This holds because for any Rosetta item, `i::i` holds by definition.

Consider the case of determining if increment is contained in identity over natural numbers. In this case, the law should not hold:

```
inc :: id
   == <*(x::natural)::natural is x+1*> :: <*(x::natural)::natural is x*>
   == dom(inc) :: dom(inc) and forall(x::natural | inc(x) :: id(x))
   == natural :: natural and forall(x::natural | x+1 :: x)
   == true                and false
```

`false` is obtained from the second expression by the counter example provided by x=0 as `0+1 /= 0`.

When `f1::d1->e1` where e1 is an expression, the following holds:

```
f1 :: <*(x::d2)::r2*> == d1::d2 and forall(n::d1 | f1(n) :: r2*>)
```

or

```
f1 :: <*(x::d2)::r2*> == d1::d2 and ran(r1) :: r2

%% Note that in both the immediately preceding equation and an early
%% equation, dom(r1) was replaced with ran(r1).  There may be too
%% much coincidence here to let that go unchecked.
```

The function containment law gives the criteria by which one function may be said to be included within a function type. Each function type defines a bunch of functions consisting of all those functions included in it. This means that any function can be used as a type, or bunch, and all the containment laws for bunches apply to them. This is particularly useful when using a function that returns a bunch rather than a single value. Consider the function `<*(n::natural)::natural*>`. This function defines the bunch of all functions that take a natural number as an argument and return a natural number. Rosetta allows the user to ask if a given function is contained in that bunch (is a member of that type). For example, consider:

```
succ(n::natural)::natural is n+1;
```

3.5.3 Function Types and Inclusion

The function type former `<*(d::domain)::range*>` defines the bunch of functions mapping `domain` to `range`. This bunch is in all ways a Rosetta type and can be manipulated as a bunch. Thus, operations such as bunch containment are defined over functions.

Bunch containment applied to functions is referred to as function containment. Function containment, `f1::f2`, holds when a function is fully contained in a function or function type. Assuming `f1(x::d1)::r1` and `f2(x::d2)::r2` where `d1`, `d2`, `r1` and `r2` are expressions (potentially bunches):

```
%% The original definition prior to rewrite:
%% f1 :: f2 == d1::d2 and forall(<* x::d1 -> r1 x :: r2 x *>)


f1 :: f2 == d1::d2 and forall(x::d1 | r1(x) :: r2(x))
```

where `r1(x)` is the result of instantiating the expression associated with `f1` with `x` and evaluating the result.

`f1` is contained in `f2` if and only if the domain of `f1` is contained in the domain of `f2` and for every element of `f1`'s domain, `f1(x)` is contained in `f2(x)`. Exploring function inclusion's several cases reveals how it applies in several situations.

The simplest case is when `r1` and `r2` are specified as bunches where the parameter `x` is not involved in the definition. Examining the function inclusion law, the universal quantifier falls out and the following relationship results:

```
f1 :: f2 == d1::d2 and r1 :: r2
```

In this case, `f1` is included in `f2` when: (i) its domain is contained in `dom(f2)`; and (ii) its range is contained in `ran(f2)`.

A second case occurs when `r1` is an expression and `r2` is a bunch. Instantiating the function inclusion law results in the following statement:

```
f1 :: f2 == d1::d2 and forall(x::d1 | f1(x) :: r2)
```

`r2` is a constant value independent of `x`. Therefore, the law requires that the result of applying expression `r1` to actual parameter `x` result in an element of `r2`. This is actually equivalent to the previous result and can be simplified to:

```
f1 :: f2 == d1::d2 and ran(f1) :: r2
```

As an example, consider the increment function defined over natural numbers. It should hold that:

```
inc :: <*(x::natural)::natural*>
```

Instantiating the function inclusion law gives:

```
inc :: natural -> natural
     == dom(inc) :: natural and forall(x::natural | inc(x) :: natural)
     == natural :: natural and forall(x::natural | x+1 :: natural)
     == true                and true
```

and returns `true` if every `x` selected from `B` is greater than 0. The notation allows specification of the domain on the left side of the bar and the expression on the right. The domain of the expression is assumed to be boolean for `forall`, `exits`, and `sel`. For `min` and `max`, the domain is taken from the expression. This notation is substantially clearer and easier to read than the pure functional notation. Note that the original notation is still valid for specifying quantified functions.

The notation extends to n-ary functions by allowing parameter lists to appear before the "|" to represent parameter lists. The format of these lists is identical to the format of function parameter lists. Specific examples include:

```
forall(x,y::integer | x+y>0)
exists(x,y::integer | x+y>0)
sel(x,y::integer | x+y>0)
```

It is important to remember that like `forall` and `exists`, `sel` observes the function range and selects appropriately. Interpreting the notation in the standard way results in the definitions:

```
forall(<* (x,y:integer)::boolean is x+y>0 *>)
exists(<* (x,y:integer)::boolean is x+y>0 *>)
sel(<* (x,y:integer)::boolean is x+y>0 *>)
```

Summary: Quantifier functions operate on other functions. Each generates the range of their function argument and returns a specific value associated with that range. `min` and `max` return the minimum and maximum range values respectively and are synonymous with `forall` and `exists`. `sel` and `ran` provide comprehension and image functions respectively. `sel` applies a specific boolean expression to a function's range and returns a bunch of domain elements satisfying the expression. `dom` returns the domain of a function defined as the application of the result expression to every domain element.

3.5.2 Set Constructors and Quantifiers

There are no direct mechanisms for defining sets by comprehension or for providing quantification. Quantification functions over bunches in conjunction with the set former and elements functions provide an effective means for achieving such definitions. For example, Given a set `S` and a boolean expression `P`, the set of objects from `S` satisfying `P` is defined as:

```
{sel(x::~ S | P(x))}
```

The result of evaluation is the set formed from the bunch of items from `S` satisfying `P`. Effectively, the select operation forms the bunch of set elements by: (i) extracting the bunch of elements from `S` using the contents operator; (ii) filtering the resulting bunch for those elements satisfying `P`; and (iii) using the set former to construct a set from the bunch. Using a similar approach, other operations from bunches can be quickly applied to set objects. Assuming that `S::set(T)`, `F::<*(t::T)::univ*>`, and `P::<*(t::T)::boolean*>`

Set Operation	Definition
`max(x in S|F(x))`	`{max(<*(x::~ S)::univ is F(x)*>)}`
`min(x in S|F(x))`	`{min(<*(x::~ S)::univ is F(x)*>)}`
`ran(x in S|F(x))`	`{ran(<*(x::~ S)::univ is F(x)*>)}`
`sel(x in S|P(x))`	`{sel(<*(x::~ S)::boolean is P(x)*>)}`
`forall(x in S|P(x))`	`{forall(x::~ S)::boolean is P(x)*>)}`
`exists(x in S|P(x))`	`{exists(x::~ S)::boolean is P(x)*>)}`

Note that these shorthands are defined in terms of the dual bunch operations. Users are advised to use the shorthand notations whenever possible rather than risk definition error.

```
forall(<* (x::B)::boolean is x>0 *>)
```

Here, the domain of the argument function is the bunch B and the result expression x>0. To determine the range of the argument function, x>0 is applied to each element of B. Assume that B=1++2++3. Substituting into the above expression results in:

```
forall(<*(x::1++2++3)::boolean is x>0 *>)
```

Applying the result expression to each element of the domain, the range of the function becomes:

```
true++true++true
```

As `true` is greater or equal to all boolean values, the minimum resulting value is `true` as expected. Assuming B=-1++0++1 demonstrates the opposite effect. Here, the range of the internal function becomes:

```
false++false++true
```

As `false` is less than `true`, the minimum resulting value is `false`. Again, this is as expected.

It is important to note here that `forall` and `exists` behave identically to `min` and `max` for boolean valued functions. The `min` and `max` functions applied in the same way would result in the same outcome. `forall` and `exists` provide useful shorthands and are not absolutely necessary in the larger language. Thus, `max` and `min` are referred to as quantifiers.

The function `sel` provides a comprehension operator over boolean functions. The signature for `sel` is defined as follows:

```
sel(<*(x::univ)::boolean*> is bunch(univ)*>)
```

Like `min` and `max`, `sel` observes the range of the input function. However, instead of returning a single value, `sel` returns a bunch of values from the domain that satisfy the result expression. Consider the following example where `sel` is used to filter out all elements of B that are not greater than 0:

```
sel(<*(x::B)::boolean is x>0*>)
```

Assuming B=1++2++3, x>0 is true for each element. Thus, the above application of comprehension returns 1++2++3. If B=-1++0++1 then x>0 holds only for 1 and the application of comprehension above returns 1.

3.5.1 Notation Issues

A shorthand notation is provided to make specifying `forall`, `exists`, `sel`, `min` and `max` expressions simpler. Notationally, the following statement:

```
forall(x::B | x>0);
```

is equivalent to:

```
forall(<*(x::B)::boolean is x>0*>);
```

The `ret` function takes a function and returns its defined return type. This is the type specified in the function definition following parameters and limits the values that can be returned by the function definition. The `ran` function is similar to a bunch comprehension function and returns the image of a function with respect to its domain. It returns the bunch resulting from applying the parameter function's expression to each element of the domain. By definition, `ran(F)::ret(F)`, but it is not necessary for the range of a function to be equal to its return type. Consider the following example where `ran` is used to add one to each element of bunch B:

```
ran(<*(x::B)::natural is x+1*>)
```

Given that `B=1++2++3`, the expression above evaluates to `2++3++4`. This is precisely the application of `x+1` to each element of the range bunch. Applying `ret` to the same function would return `natural` as the return type.

The `dom` function is defined similarly to the range function but instead returns the domain associated with its function argument. For example:

```
dom(<*(x::B)::natural is x+1*>)
```

evaluates to the bunch B.

The domain and range functions present a greater challenge when dealing with functions of arity other than 1. The domain of a nullary function is defined as the `null` type while the range of a nullary function is the result of its evaluation:

```
dom(<* 3+2 *>) == null
ran(<* 3+2 *>) == 5
```

Using this identity, one can define evaluation of a fully instantiated function as taking the range of that function. Specifically, if all arguments to a function are known, then the range of that instantiated function is the same as evaluating the function.

The domain of functions with arity greater than one is defined as the type of the first parameter. The range of such a function is defined as the result of evaluating the function over all possible parameter values. Specifically, given the canonical `add` function:

```
add(x,y::natural)::natural is x+y;
```

domain and range are defined as:

```
dom(add) == natural
ran(add) == natural
ret(add) == natural
```

As previously defined, the functions `min` and `max` provide minimum and maximum functions as defined previously. Interestingly, `min` and `max` provide quantification functions `forall` and `exists` when F is a boolean valued function. Recall that `true` and `false` are defined as the maximum and minimum integer values respectively. As noted earlier, `and` and `or` correspond to the binary relations `min` and `max` respectively. This fact is born out by introduction of truth tables. As `forall` and `exists` are commonly viewed as general purpose `and` and `or` operations, this makes some sense.

Consider the following application of `forall` to determine if a bunch, B, contains only integers greater than zero:

```
    let (v1::T1 be e1, v2::T2 be e2, ..., vn::Tn be en) in exp ;
```

where `v1` through `vn` define variables, `T1` through `Tn` define the types associated with each variable, and `e1` through `en` define expressions for each variable.

Evaluating the `let` expression results in the expression *exp* with each variable replaced by its associated expression. The semantics of the let expression are defined using function semantics. It is sufficient to realize that the `let` provides local definitions for expressions.

The `if` expression provides a simple mechanism for expressing choice. The general form:

```
   if expression then cond1 else cond2 end if;
```

evaluates to *cond1* if *expression* evaluates to `true` and *cond2* if *expression* evaluates to `false`.

3.5   Bunch Constructors and Quantifiers

Other important operations available for functions include what are traditionally called quantifiers. In Rosetta, all quantifiers are functions defined over other functions. As an example, consider the `min` function. The signature for the `min` function is:

```
   min(f::<*(x::univ)::univ*>)::univ
```

The `min` function accepts an arbitrary function and returns the minimum value associated with the range of the argument function. Recall that the range of the argument function is the result expression applied to each element of the domain. Consider the following function application:

```
   min(<*(x::1++2++3++4)::natural is x*2 *>)
```

The domain of the argument is the bunch `1++2++3++4`. Although it is unusual to define a bunch by extension in these circumstances, it is perfectly legal. The range of the argument function is the expression applied to each element of the domain. Specifically, `2++4++6++8`. The `min` function then returns the minimum value associated with `2++4++6++8` or 2.

If the unaltered minimum value associated with the input bunch is desired, the `min` function can be applied using an identity function as in:

```
     min(<*(x::1++2++3++4)::natural is x *>)
```

The `max` function is defined similarly and operates in the same manner.

A number of second order functions such as `min` and `max` are defined and will be presented here. Given that `F(x::univ)::univ` and `P(x::univ)::boolean`, the following quantifier functions are defined:

Operation	Format	Traditional name
Function Former	`<* (rank)::return is exp *>`	Forms a function value or type.
dom	`dom(F)`	Domain
max	`max(F)`	Maximum value
min	`min(F)`	Minimum value
sel	`sel(P)`	Comprehension
ret	`ret(F)`	Return type
ran	`ran(F)`	Range or Image
exists	`exists(P)`	Existential quantifier
forall	`forall(P)`	Universal quantifier
+	`+ B`	Summation
*	`* B`	Product

Summary: A Rosetta function is defined by specifying a domain, range and an expression defining a relationship between domain and range elements. The notation:

```
f(d::domain)::range is exp;
```

Defines a function mapping `domain` to `range` where `exp` is an expression defined over domain parameters and gives a value for the associated range element. The notation:

```
f(d::domain)::range;
```

defines `f` as an element of the bunch of all functions relating *domain* to *range* without specifying the precise mapping function.

As an example, the increment function is defined over naturals using the notation:

```
inc(x::natural)::natural is x+1;
```

The angle brackets (`<* *>`) define the scope of the named parameter `x` while `x+1` defines the result.

Applying a function is simply substitution of an actual parameter for a formal parameter. Evaluating `inc(2)` involves replacing `x` with 2 and applying the definition of a function. Specifically:

```
inc(2) = 2+1 =
inc(2) = 3
```

Function types are specified as anonymous functions using the notation:

```
<*(d::domain)::range*>
```

This function type specifies the bunch of all functions mapping *domain* to *range*. The definition:

```
f::<*(d::domain)::range*>
```

says that `f` is a function that maps *domain* to *range*.

An anonymous function is a function having no assigned label. It is treated like a lambda function in Lisp programming languages in that it can be evaluated like any other function, but has no name by which to reference it.

```
<*(d::domain)::range is exp*>
```

This anonymous function specifies the function mapping *domain* to *range* using the expression *exp*. It is semantically the same as `f(d::domain)::range`. The definition:

```
f::<*(d::domain)::range is exp*>
```

says that `f` is a function that maps *domain* to *range* using the expression `exp`. It is semantically the same as `f(d::domain)::range is exp`.

The `let` expression provides a mechanism for defining local variables and assigning expressions to them. This provides shorthand notations that can dramatically simplify complex specifications by reusing specification fragments. The syntax of the general `let` expression is:

```
let (x::integer be i+1) in i'=x;
```

The semantics of this `let` expression is:

```
<*(x::integer)::univ is i'=x*>(i+1)
```

Evaluation of the function application gives:

```
i'=i+1
```

The usefulness of `let` becomes apparent when an expression is used repeatedly in a specification. Consider a facet with many terms that reference the same expression. The `let` construct dramatically simplifies such a specification.

Let expressions may be nested in the traditional fashion. In the following specification, the variable x of type T has the associated expression *ex1* while y of type R has the expression *ex2*. Both may be referenced in the expression *ex3*.

```
let (x::T be ex1) in let (y::R be ex2) in ex3;
```

This expression may also be written as:

```
let (x::T be ex1, y::R be ex2) in ex3;
```

The semantics of this definition are obtained by applying the previously defined semantics of `let`:

```
<* (x::T)::univ is <*(y::R)::univ is ex3 *>*>(ex1)(ex2)
```

Consider the following specification assuming that `i` is of type integer and `fcn` is a two argument function that returns an integer:

```
let (x::integer be i+1, y::integer be i+2) in i'=fnc(x,y);
```

When evaluated, the following function results:

```
<*(x::integer)::univ is <*(y::integer)::univ is i'=fnc(x,y)*>*>(i+1)(i+2)
```

The result of evaluating this function is:

```
<*(x::integer)::univ is <*y::integer is i'=fnc(x,y)*>*>(i+1)(i+2) ==
<*(y::integer)::univ is i'=fnc(i+1,y)*>(i+2) ==
i'=fnc(i+1,i+2)
```

In order for normal argument substitution to work, the expressions in the Rosetta `let` expression must not be mutually recursive. If recursion is necessary, the expressions must be represented as normal Rosetta rules or predicates.

```
case exp0 is
   b1 -> exp1 |
   b2 -> exp2 |
   b3 -> exp3 |
   ...
   bn -> exprn
```

where *b1-bn* are bunches and *exp0-expn* are expressions. The `case` statement evaluates to *expk* when exp0
`::` expk holds. If this relationship is satisfied by multiple bunches, the `case` expression evaluates to the
expression associated with the first such bunch. A default case can be achieved using `univ` as the bunch
expression. The equivalence check performed in most traditional languages is performed by using singleton
bunches in the expression. For example:

```
case x is
   sel(x::integer | x > 0) -> false |
   0 -> true |
   sel(x::integer | x < 0) -> false
```

implements a zero test on the x.

Please note that the case statement is semantically equivalent to a unary function defined using selective
union. Specifically:

```
(<*(x::sel(i::integer | x > 0))::boolean is false *> |
 <*(x::0)::boolean is true *> |
 <*(x::sel(i::integer | x < 0))::boolean is false *> )
```

is a function with identical semantics to the previous `case` statement.

3.4.7 The Let Expression

Function application provides Rosetta with a mechanism for defining expressions over locally defined vari-
ables. An additional language construct, the `let` expression generalizes this providing a general purpose
`let` construct. The Rosetta `let` is much like a Lisp `let` in that it allows definition of local variables with
assigned expressions. The general form of a let expression is:

```
let (x::T be ex1) in ex2;
```

This expression defines a local variable x of type T and associates expression *ex1* with it. The expression
ex2 is an arbitrary expression that references the variable x. Each reference to x is replaced by *ex1* in
the expression when evaluated. The `be` keyword is semantically equivalent to `is` and is included only for
readability. Like traditional parameter definitions, `let` parameter definitions may omit the `be` clause and be
variable.

The syntax of the `let` expression is defined by transforming the expression into a function. Specifically, the
semantic equivalent of the previous `let` expression is:

```
<*(x::T)::univ is ex2*>(ex1)
```

When the function is applied, all occurrences of x in `ex2` are replaced by *ex1*. This process is identical to
the application of any arbitrary function to an expression. Assume the declaration `i::integer` and consider
the following `let` expression:

3.4.5   The If Expression

The Rosetta `if` expression is a polymorphic function that supports choice between options. The syntax of the `if` expression is:

```
%% Check to see if the compiler supports ''end if'' or ''endif''

%% Ans: Parser only supports 'endif' --D.S., June 3, 2001



   if exp1 then exp2 else exp3 endif;
```

where `exp1` must be of type `boolean` while `exp2` and `exp3` may be of arbitrary types.

The rules for evaluating an `if` expression are quite simple, but differ from the `if` statement in an imperative language. When `exp1` is `true`, the expression evaluates to `exp2`. When `exp1` is `false`, the expression evaluates to `exp3`. Specifically:

```
   if true then a else b endif == a
   if false then a else b endif == b
```

The domain of an `if` expression is simply `boolean` while the range is `ran(a)++ran(b)`.

For convenience, an `elsif` construct is provided to nest `if` statements. The notation:

```
   if exp1 then exp2
     elsif exp3 then exp4
     elsif exp5 then exp6
     ...
     else expn
   endif
```

is semantically equivalent to:

```
   if exp1 then exp2
     else if exp3 then exp4
             else if exp5 then exp6
                     ...
                     else expn endif
          endif
   endif
```

3.4.6   The Case Expression

The `case` expression supports selection from multiple options in a manner similar to using the `if` construct with the `elsif` extension.

The general form of a case statement is:

The definition of a function whose value is $(x + 1)^2$ can be defined as:

```
<*(z::integer)::integer is sqr(inc(z))*>
```

Expanding the definitions of `sqr` and `inc` gives the following function:

```
<*(z::integer) is <*(y::integer) is y^2*>(<*(x::integer) is x+1*>(z))*>
```

The only available simplification is to substitute y's actual parameter in to the expression for `sqr` giving:

```
<*(z::integer) is <* (<* (x::integer) is x+1*>(z))^2 *> *>
```

Continuing to substitute, replacing formal parameters with actuals and eliminating function formers when parameters are replaced gives:

```
<*(z::integer) is <* (<* z+1 *>)^2 *> *>
  == <*(z::integer) is <* (z+1)^2 *> *>
  == <*(z::integer) is (z+1)^2 *>
```

The result is a new function defined over `z` that gives the result of composing `inc` and `sqr`.

Selective Union

Selective union of two functions `f` and `g` is defined formally as:

```
%% Should dom be rank in the following to allow multi-parameter
%% functions?

(f|g)(x) = if x::dom(f)
              then f(x)
              else if x::dom(g)
                      then g(x)
           endif;
```

Using the `if` construct insures that the function associated with the first including domain will be called. In the above example, if `dom(f)=integer` and `dom(g)=real`, then an integer value will cause `f` to be selected while only a `real` value that is not an `integer` will cause `g` to be selected. If the domains are reverse, *i.e.* `dom(f)=real` and `dom(g)=integer`, `g` will never be selected because any element of `integer` is also in `real`. If the type of `x` is in none of the domains specified, then the function is undefined.

The domain and range of `(f|g)` is defined as:

```
dom(f|g) == dom(f)++dom(g);
ran(f|g) == ran(f)++ran(g);
```

Selective union is highly useful for implementing a form of polymorphism. An example of a function defined by selective union is the simple
ttfamily non_zero function:

```
non_zero(n::number)::boolean is
   (<*(n::0) is true*> |
    <*(n::sel(x::integer | x>0))::boolean is false*> |
    <*(n::sel(x::integer | x<0))::boolean is false*>)
```

This is a rather pedestrian use of selective union and there are better definitions of the `non_zero` function. However, it does demonstrate how domain values can be used to select from among different function definitions.

Partial Evaluation

Partial evaluation is the process of taking a function and instantiating only a subset of it's parameters. The semantics of partial evaluation are defined using currying as described previously. Here, only the usage and application of partial evaluation are discussed.

Consider the following definition of f over real numbers:

```
f(x::real;y::real;z::real)::real is (x+y)/z;
```

Application of f follows the traditional rules of substituting actual parameters for formal parameters in the expression and substituting the expression for the function. Partial evaluation will perform the same function, but will not require instantiating all parameters. Consider a situation where f is applied knowing that in all cases the value of z will be fixed at 2 to perform an average. The following syntax partially evaluates f and assigns the resulting function to the new function name avg:

```
avg(x::real;y::real)::real;

avg=f(_,_,2);
```

In this definition, the "_" symbol is used as a placeholder for a parameter that will not be instantiated. To calculate the value of f(_,_,2), we simply follow the instantiate and substitute rule:

```
f(_,_,2) == <*(x::real;y::real)::real is (x+y)/2*>;
```

The result is a 2-ary function that returns a real value. As noted, this function calculates the average of its two arguments. An alternate, more compact notation for the definition is:

```
%% There is a problem here and earlier.  An expression follows the
%% is, not a function.  We must modify the definition of is or
%% provide and alternate notation when we want to explicitly specify
%% the function value rather than the expression.  The latter is far
%% more common, so we'll stick with the nice notation for that.

avg(x::real;y::real)::real is f(_,_,2);
```

This general approach is applicable to functions of arbitrarily many value.

Function Composition

Function composition is an application of function definition capabilities. Assume that two functions, f and g exist and that ran(g)::dom(f). We can define a new function h as the composition of f and g using the following definition style:

```
h(x::R)::T is f(g(x));
```

The approach extends to other definition styles in addition to the direct definition style.

Consider the definition of inc and a function sqr defined as:

```
sqr(y::integer)::integer is y^2;
```

41

Currying Multi-Parameter Functions

Technically, evaluation of multi-parameter function is achieved by a process based on the concept of a curried function. This process provides the basis of the evaluation process used previously. All Rosetta functions can be expressed as functions of a single argument, or curried functions. Specifically, the function:

```
<* (x::R;y::S)::T is exp *>
```

can be expressed equivalently as:

```
<* (x::R)::<*(y::S)::T is exp*> *>
```

The new function is expressed is unary function over items of type `R` that returns another unary function that maps items of type `S` to type `T`. Given a function `f(x::R,y::S)::T` and `r::R` and `s::S`, the following equivalence holds:

```
f(r,s) == f(r)(s)
```

Given the definition above, this equivalence generalizes to functions of arbitrarily many variables.

Consider again the definition of `add` defined over two integer numbers:

```
add(x,y::integer)::real is x+y;
```

Using the previous notation, `add` can be expressed in a curried fashion as:

```
add(x::integer)::<*(y::integer)::integer is
    <*(y::integer)::integer is x+y *>;
```

The use of `x` in the expression is perfectly legal as the second function definition is done in the scope of `x`.

The `add` function is now defined over a single parameter of type `integer`. Its return type is no longer an `integer` value, but a function that maps one `integer` onto another. Using this notation, adding two values `a` and `b` is achieved using `add(a)(b)` - exactly the notation discussed previously.

Now consider application of the `add` function using the curried function approach:

```
add(1,2) == add(1)(2)
add(1,2) == <*(x::real)::<*(y::real)::real is x+y*>*>(1)(2)
add(1,2) == <*(y::real)::real is 1+y*>(2)
add(1,2) == <*1+2*>
add(1,2) == 3
```

The function is evaluated by substituting an actual parameter for a formal parameter in first the original add function and then in the unary function returned by add, exactly as it was done in the previous section.

It is particularly interesting to note the following equivalence:

```
add(1) == <*(x::real)::<*(y::real)::real is x+y*>*>(1)
add(1) == <*(y::real):real is 1+y*>
```

This process, called currying, defines the semantics of multi-parameter functions and is the basis for all function evaluation.

```
<* (z::integer)::integer is z+1 *>(<* (x,y::natural)::natural is x+y *>(4,5))
```

Instantiating the argument to what was the increment function results in the following definition:

```
<* <* (x,y::natural)::natural is x+y *>(4,5) +1 *>
```

As the resulting function has no arguments, the outer function former can be dropped resulting in the new anonymous function:

```
<* (x,y::natural)::natural is x+y *>(4,5) +1
```

Instantiating the x parameter of the new function by replacing the formal parameter with its associated actual parameter results in:

```
<* (y::natural)::natural is 4+y *>(5) + 1
```

Finally, instantiating the y parameter in the same manner results in:

```
<* 4+5 *> + 1
```

Again, as there are no variables associated with the resulting function, the function former can be dropped resulting in the expected result:

```
4+5+1 == 10
```

The same result occurs regardless of the order of instantiation. The following shows a different order resulting in the same result:

```
<* (z::integer)::integer is z+1 *>(add(4,5))
    == <* add(4,5)+1 *>
    == <* <*(x,y::integer)::integer is x+y *>(4,5)+1 *>
    == <* <*(x::integer)::integer is x+5 *>(4) + 1 *>
    == <* (x::integer)::integer is x+5 *>(4) + 1
    == <* 4+5 *> + 1
    == 4+5+1
    == 10
```

3.4.4   Partial Evaluation, Function Composition and Selective Union

Thus far, all Rosetta functions have been defined by using an expression in the function definition. Rosetta provides three additional mechanisms for function construction: (i) curried functions and partial evaluation; (ii) function composition; and (iii) selective union. Partial evaluation generates new functions by substituting values for some parameters and simplifying. Function composition is simply an application of function definition techniques that allow a new function to be constructed from existing functions. Selective union allows functions to be specified by extension.

```
<* (x,y::natural)::natural is x+y*>
```

defines an anonymous function identical to the function `add` above, except the anonymous function has no label. Such function definitions can be used as values and are evaluated in exactly the same manner as named functions. Specifically:

```
<* (x,y::natural)::natural is x+y *>(1,4) == <* 1+4 *> == 5
```

Anonymous function signatures can also be defined in a similar manner. The definition:

```
add::<* (x,y::natural)::natural *>
```

is equivalent to the function signature defined above without associating the signature with a name. We call this a function type because it defines a collection of functions mapping pairs of natural numbers to natural numbers. Technically, it defines a bunch of functions that map two natural numbers to a third natural number. This is true because the function's expression is left unspecified. Because the definition represents a bunch of functions, it can be used as a type in formal definitions. The definition:

```
add::<* (x,y::natural)::natural *>
```

is semantically equivalent to the earlier `add` signature definition. The definition says that `add` is of type `<*(x,y::natural)::natural*>` or that `add` is a function that maps two natural numbers to a third natural number. The earlier signature definition is a shorthand for this notation provided to make definitions easier to read and write.

The definition:

```
add::<* (x,y::natural)::natural *> is <* (x,y::natural)::natural is x+y*>
```

is equivalent to the first definition of `add` above and semantically defines the function definition shorthand. The `add` function is defined as a variable of type function. The declaration asserts that the `add` function is equivalent to the anonymous function value mapping two integers to their sum. Any function label can be replaced by its value, if defined.

The notation `<* *>` is refereed to as a function former because it encapsulates an expression with a collection of local symbols to form a function. The brackets form a scope for locally defined parameters. When no parameters are present, the function former brackets can be dropped as there is no need to define parameter scope.

Rosetta provides the special type `function` that contains all functions definable in a specification. Stating that `f::function` says that `f` is a function, but does not specify its `domain` or `range` values.

3.4.3 Function Evaluation

Evaluation of Rosetta functions allows for both currying and partial evaluation. All Rosetta functions are evaluated in a lazy fashion. Arguments can be instantiated and functions evaluated in any order. Consider the following use of `inc` and `add`:

```
inc(add(4,5))
```

Rather than evaluating in the traditional style, expand the function definitions Using the canonical definitions of `inc` and `add` results in the following anonymous function:

3.4 Functions

Defining functions in Rosetta is a simple matter of defining mappings between bunches. Functions are extensive mappings between two different types, called the *domain* and *range* of the function. Functions are defined by defining their domain and stating an expression that transforms elements of the domain into elements of the range. This is achieved by introducing variables of the domain type whose scope is confined to the function definition and defining a result expression using that variable. The domain is given by an expression that describes the domain bunch. The range is given by an expression using the variable introduced by the function, called the result or result function.

3.4.1 Direct Definition

The direct definition mechanism for defining functions is to define the function's signature and an expression that relates domain values to a range value. Functions are typically defined by providing a signature and an optional expression. The syntax:

```
add(x,y::natural)::natural is x+y;
```

In this definition, `add` is the function name, `x,y::integer` defines the domain parameters, and `integer` is the return type. Together, these elements define the signature of `inc` as a function `add` that accepts two arguments of type `natural` and evaluates to a value of type `natural`.

Following the signature, the keyword `is` denotes the value of the return expression, in this case `x+y`. The return expression is a standard Rosetta expression defined over visible symbols whose type is the return type. Literally, what the function definition states is that anywhere in the defining scope `inc(x,y)` can be replace by `x+y` for any arbitrary integer values. Whenever a function appears in an expression in a fully instantiated form, it can be replaced by the result of substituting formal parameters by actual parameters and evaluating the resulting expression. Specifically, if the function instantiation `add(3,4)` appears in an expression, it can be replaced by the expression `3+4` and simplified to `7`. Any legal expression can be encapsulated into a function in this manner.

Like any Rosetta definition, `add` is an item with an associated type and value. In this case, the type of `add` is a function type defining a mapping from two integers into the integers. The value of `add` is known and is a function encapsulating the expression `x+y`. The specific value resulting from function evaluation is determined by its associated expression.

A function signature can be defined separately by specifying its arguments and return type without an expression. The notation:

```
add(x,y::natural)::natural;
```

defines the signature of a function `add` that maps pairs of `natural` into `natural`. Because this `add` definition has no associated expression, it is referred to as a function signature. Allowing the definition of a signature without an associated expression supports flexibility in the definition style. In this case the expression associated with `add` can be defined directly using equality or indirectly by defining properties. Function signatures help dramatically in reducing over-specification in definitions.

3.4.2 Anonymous Functions and Function Types

Anonymous functions, frequently called lambda functions, are defined by excluding the name and encapsulating the function definition in the function former `<* *>`. The definition:

Thus, `array(univ)` is any array while `array(B)` restricts possible arrays to the elements of B.

A special case of an array is the `bitvector` type. Formally, `bitvector` is defined as:

```
bitvector = array(bit);
```

Operations over `bit` are generalized to `bitvector` by performing each operation on similarly indexed bits from the two bit vectors. Assuming that `op` is any `bit` operation, the `bitvector`, C, result of applying the operation over arbitrary `bitvector` items A and B is defined by:

```
C(n) = A(n) ∘ B(n)
```

If either A or B is longer, then the shorter bitvector is padded to the left with 0s. to achieve the end result.

In addition, the following operations are defined over items of type `bitvector`: (Assume `A::bitvector`)

Operation	Format	Definition
bv2n, n2bv	bv2n(A)	Convert between bitvectors and naturals
twos, swot	twos(A)	Generate two's compliment
lshr,lshl,ashr,arshl	ashr(A), lshl(A)	Logical and arithmetic shift right and left
rotr,rotl	rotr(A), rotl(A)	Rotate right and left
padr, padl	padr(A,1,n), padl(A,0,n)	Pad right and pad left with value to n bits.

The operations `bv2n` and `n2bv` provide standard mechanisms for converting between binary and natural numbers. It is always true that `bv2n(n2bv(x))=x`.

The operations `twos` and `swot` take and invert the two's compliment of a binary value. The `lshr` and `lshl` operations provide provide logical shift right and left while `ashr` and `ashl` provide logical and arithmetic shifts right and left. The distinction being that logical shift operations shift in 0s while arithmetic shift operations shift in 1. The `rotr` and `rotl` operations provide rotation or circular shift. Note that none of the compliment, shift, or rotate operations change the length of the bit vector.

The `padr` and `padl` operations pad or catenate a bit vector. Each takes three arguments: (i) the bitvector being manipulated; (ii) the pad value (1 or 0); and (iii) the resulting length. If the length value is less than the length of the argument vector, `padr` removes bits to the right and `padl` removes bits to the left resulting in a vector of length n. In this case, the pad value is ignored.

Summary: The following predefined composite types are available in a Rosetta specification:

- `bunch` — The basic collection representation in Rosetta. A `bunch` is an unpackaged collection of items that resembles a `set` without the power-set concept. Rosetta types are bunches. The keyword `bunch` is synonymous with the keyword `type`. The declaration `bunch(univ)` refers to any collection of Rosetta items. The notation `bunch(B)` (or `type(B)`) refers to any bunch of Rosetta items formed from the elements of bunch B.

- `set` — All possible sets that may be defined in Rosetta. A set is a packaged collection of items and is formed by packaging a bunch. They declaration `set(univ)` refers to any set of Rosetta items. The notation `set(B)` refers to any set formed from the elements of bunch B.

- `sequence` — The basic ordered structure representation in Rosetta. A `sequence` is an unpackaged, indexed collection of times. The declaration `sequence(univ)` refers to any indexed collection of Rosetta items. The notation `sequence(B)` refers to any sequence of Rosetta items formed from the elements of bunch B. The type `string` is a special sequence defined as `string=sequence(character)`.

- `array` — All possible arrays that may be defined in Rosetta. An array is a packeted collection of indexed items and is formed by packaging a sequence. The declaration `array(univ)` refers to any indexed collection of Rosetta items. Then notation `array(B)` refers to any array of Rosetta items formed from the elements of bunch B. The type `bitvector` is a special array defined as `bitvector=array(bit)`.

refers to any sequence constructed from the elements of B.

To define an item of type `sequence`, the following notation is used:

```
x::sequence(univ);
```

where `x` is the new item and `sequence(univ)` refers to the set of all possible Rosetta sequences. Thus, `sequence(univ)` is any sequence while `sequence(B)` restricts possible sequences to the elements of B.

A special case of a sequence is the `string` type. Formally, `string` is defined as:

```
string = sequence(character)
```

A shorthand for forming strings is the classical notation embedding a sequence of characters in quotations. Specifically, `"ABcdEF"` = `'A';;'B';;'c';;'d';;'E';;'F'`. This definition implies that strings are ordered, but are not packaged.

Functions defined over strings include those defined for general sequences. In particular, the notation `''abc''` `;;` `''def''` is appropriate for concatenation of strings. It is important to note that the ordering operations for sequences provide lexigraphical ordering for strings. No additional function definitions are required.

3.3.4 Arrays

Arrays are indexed, packaged collections of elements; thus, arrays of arrays (multi-dimensional arrays) are allowed. Arrays are enclosed in brackets; we allow `[]` as an abbreviation for `[nil]`. If `A` and `B` are arrays packaging sequences `S` and `T` respectively, `n` is a natural number, `I` is a sequence of natural numbers, and `e` is an element, the following operations are defined on arrays:

Operation	*Format*	Definition
Formation	`[S], [1;;2;;3]`	Array from sequence
Contents	`~ A`	S
=,/=	`A = B, A /= B`	`S = T, S /= T`
access	`A(n)`	`S(n)`
subscription	`A sub I`	`[S sub I]`
catenation	`A ^ B`	`[S;;T]`
replacement	`n->e|A`	`[n->e|S]`
>=, =<, >, <	`A >= B, A > B`	`S >= T , S > T, and S /= T`
size	`# A`	`$ S`
max,min	`A max B, A min B`	`[S max T], [S min T]`
empty	`[]`	`[nil]`

Arrays are to sequences what sets are to bunches. Packaging a sequence forms an array. Further, array operations have dual sequence operations in the same manner that sets have dual bunch operations. These functions are defined formally in the Rosetta Semantics document.

To define an item of type `array` whose contents are taken from type B, the following notation is used:

```
x::array(B);
```

To define an item of type `array` whose contents are taken from any arbitrary type, the following notation is used:

```
x::array(univ);
```

Operation	Format	Definition
=,/=	S = T, S /= T	Equality defined on elements
access	S(n)	nth element of S from 0
subscription	S sub I	Sequence from S corresponding to I
catenation	1;;2;;S;;T	Concatenation
integer sequence	i;;..j	Sequence of integers from i to j
replacement	n->e|S	Copy S with element n replaced by value e
<,=<,>,>=	S < T, S > T	Corresponding order on elements
size	$ S	Unpackaged size
max,min	S max T, S min T	Order defined on elements
nil	nil	The empty sequence

```
%% Note that head and tail have been removed in deference to the
%% more general subscription operation that is now defined.
```

Most operations operated as expected based on common usage. Equal and not equal take their canonical meanings. The catenation operator, ;;, concatenates two sequences. Like bunches, any singleton item can be a sequence. Thus, ;; can form all sequences from the empty sequence.

Subscription is an extraction mechanism where elements from a sequence are extracted to form a new sequence. Give S and an integer sequence I, S sub I extracts the elements from S referenced by elements of I and forms a new sequence. For example:

```
A;;B;;C;;D sub 0;;2;;1 == A;;C;;B
```

The notation i;;..j forms an integer sequence from i running to j. As an example, The functions head and tail can be defined using subscription and integer sequence as follows:

```
head(S) = S(0)
tail(S) = S sub 1;;..$S
```

S(0) returns the first element in the sequence. The integer sequence former 1;;..$S forms the integer sequence from 1 to the length of S. Extracting elements of S associated with 1 through $S includes all elements except the first and thus defines tail in the canonical fashion.

The replace operator allows replacement of an element within a list. The notation n->e | S generates a new sequence with the element in position n replaced by e. For example:

```
2 -> 5 | 1;;2;;3;;4;;5 == 1;;2;;5;;4;;5
```

The ordering operations as well as min and max are defined much like lexagraphic ordering. If x;;S<y;;T, then either x<y or x=y and S<T. Note that for any sequence, S /= nil, S > nil. S=<T is defined as S<T or S=T. The S min T and S max T operators return the minimum sequence of S and T and the maximum sequence respectively.

```
%% Note: because sequences are not packaged, they are not elements and
%% cannot belong to types (bunches); members of types must be arrays or
%% sets.
```

To define an item of type sequence containing only elements from type B, the following notation is used:

```
x::sequence(B);
```

good time to point out that `A::B` is defined for any bunch `A`, not just a singleton bunch. This is in contrast to the `in` set operation that is distinct from the subset operation. The containment operator for bunches operates over all bunches.

Size returns the cardinality of the set while `empty` names the empty set. The `contents` function returns the unpackaged bunch and is the dual of the set former.

Note that laws for sets do not eliminate the possibility of paradox; it is not syntactically or semantically impossible to encode Russell's Paradox, for example. In some cases, the Rosetta user will have to include a cardinality argument, or a proof that a given set is well founded.

Defining items of type is achieved using similar notations to defining items of type bunch. The following notation defines `x` to be an element of all possible sets composed of items from bunch `B`:

```
x::set(B);
```

The declaration may intuitively be read as 'x is a set of items from `B`'.

As for bunches, an item that is an arbitrary set containing any legal Rosetta values is defined:

```
x::set(univ);
```

Thus, `set(univ)` is any set while `set(B)` restricts possible sets to the elements of `B`.

The notation `x::set(B)` restricts the value of item `x` to single sets of items from `B`. It is also useful to define new set types where we restrict the value of a new item to bunches of sets from `B`. This is achieved using the following notation:

```
X::bunch(set(B));
```

or

```
X::subtype(set(B));
```

Both notations are equivalent and define a new bunch that includes possible subsets of `B`. The original notation `x::set(B)` declares that `x` is a singleton bunch contained in the bunch of all subsets of `B`. The new notation declares that `X` is *any* bunch contained in the bunch of all subsets of `B`. Although mathematically not equivalent to a powerset, the result is similar to type theoretic systems that use powersets to define new types. Using the above notations, `X` is restricted to bunches of subsets of `B` that can then be used as types. These bunches can be further restricted as in:

```
set4 :: subtype(set(B)) is sel(x::bunch(set(B)) | forall(t::x | #t=4));
```

where `set4` is the bunch of subsets of `B` that contain exactly four elements. The notation `z::set4` declares `z` to be a singleton element of the bunch `set4`, or simply a subset of `B` containing four elements.

3.3.3 Sequences

Sequences are indexed, but unpackaged, collections of elements. As with bunches, sequences of sequences are not defined. Because the elements of sequences are ordered by their mapping, a lexicographic total order exists on the sequence whenever a total order exists between all of the members of the sequence. The simplest sequence is `nil`, the empty sequence. If `S` and `T`, are sequences, `n` a natural number, `e` an element, and `I` a sequence of natural numbers, the following operations are defined on sequences:

When used in this context, the type `B` is said to be the supertype of type `X` as by definition `X::B` in the general sense.

It is also possible to define a bunch whose values are completely unconstrained at declaration time using the notation:

```
x::type;
```

This notation is a shorthand for making the supertype `univ`, the bunch containing all possible Rosetta items. Semantically, it is the same as the notation:

```
x::bunch(univ);
```

or

```
x::subtype(univ);
```

where `bunch(univ)` and `subtype(univ)` refer to the Rosetta bunch containing all items. This notation and semantics is particularly useful when defining uninterpreted types whose values are not enumerated, but are constrained by later specifications. Such definitions are not typical in traditional programming languages, but can be useful in abstract specification situations where types are defined via properties rather than via extension.

3.3.2   Sets

Sets are packaged bunches that exhibit properties traditionally defined in classical set theory. Most operations on sets have their dual in bunches with the exception of power-set. It is impossible to have a bunch of bunches while it is possible to have a set of sets. This fact can be observed in the distinction between the set relations `in` and `subset` and the bunch relation contained in ("`::`"). As bunches are not packaged, there is no distinction between a bunch and an element. Singleton elements are distinct from sets of size one while singleton elements and bunches of size one are not.

In the following, `B` and `C` are bunches, and `S` and `T` are the sets that package them, respectively, i.e. `S = {B}` and `T = {C}`. The first table lists functions that form sets from bunches or other sets:

Operation	Format	Definition
Formation	{B}, {1++2++3}	
*, +, -	S*T, S+T, S/T	{B**C},{B++C},{B--C}
power	power(T)	({B} in power(T)) = B::C

The basic set former takes an arbitrary bunch and forms a set containing its elements. Operations for intersection, union and difference are defined by forming a set from the bunches resulting from the bunch duals of these operations. The power-set operation creates all possible sets from all bunches contained in the bunch representing the contents of the original set. Note again that power-set has no dual in bunches.

Operation	Format	Definition
=,/=	S = T, S /= T	B = C, B /= C
in	a in S	a::B and $a=1
<, =<, >=, >	S >= T, S > T...	B::C, B >> C...
size	# S	$ B
contents	~ S	A
empty	empty	{null}

Other set functions listed above observe properties of sets and typically have duals in bunches. Equivalence is equivalence of contents, subset and proper subset are defined in terms of inclusion. This is a particularly

be singleton bunches. Unlike sets where both union and add are needed to define composing two sets and adding an element to a set, bunches require no such distinction. Thus, the notation 1++2++3 forms a bunch containing 1, 2 and 3.

When we say x::A, x may be either a singleton or multiple element bunch. In contrast for sets, $\{a\} \in S$ and $a \in S$ are two very different things. Bunches make no such distinction, thus both concepts are embodied in the notation a::S. Every bunch contained in A or B is also contained in A++B. Formally, x::A or x::B => x::A++B. Thus, ++ is referred to as a union operation. The ** operation, bunch intersection, is defined similarly as is bunch difference, --.

An interesting result related to variable manipulation is that given a variable declaration x::B and a bunch y, the following holds:

```
y::x and x::y == x=y
```

If we know we are dealing with singleton values, then we further know that inclusion and equals are similar. Specifically, the declarations x::T and y::R both implicitly assert that x and y are singleton bunches. In this special case, inclusion is equals. Specifically:

```
x::y == y::x == x=y
```

This does *not* hold for other situations. If the cardinality of x or y is not 1, then the definition does not apply.

A shorthand notation, i++..j, is provided for generating bunches of integers over a range. The notation generates a bunch containing all integers from i to j inclusive. Bunches are not ordered, so the result is not equivalent to a sequence.

The values of all Rosetta types are bunches. The terms type and bunch are synonyms and may be substituted for each other at any place in this document. Further, as bunches are values, type values may be manipulated in the same manner as any Rosetta value. This includes reference and dynamic formation within a Rosetta specification. Although this dramatically complicates type checking, the added flexibility is useful and well behaved language subsets can (and will be) be defined.

To declare an item whose value is selected from bunch B, the following declaration notation is used:

```
x::B;
```

Used in the declaration section, this statement defines an item x whose value is a single element taken from the bunch B.

To define an item whose value is contained in bunch B and is not limited to a single element, the following declaration notation is used:

```
x::bunch(B);
```

The item x now refers to any bunch constructed from the elements of B.

The primary use for bunch definition is when defining a new type. Thus the subtype synonym is introduced. The following notation is semantically identical to the previous bunch declaration:

```
x::subtype(B);
```

- `number` — Any legally defined Rosetta number.

- `character` — All unicode values associated with characters.

- `element` — All elementary values, including all elements of all types named above and all values.

- `enumerations` — The enumeration construct allows the definition of new value items. The result of an enumeration is the addition of the value item to the `element` type. The notation `enumeration[x,y,z]` is used to define three new element values and assembles them into a bunch that can be used as a type.

3.3 Composite Types

Composite types make complex values by combining simpler values. There are two mechanisms for structuring: (i) packaging; and (ii) indexing. Packaging transforms a structure into an element that can be included in other structures. Specifically, bunches are constructed using the notation `A++B` while sets are constructed with element and union. Sets distinguish between a singleton set and an element of a singleton set. As bunches are unpackaged, they do not. Indexing establishes a function from the natural numbers (from zero to the size of the structure minus one) to the elements of the structure. There are four kinds of structures, depending on the mechanisms used to build them:

Packaged?	Indexed?	Structure
no	no	bunch
yes	no	set
no	yes	sequence
yes	yes	array

We will define each of these in turn.

3.3.1 Bunches

Bunches are collections of elements. Any element is a bunch of size one. Because bunches are not packaged, bunches may not consist of other bunches. A bunch is defined by creating it with operations, as given in the following list of operations (where `A` and `B` are bunches):

Operation	Format	Definition
union	A++B	x::A++B = x::A or x::B
intersection	A**B	x::A++B = x::A and x::B
difference	A--B	x::A++B = x::A and - (x::B)
sub-bunch	A<<B, A>>B	A::B and A/=B, B::A and B/=A

In most circumstances, bunch operations are distinguished from non-bunch operations by simply repeating the operator symbol twice. For example, "+" is traditionally sum or union while "++" is bunch union. Such distinction is necessary as all Rosetta items are bunches. Thus, operator symbols are used to distinguish bunch from non-bunch operations.

Additionally, the following functions are defined over bunches:

Operation	Format	Definition
=,/=	A = B, A /= B	Equivalence
$	size C	Number of bunch equivalence
inclusion	A::B	Contained in
sequence	i++..j	Integer bunch

The key to understanding bunches is realizing the effects of the union and inclusion operators. The primary bunch former is the `++` operator. Given any two bunches `A` and `B`, the operator `A++B` defines the union or composition. Specifically, every element of `A` and `B` is also a member of `A++B`. Remember that `A` and `B` may

is semantically equivalent to adding the new value items `apples`, `oranges`, and `pears` to `univ` and assembling them into a new bunch.

Enumerations can be used to define new types using the canonical Rosetta notation:

```
fruit :: subtype(element) is enumeration[apple,orange,pear];
```

This declaration create a new item whose supertype is a bunch contained in `univ` and whose value is the bunch containing the value items `apple`, `orange`, and `pear`. A variable defined as:

```
x :: fruit;
```

must take its value from the bunch `apple++orange++pear`.

Only value items may be included in enumeration declarations. Thus, the pair of declarations:

```
x :: integer;
c :: enumeration[ x,1,z ];
```

is illegal because `x` is an item, but not a value item. Assuming that if `z` has been defined, it is defined as a value item, the declaration:

```
c :: enumeration[ 1,z ];
```

is semantically legal. A new value item `z` is created in the current scope and the enumeration evaluates to the bunch containing `1++z`. Note that having been defined as a value item by the enumeration, `z` cannot be used as anything else in the current scope. Upon leaving the enumeration's scope, `z` is again available for use as any item name.

Summary: The following predefined elemental types are predefined for all Rosetta specifications:

- `null` — The empty bunch containing no elements.

- `bit` — The enumerated type containing only the values 0 and 1. Bit is a subtype of `natural`. Operations on `bit` elements correspond to operations on the booleans in the canonical fashion. Note that this implies that operators such as "+" are specialized for bits such that `1+1=1` (bit) rather than `1+1=2` (natural).

- `boolean` — The named values `true` and `false`. Boolean is a subtype of `integer`. True is the greatest integer value while `false` is the smallest. Thus, the boolean operators `and`, `or` and `not` correspond to `min`, `max` and "-" respectively.

- `integer` — Integer numbers, from `false` to `true`, including all integral numeric values. Integer is a subtype of `real`, specifically `real` restricted to integral values.

- `natural` — All natural numbers, from zero to `true`. Natural is a subtype of `integer`, specifically `integer` restricted to non-negative values.

- `rational` — Rational numbers, consisting of two integers, a numerator and a denominator. Rational is a subtype of `real`, specifically `real` restricted to those values that can be expressed as the ratio of two integers.

- `real` — Real numbers consisting of all real valued numbers expressed as strings of decimals in traditional decimal or exponential form.

3.2.5 Characters

Character constants are defined in the traditional. Character values are referenced using the notation `'x'`.

Type	Format	Subtype Of
character	`'A'`,`'a'`	element

Given `a::character` and `b::character`, operators on characters include:

Operator	Format	Definition
ord,char	`ord(a)`, `char(a)`	Unicode value
=,/=,<,=<,>=,>	`a<b`, `a=<b` ...	`ord(a)<ord(b)`, `ord(a)=<ord(b)` ...
uc,dc	`uc(a)`, `dc(a)`	Raise/lower case

`uc` and `dc` change case of a character to upper or lower case respectively. These operators return their argument unchanged for Unicode values other than alphabetic characters. The `ord` operator returns Unicode values associated with characters. The `char` operator is its dual. Thus, there is a mapping between Unicode values and unicode characters; `ord` and `char` express the mapping, which is one to one. Unicode literals are expressed using the standard notation 'U+XXXX' where XXXX is a 4-digit, hexadecimal number. All other laws are the same as for their natural number counterparts with respect to `ord`, and no further laws will be given here. For example, assuming that `x` and `y` are characters, `x<y` is defined:

```
x < y == ord(x) < ord(y)
```

3.2.6 Elements

The type `element` is provided to contain all items of any elemental type. It is formally defined as all the bunch union of all elemental type definitions. The type is defined by taking the bunch union of all types defined in this section. Formally:

```
element == number ++ character ++ enumerations
```

3.2.7 Null

A special bunch called `null` is defined to be the bunch containing no elements. Formally, `null` is defined as:

```
null::bunch(univ);
forall(x::univ | -(x::null));
```

Thus, the type `null` is the type containing no elements. It is rarely used in the specification process, but must be included for completeness. As it contains no elements it is neither an element or composite type. Its definition is included here for convenience.

3.2.8 Enumerations

Enumerations provide a mechanism for defining new elemental values and types by extension. When an enumeration is declared, two semantic operations are performed on the list of value items associated with the enumeration. First, the list of values associated with the enumeration are added to `univ` as elemental types if they are not already present. Second, the list of value items assembled into a bunch associated with the construction. For example, the following notation:

```
enumeration[ apple,orange,pear ]
```

A	B	-A	-A max B
false	false	true	true
false	true	true	true
true	false	false	false
true	true	false	true

This is precisely the definition of implication. Reverse implication works similarly and the definition of equivalence (A=B) is consistent with the above definition. Further, when values are restricted to `boolean`, the following equivalences hold:

```
A => B == A =< B
A <= B == A >= B
A <= B and B <= A == A = B
```

It should be noted that Rosetta does not define logical equivalence, `iff`, separately from numerical equivalence. Given the mathematical definition of booleans, the normal equivalence operations are sufficient.

3.2.3 Bit

Bits are a subtype of the natural numbers. The type `bit` consist of the two values 0 and 1 and is defined as `bit=0++1`.[1] Operations over `bit` mimic operations over `boolean`. However, `bit` and `boolean` are distinct. Specifically, an item of type `boolean` cannot have a value of type `bit` and an item of type `bit` cannot a value of type `boolean`. Conversion operators are defined to provide definitions for various operators. Specifically, the following operation maps between `bit` and `booleans`:

Operation	Format	Definition
%	% 1, % true	Converts between bits and boolean in the canonical fashion

The following operations are defined over bits where `x::bit`, `y::bit`, `a=% x`, and `b=% y`:

Operation	Format	Definition
=,/=	x = y, x /= y	Axiom
=>, <=, >, <, >=, =<	x => y, x <= y, ...	a => b, a <= b, ...
max,min,and,or	x max y, x min y, x and y, ...	a max b, a min b, a and b, ...
not	- x	% -a
xor,xnor	x xor y, x xnor y	a xor b, a xnor b
nmin,nmax	x nmin y, x nmin y	a nmin b, a nmin b

All `bit` operations are defined by converting to `boolean`, applying the appropriate `boolean` operation and converting back. This suggests that the definition of `bit` is redundant. However, `bit` is of such importance in digital design that a representation specific to it is important. The true distinction between `bit` and `boolean` arises when considering the `bitvector` type later.

3.2.4 Numbers Revisited

The type `number` is provided to represent any legal Rosetta number. It is defined as the bunch union of all types defined thus far and is effectively equivalent to `real` in the Rosetta numbering system. Formally, `number` is defined as:

```
% Make sure this definition is inclusive of all number types

number == boolean ++ integer ++ rational ++ real ++ bit;
```

where `A++B` defines the union of bunches `A` and `B`.

---

[1]Note that the "++" operator forms bunches from other bunches. It will be explained further in a following section.

3.2.2 Boolean

Although `boolean` is a number type and obeys number conventions and axioms, its behaviors are important enough to justify separate discussion. The Rosetta `boolean` type is defined by the two element bunch `true,false` and is a subtype of `number`. Thus, anything of type `boolean` must take either `true` or `false` as it's value. Operators on booleans include all operators on numbers plus the following:

Operator	Format	Valid For
=,/=	A = B, A /= B	number
and, or, xor	A and B, A or B, A xor B	number
min, max	A min B, A max B	number
non-min, non-max	A nmin B, A nmax B	number
implies, implied by	A => B, A <= B	number
less,greater,greater or equals	A < B, A > B, A >= B, A =< B	number
not	- A	number

Rules of mathematics from infinite numbers apply to `true` and `false`, but they are not technically infinite. Nor are booleans are equivalent to machine limitations on numbers. Specifically, when treated as numeric values, `true` and `false` follow the following rules:

Property	Meaning
false = -true	`false` is equivalent to not `true`
-false = true	`true` is equivalent to not `false`
forall(x::number | x /= true => x < true)	`true` is the greatest number
forall(x::number | x /= false => x > false)	`false` is the least number
forall(x::number--true | x+1 < true	True cannot be generated from other numbers
forall(x::number--false | x-1 > false	False cannot be generated from other numbers

These properties assert that true and false behave like infinite values in that they cannot be generated from other numbers using increment or decrement functions. They do obey ordering operations and are greater (or less) than every other number.

An interesting consequence of defining `false` and `true` as the minimum and maximum possible numeric values is that the boolean `and` and `or` operators are equivalent to the `min` and `max` operators. Given the axiom defining `true` and `false`, we know that `false < true`. Thus, we have the following table defining `min` and `max` over `true` and `false`:

A	B	A min B	A max B
false	false	false	false
false	true	false	true
true	false	false	true
true	true	true	true

This is identical to the truth table defining `and` and `or`. The negation operator, `-`, also follows directly from the numeric interpretation of `boolean`. The greatest positive number negated is the least negative number. Thus, `-true = false`. As negation is its own inverse, we know that `-(-x)) = x` for any boolean value `x`. Thus, `-false = true`. The resulting truth table has the form:

A	-A
false	true
true	false

Definitions for other logical operations follow directly. Of particular interest is the definition of implication as:

```
A => B == -A or B
```

By definition, this is equivalent to `-A max B`. Again, consider the truth table generated by the definition of `true` and `false` as numeric values:

3.2 Elements

By definition, elements are values that cannot be described as being made up of other elements and element types are bunches of such values. Numbers such as 1, 5.32, and -32, characters such as 'a', 'B', and '1', and boolean values such as `true` and `false` represents such undecomposable values. In contrast, composite values such as records, sequences and sets are not elemental in that each is defined by describing its contents. Element values are also called *atoms* in Rosetta and *scalars* in traditional programming languages.

3.2.1 Numbers

Numeric types include standard bunches of values associated with traditional number systems. Predefined numeric types include `real`, `integer`, `natural`, `bit` and `boolean` and are listed in the following table:

Type	Format	Subtype Of
boolean	true, false	number
bit	1,0	natural
natural	123, true,	integer
integer	123,-123, false	real
rational	123/456	real
real	123.456, 1.234e56	number
number	*Any element of the above types*	element

Numeric values are represented by a strings of digits and optional sign, decimal point and/or exponential indicators. The format of a Rosetta number follows the canonical number format from traditional programming languages and will not be referred to here. The type `number` refers to any legal Rosetta number.

Predefined operators on numbers include: (Assuming `A` and `B` are numbers)

Operator	Format	Valid For
negation	- A	number
min,max	A min B, A max B	number
nmin,nmax	A nmin B, A nmax B	number
+,-,*,/,^	A+B,A-B,A*B,A/B,A^ B	number
=,/=,<,=<,>=,>,==	A<B,A=<B,A>B,A>=B	number
abs,sqrt	abs(A), sqrt(A)	number
sin,cos,tan	sin(A), cos(A), tan(A)	number
exp,log	exp(A), log(A)	number

All operators are defined in the traditional manner. The `nmin` and `nmax` operators are interpreted as "not min" and "not max" respectively and provide an inverse for those operators. The `^` operator represents raising a value to a power.

The numeric values `true` and `false` obey the following rule:

```
forall(x::number | false =< x and x =< true)
```

`true` acts as positive infinity and `false` acts as negative infinity. Consequences of this convention are numerous and useful. They include: (i) `max` is the same operation as `or`; (ii) `min` is the same operation as `and`; and (iii) `=<` is the same operation as implication (`=>`) and `>=` is the same operation as implied by (`<=`). The same laws apply to these operations in all cases, and the different signs are taken to be synonyms of each other, maintained here for the sake of historical recognition. Note that `true` and `false` are not associated with 1 and 0 as is the case with many traditional programming languages. Adopting this convention is appealing in a traditional sense, it prohibits semantic simplicity.

Chapter 3

Items, Variables, Values and Types

3.1 Items and Values

Rosetta's basic semantic unit is called an *item*. Item structures result when Rosetta descriptions are parsed prior to manipulation. Although most users will never deal directly with items, they present an effective way to describe the relationships between variables, values and types.

Informally, an item consists of a *label* naming the item, a *value* the item represents, and a *type* from which specific item values must be chosen. When any structure is defined in a Rosetta specification, an item is created with the specified label. Variables, constants, terms, even facets themselves are items in a Rosetta specification. When a label is referred to in a specification, it refers to the value of the item it is associated with. An items set of potential values is delineated by it's associated type. In a legal Rosetta specification, every item's value is an element of it's associated type. A more complete description of items can be found in the *Rosetta Semantics Guide*.

An item whose value is known is referred to as a constant item while a variable whose value is unspecified is referred to as a variable item or simply a variable. Any item of any type may be constant or variable including facets, terms and other traditionally constant items.

Values represent objects that can used as values for items. There are three general classes of values: (i) elemental; (ii) composite; and (iii) functional. Elemental values represent primitive, atomic values that are directly manipulated by Rosetta. Elemental values include such things as integers, naturals, characters, bits and boolean values. Traditional programming languages refer to elemental values as scalar. Composite values are constructed from other values. Composite values include such things as records, tuples, vectors, sets, bunches and facets. Functional values represent functions and relations. The name `universal` to refer to any value at all. Universal is **not** itself a bunch, as it must contain bunches.

All Rosetta types are *bunches*. A bunch is simply an unpackaged collection of items. Functions and properties for bunches are defined completely in Section 3.3.1. Throughout this document, the terms `bunch`, `type` and `subtype` are used interchangeably. The notation `a::T` represents bunch containment is interpreted as "a is contained in T". Although this notation is formally a logical relation, it is also used as a syntax for declaring variables and constants. Specifically, `v::T` in a declaration area of a facet declares a variable item of type `T` whose value is a *single* element of bunch `T`. This distinction will become clear when bunches and bunch relationships are described fully. Similarly, the notation `v::T=c` defines a constant item of type `T` whose value is determined by the expression `c`. This notation will be used and explained extensively in the following sections.

```
package AlarmClock is

  use timeTypes;

  facet mux(timeIn::in time; displayTime::out time; clockTime::in time;
            setAlarm::in bit; setTime::in bit) is
  begin state_based
    l1: %setAlarm => displayTime' = timeIn;
    l2: %setTime => displayTime' = timeIn;
    l3: %(-(setTime xor setAlarm)) => displayTime' = clockTime;
  end mux;

  facet store(timeIn::in time; setAlarm::in bit; setTime::in bit;
              toggleAlarm::in bit; clockTime::out time; alarmTime::out time
              alarmOn::out bit) is
  begin state_based
    l1: alarmTime' = if %setAlarm then timeIn else alarmTime endif;
    l2: clockTime' = if %setTime then timeIn else clockTime endif;
    l3: alarmOn' = if %toggleAlarm then -alarmOn else alarmOn endif;
  end store;

  facet counter(clockTime :: inout time) is
  begin state_based
    l4:: clockTime' = increment-time clockTime;
  end counter

  facet comparator(setAlarm:: in bit; alarmTime:: in time;
                   clockTime:: in time; alarm:: out bit) is
  begin state_based
    l1: alarm = %(setAlarm and (alarmTime=clockTime)) endif
  end comparator;

  facet alarmClockStruct(timeIn::in time; displayTime::out time;
                         alarm::out bit; setAlarm::in bit; setTime::in bit;
                         alarmToggle::in bit) is
    clockTime :: time;
    alarmTime :: time;
    alarmOn :: bit;
  begin logic
    store_1 : store(timeIn,setAlarm,setTime,toggleAlarm,clockTime,
                    alarmTime,alarmOn);
    counter_1 : Counter(clockTime);
    comparator_1 : comparator(setAlarm,alarmTime,clockTime,alarm);
    mux_1 : mux(timeIn,displayTime,clockTime,setAlarm,setTime);
  end alarmClockStruct;

begin logic
end AlarmClock;
```

Figure 2.1: The complete alarm clock specification

```
                              alarm::out bit; setAlarm::in bit; setTime::in bit;
                              alarmToggle::in bit) is

    clockTime :: time;
    alarmTime :: time;
    alarmOn :: bit;
  begin logic
    store_1 : store(timeIn,setAlarm,setTime,toggleAlarm,clockTime,
                    alarmTime,alarmOn);
    counter_1 : Counter(clockTime);
    comparator_1 : comparator(setAlarm,alarmTime,clockTime,alarm);
    mux_1 : mux(timeIn,displayTime,clockTime,setAlarm,setTime);
  end alarmClockStruct;
```

2.4.4   The Specification

The final specification enclosed in a Rosetta `package` is shown in Figure 2.1.

```
            then alarmTime' = timeIn
            else alarmTime' = alarmTime
        endif;
  l2:: if %setTime
            then clockTime' = timeIn
            else clockTime' = clockTime
        endif;
  l3:: if %toggleAlarm
            then alarmOn' = -alarmOn
            else alarmOn' = alarmOn
        endif;
  end store;
```

Counter

The `counter` component is the simplest component involved in the definition. It states that each time the clock is invoked, its internal time is incremented.

```
// counter increments the current time
use timeTypes;
facet counter(clockTime :: inout time) is
begin state_based
  l4:: clockTime' = increment_time(clockTime);
end counter
```

Comparator

The `comparator` implements the guts of the alarm clock's alarm function. It determines the appropriate value for the alarm output given the state of the alarm set bit and the values of the alarm time and the clock time. If the alarm is set and the alarm time and clock time are equal, then the alarm output is enabled. Again, the `%` operator is used to convert a boolean value into the bit value associated with the alarm output.

```
// comparator decides if the alarm should be sounded based on the
// setAlarm control input and if the alarmTime and clockTime are
// equal.
use timeTypes;
facet comparator(setAlarm:: in bit; alarmTime:: in time;
                 clockTime:: in time; alarm:: out bit) is
begin state_based
  l1: alarm = %(setAlarm and (alarmTime=clockTime)) endif
end comparator;
```

2.4.3   Structural Definition

The actual structural definition instantiates each component and provides appropriate interconnections.

```
// The alarm clock structure is defined by assembling the components
// defined previously.
use timeTypes;
facet alarmClockStruct(timeIn::in time; displayTime::out time;
```

```
            endif
        else h(t)
      endif;

  end timeTypes;
```

2.4.2 Structural Definition

The structural definition begins by defining facets representing each of the alarm clock components. Specifically, this includes: (i) a multiplexor for defining what values are displayed; (ii) a store for internal state values; (iii) a counter for incrementing the current time; and (iv) a comparator for determining when the alarm should be sounded.

Multiplexor

The `mux` definition describes a component that determines which of its data inputs, `timeIn` or `clockTime`, should be displayed by the clock. This determination is made by examining the control signals `setAlarm` and `setTime`. Three terms are defined that select an output based on the control inputs.

```
  // mux routes the proper value to the display output based on the
  // settings of the setAlarm and setTime inputs.
  use timeTypes;
  facet mux(timeIn::in time; displayTime::out time; clockTime::in time;
            setAlarm::in bit; setTime::in bit) is
  begin state_based
    l1: %setAlarm => displayTime' = timeIn;
    l2: %setTime => displayTime' = timeIn;
    l3: %(-(setTime xor setAlarm)) => displayTime' = clockTime;
  end mux;
```

Recall that the Rosetta operator `%` converts bit values into boolean values allowing bits to be used in implications directly.

Store

The `store` component is the store for the alarm clock's internal state. It operates by examining the control bit associated with each stored value. If the control bit is set, a new value is loaded from an appropriate input, or in the case of `alarmOn`, toggling the existing value. If the associated control bit is not set, then the stored value is retained.

```
  // store either updates the clock state or makes it invariant based
  // on the setAlarm and setTime inputs.  Outputs are invariant if
  // their associated set bits are not high.
  use timeTypes;
  facet store(timeIn::in time; setAlarm::in bit; setTime::in bit;
              toggleAlarm::in bit;
              clockTime::out time; alarmTime::out time
              alarmOn::out bit) is
  begin state_based
    l1:: if %setAlarm
```

Summary: The facet algebra supports combining facet definitions into new facet definitions. The `and` and `or` operations corresponding to product and sum operations over facets combine facets under conjunction and disjunction respectively. The `and` operation defines new facets with all properties from both constituent facets. The `or` operation defines new facets with properties from either facet.

```
%% Things are okay beyond this point...
```

2.4 The Alarm Clock Example

In Section 2.1, the alarm clock example was introduced as an example systems level specification. In this section, the alarm clock example is examined more carefully and a structural definition introduced. The example is completely specified to provide an overall view of a Rosetta functional specification.

2.4.1 The `timeTypes` Package

`timeTypes` is a general purpose package introduced and explained in Section 2.1. It contains basic data types and functions used in the definition of the alarm clock system and structural definition. The only construct used in this definition that may require some explanation is the comprehension quantifier, `sel`. This function implements set comprehension for bunches. It does so by taking as its argument a function that maps a bunch onto the booleans and returning all domain elements for which the function is true. Thus, the statement:

```
sel(x::natural | x =< 12)
```

examines all elements of the natural numbers and returns those that are less than 12. Because its return type is bunch, its use in defining a type is perfectly legal. Further note that both `hours` and `minutes` are subtypes of `type(natural)`. This indicates that both have bunches as values, not singleton elements.

```
package timeTypes is
begin logic

  hours :: type(natural) is sel(x::natural | x =< 12);
  minutes :: type(natural) is sel(x::natural | x =< 59);
  time :: type(universal) is record [h::hours; m::minutes];

  increment_time(t::time) :: time is
    record [h = increment-hours(t);
            m = increment-minutes(t)];

  increment_minutes(t::time) :: minutes is
    if m(t) =< 59
       then m(t) + 1
       else 0
    endif;

  increment_hours(t::time) :: minutes is
    if m(t) = 59
       then if h(t) =< 12
               then h(t) + 1
               else h(t)
```

19

Example 8 (Requirements and Constraints) *Reconsider the previously defined facets* `sort_req` *and* `sort_const`. *Recall that* `sort_req` *defined requirements for a sorting component while* `sort_const` *defined a power constraint over the same component. A sorting component can now be defined to satisfy both facets:*

```
sort :: facet is sort_req and sort_const;
```

Informally, `sort`: *(i) outputs a sorted copy of its input; and (ii) consumes only 5mW of power. Formally, the new facet* `sort` *is the* product *of properties from* `sort_req` *and* `sort_const`. *In this example, the interaction between constraints domain and other requirements domains are unspecified. Therefore, analysis of interactions will reveal little additional information. However, it is certainly possible to define a relationship between the* `constraints` *and* `state_based` *domains if desirable.*

Example 9 (Postcondition Specifications) *Consider again the specifications for* `sort_req` *and* `sort_op`. *The first facet specifies the requirements for a sorting component using a black-box, axiomatic style. The second facet defines sorting using a specific, operational algorithm. Like the constraint model and requirements models previously,* `sort_req` *and* `sort_op` *can be combined into a single sorting definition:*

```
sort :: facet is sort_req and sort_op;
```

Here, the composition behaves much differently. The state-based and models do interact in interesting ways. The composition of `sort_req` *and* `sort_op` *provides a pre- and post-condition for the operational sorting definition. The net effect is like an assertion in VHDL. However, the requirements are specified distinctly and are not intermingled in the operational definition. Thus, for this composition to be consistent, the operational specification must hold along with it's real time constraints and the axiomatic specification must hold defining pre- and post-condition requirements on the composition.*

Similarly, a sort specification can be developed that combines requirements, operational and constraint models:

```
sort :: facet is sort_req and sort_op and sort_const;
```

`F1 or F2` Facet disjunction states that properties specified by either terms T1 or T2 must be exhibited by the composition. Note that this is logical or, not exclusive or. The most obvious use of facet disjunction is combining different component models into a component family. The following example illustrates such a situation.

Example 10 (Component Version) *Consider the following definitions using sort facets defined previously:*

```
multisort::facet is sort_req and (bubble_sort or quicksort);
```

The new facet `multisort` *describes a component that must sort, but may do so using either a bubble sort or quicksort algorithm. While* `and` *is a product operator,* `or` *is a sum operator over facets.*

Other facet operations are defined and include negation, implication and equivalence. These will be presented in detail in a later chapter. The objective here is simply to demonstrate various facet composition operations and where they might apply in a specification.

In this definition, the facet `register` is "copied" and relabeled twice. In the first case, the new facet is named `r1` and in the second, `r2`. The `memory` variable associated with `r1` is referenced via `r1.memory` and similarly for `r2.memory`. Now there is no conflict and the elements of each component have unique references. This aspect of labeling is simple, but extraordinarily powerful.

Summary: Including facet definitions as terms supports structural definition through facet aggregation. Including and instantiating facets in definitions is achieved using relabeling. Instantiating facets replaces formal parameters with actual items. Unique naming forces these items to be shared among facets providing for communications. When a facet is renamed, all of its internal items are renamed making each instance of that included facet unique.

```
%% The following section is way out of date given the updates to the
%% facet algebra.  We need to rethink facet declaration to include
%% parameters (using a notation like functions) to make this happen
%% correctly.  I think we can use the same semantics.
```

2.3   Facet Composition

The essence of systems engineering is the assembling of heterogenous information in making design decisions. Rosetta supports this type of specification directly with operations collectively known as the *facet algebra*. The facet algebra provides mechanisms for defining new specifications by composing existing specifications using the standard operators `and`, `or`, and `not`.

In the context of facets, these are not logical operators. The operation `F1 and F2` does not have a boolean value. Instead, it defines a new facet with properties from both `F1 and F2`. Looking ahead, this operation provides us a mechanism for combining properties from several facets into a single facet.

Facets under composition must maintain the logical truths as specified by standard interpretations of logical connectives. For example, if `F3 = (F1 and F2)`, then `F3` is consistent if and only if `F1 and F2` is consistent (Note: `F1 and F3` are enclosed in parentheses because = has higher precedence than `and`). Facet composition is useful for specifying many systems level properties by combining properties from various facets. A new facet can be defined via composition by an expression of the following form:

    <name>(<paramlist>) is *<facet_expression>*;

where $< name >$ is the new name, $< paramlist >$ is an optional parameter list, and $< facet\_expression >$ is an expression comprised of facet algebra operations.

The following examples describe several prototypical uses of facet composition. Please note that domains used in these examples are defined in an accompanying document.

`F1 and F2`   Facet conjunction states that properties specified by terms T1 and T2 must by exhibited by the composition and must be mutually consistent. Further, the interface is $I_1 \cup I_2$ implying that all symbols visible in `F1` and `F2` are visible in the composition.

The most obvious use of facet conjunction is to form descriptions through composition. Of particular interest is specifying components using heterogeneous models where terms do not share common semantics. A complete description might be formed by defining requirements, implementation, and constraint facets independently. The composition forms the complete component description where all models apply simultaneously.

by each are true. Thus, the first term instantiates `andgate` with items `a`, `b` and `i` where `i` is an internally defined variable of type `bit`. Thus, the facet asserts that `i` is equal to `a and b`. The second term does the same except it asserts that `d` is equal to `i and c`.

Communication between facets is achieved by sharing items. Here, the items are variable items defined either in the parameter list or in the body of the including facet. This models instantaneous exchange of information between facets via variables. Later, channels will be introduced to provide means for defining connections with properties such as storage and delay.

Although similar to VHDL structural definition, this Rosetta definition style is semantically quite different. To understand this requires some understanding of labels and item labeling. The notation `l:   term` defines `term` and associates label `l` with it. Thus, the definition:

```
l1: andgate(a,b,i);
```

asserts `andgate(a,b,i)` as a term and associates label `l1` with it. Effectively, the definition renames `andgate` locally to `l1`. Thus, the terms `l1` and `l2` define facets equivalent to `andgate`, but with new names. The reasoning for this is demonstrated in any definition where components that locally define variables and constants have multiple instances. For example, consider the following incorrect specification:

```
facet register(i::in bitvector; o::out bitvector; load::in bit) is
  memory::bitvector;
begin state_based
  load1: if %load then memory'=i else memory'=memory endif;
  output: o'=memory;
end register;

facet registerx2(i1::in bitvector; i2::in bitvector;
                 o1::out bitvector; o2::out bitvector;
                 load::in bit) is
begin state_based
  register(i1,o1,load);
  register(i2,o2,load);
end registerx2;
```

Consider the `memory` variable associated with each register. In the above definition, `register.memory` reference to the `memory` variable in facet `register`. Unfortunately, there's no way to learn which register. Further, because the register variables share the same name in the facet, they must be equal.

The proper definition is:

```
facet register(i::in bitvector; o::out bitvector; load::in bit) is
  memory::bitvector;
begin state_based
  load1: if %load then memory'=i else memory'=memory endif;
  output: o'=memory;
end register;

facet registerx2(i1::in bitvector; i2::in bitvector;
                 o1::out bitvector; o2::out bitvector;
                 load::in bit) is
begin state_based
  r1:register(i1,o1,load);
  r2:register(i2,o2,load);
end registerx2;
```

```
end stackAsSeq;
```

This stack definition uses the `package` construct to present a series of direct definitions. No terms are needed to describe the behavior of the provided type. The stack type, S, is not an uninterpreted type but is defined as a sequence of type E. The basic stack operations are now defined on the stack type using concrete operations.

An interesting exercise is to consider the meaning of:

```
stack(E,sequence(E)) and stackAsSeq(E)
```

As we shall see later, facet composition states that properties of both `stack` and `stackAsSeq` must apply in the facet formed by `and`. Effectively, this new definition is consistent only if `stackAsSeq` obeys the axiomatic definition provided by `stack`. In essence, `stack` represents requirements while `stackAsSeq` represents an implementation of stack.

Summary: A facet is the basic unit of Rosetta specification. It consists of a label, optional parameter list, optional declarations, a domain and terms that extend its domain. Variable declaration is achieved using the notation `v::T` interpreted as *the value of v is contained in T*. Constants are similarly defined using the notation `v::T is c` interpreted as *the value of v is contained in T and is equal to c*. Domains provide a vocabulary for defining specifications. Terms extend domains to provide definitions for the specific components. Terms are declarative constructs that are accompanied by a label. Any label defined in a Rosetta specification may be exported and referenced using the canonical `facet-name.label` notation. By default, all labels are exported. However, an explicit export statement may be used in the declaration section to selectively control label export.

2.2 Facet Aggregation

An important system level specification activity is aggregation of facets into general purpose architectures. Rosetta supports this directly using *facet inclusion* and *facet labeling*. Facet inclusion occurs when a facet name is referenced in a facet term. Facet labeling occurs when a facet is given a new label.

Consider the trivial example of defining a three input and gate from two input and gates:

```
facet andgate(x::in bit; y::in bit; z::out bit) is
begin logic
  l1: z = x and y;
end andgate;

facet andgate3(a::in bit; b::in bit; c::in bit; d::out bit) is
  i:: bit;
begin logic
  l1: andgate(a,b,i);
  l2: andgate(i,c,d);
end andgate3;
```

The resulting definition is quite similar to structural VHDL without explicit component instantiation. The first facet clearly defines the behavior of a simple *and* gate while the second seems to use facets as terms. The terms l1 and l2 both reference `andgate` and are interpreted as stating that the definitions provided

```
            endif
        else t(h)
    endif;
  end timeTypes;
```

hours and `minutes` *are restricted subranges of natural number representing hours and minutes respectively.*
The notation `type(natural)` *indicates that* `hours` *and* `minutes` *are bunches, not singleton values. The* `sel`
operation provides a comprehension operator and is used to filter natural numbers. `time` *is a record containing*
an hour and minute value.

Three increment functions define incrementing time. `increment_time` *forms a record from the results of*
incrementing the current hours and minutes values. `increment_hours` *and* `increment_minutes` *handle in-*
crementing hour and minute values respectively. Note that the field names are used to reference hours and
minutes values respectively.

Example 7 (Stack definition) *For formal specification fans, a semi-constructive stack definition is in-*
cluded to describe an alternate means for function specification. Here, the traditional stack operations are
declared, but are not defined directly. The distinction with other function definitions being that no constant
definition appears in conjunction with the declaration. Assume here that there exist in the containing package
declarations for EType and SType. Then the specification takes the form:

```
facet stack is
push(E::Etype;S::Stype)::Stype;
pop(S::Stype)::Stype;
top(S::Stype)::Etype;
is_empty(S::Stype)::boolean;
empty::Stype;

begin logic
ax1:forall(e::Etype|forall(s::Stype|pop(push(e,s))=s));
ax2:forall(e::Etype|forall(s::Stype|top(push(e,s))=e));
ax3:forall(e::Etype|forall(s::Stype|not(is_empty(push(e,s)))));
ax4:is_empty(empty);
end stack;
```

This is a canonical constructive specification for a stack. In the declarations section, `push`, `pop`, *and* `top` *are*
defined to operate over stacks and elements. The axioms defined as `ax1` *through* `ax4` *constrain the values of*
functions in the traditional declarative fashion.

This specification style may prove uncomfortable for traditional VHDL users. An alternate definition uses
sequences to represent the stack:

```
package stackAsSeq(E::univ) is
begin logic

  S::subtype(sequence(univ)) is sequence(E);
  push(s::S; e::E) :: S is e;;s;
  pop(s::S) :: S is tail(s);
  top(s::S) :: E is s(0);
  empty::S is nil;
  is_empty(s::S) :: boolean is s=empty;
```

```
   setalarm: if setAlarm=1
               then alarmTime' = timeIn and displayTime' = timeIn
               else alarmTime' = alarmTime
          endif;
   displayClock: setTime = 0 and setAlarm = 0 =>
                  displayTime' = clockTime';
   tick: setTime => clockTime' = increment_time(clockTime);
   armalarm: if alarmToggle = 1
               then alarmOn' = -alarmOn
               else alarmOn' = alarmOn
          endif;
   sound: alarm' = alarmOn and %(alarmTime=clockTime);
end alarmClockBeh;
```

Inputs correspond to data and control values for the clock. `timeIn` contains the current time input and can be used to set either the alarm time or the clock time. `displayTime` is the time currently being displayed. `alarm` drives the audible alarm. `setAlarm` and `setTime` control whether the alarm time or clock time are currently being set. `alarmToggle` causes the alarm set state to toggle.

Local variables correspond to the state of the clock. `alarmTime` is the current time associated with sounding an alarm. `clockTime` is the current time. `alarmOn` is "1" when the alarm is set and "0" otherwise.

Exploring the specification indicates that each requirement is defined as a labeled term. Each term can be traced back to a requirement from the English specification. Term `setclock` handles the case where the clock time is being set. Term `setalarm` handles when the alarm time is being set. Term `armalarm` handles the toggling of the alarm set bit. `tick` causes the clockTime to be incremented. The clock time is incremented in the next state only when the clock time is not being set. Finally, the `sound` term defines the `alarm` output in terms of the `alarmOn` bit and whether the `alarmTime` and `clockTime` values are equal. The "%" notation transforms the boolean result of equals into a bit value. All terms must be simultaneously true. Thus, the specification has the same effect as using multiple processes in VHDL.

The alarm clock facet uses the following collection of time manipulation functions and types:

```
package timeTypes is
begin logic
  hours :: type(natural) = sel(x::natural | x =< 12);
  minutes :: type(natural) = sel(x::natural | x =< 59);
  time :: type(universal) is record [h::hours; m::minutes];

  increment_time(t:: time) :: time is
    record [h = increment_hours(t);
            m = increment_minutes(t)];

  increment_minutes(t:: time) :: minutes is
    if t(m) < 59
       then t(m) + 1
       else 0
    endif;

  increment_hours(t::time) :: hours is
    if t(m) = 59
       then if t(h) < 12
               then t(h) + 1
               else 1
```

```
            endif
      endif;

   facet sort_op(i::in T; o::out T) is
   begin continuous
      l1: o@(t+5ms) = [qsort ~ i];
   end sort_op;

   export sort_op;

end work;
```

Here the function specifications are removed from the facet specification. The facet and functions are included in the package work. The similarity to VHDL here is intentional. Unlike VHDL, the package is parameterized allowing specialization for arbitrary types. Note the inclusion of the export sort_op clause. This causes the sort_op facet to be visible outside the package. Other declarations such as qsort, lside and rside are hidden in the package.

Why the obsession with sort? Thus far, an axiomatic, continuous time and operational continuous time specification have been developed. Together, we can use all three specifications to define various characteristics of a single sorting component in a manner unique to Rosetta. Specifically, in the next section we will define how a designer can specify a sorting component by combining specifications from multiple domains. The result is a requirements specification, a temporally constrained requirements specification, an operational specification, and a power specification simultaneously describing a system. With the addition of facet composition operators, this provides a powerful mechanism for mixing and composing specifications.

Example 6 (Alarm Clock System) *Consider the following definition of an alarm clock taken from the Synopsys synthesis tutorial. This alarm clock provides a basic capability for setting time, setting alarm, sounding an alarm and keeping time. The specification states the following requirements:*

1. *When the setTime bit is set, the timeIn is stored as the clockTime and output as the display time.*

2. *When the setAlarm bit is set, the timeIn is stored as the alarmTime and output as the display time.*

3. *When the alarmToggle bit is set, the alarmOn bit is toggled.*

4. *When clockTime and alarmTime are equivalent and alarmOn is high, the alarm should be sounded. Otherwise it should not.*

5. *The clock increments its time value when time is not being set.*

The systems level description of the alarm clock is defined in the following facet:

```
   use timeTypes;
     facet alarmClockBeh(timeIn::in time; displayTime::out time;
                         alarm::out bit; setAlarm::in bit; setTime::in bit;
                         alarmToggle::in bit) is
       alarmTime :: time;
       clockTime :: time;
       alarmOn :: bit;
     begin state_based
       setclock: setTime=1 =>
                    clockTime' = timeIn and displayTime' = timeIn;
```

12

```
   rside(pivot::T; i::sequence(T))::sequence(T) is
      if i=nil
         then nil
         else if i(0) > pivot
                 then i(0);;lside(pivot, tail(i))
                 else lside(pivot, tail(i))
              endif
      endif;

begin continuous
   l1: o@(t+5ms) = [qsort ~ i];
end sort_op;
```

This specification is interesting due to its similarity to a VHDL specification and its equivalence to sort_req. The sort_op specification specifies that the output parameter 5ms in the future is equal to the result of applying quicksort to the input parameter i. The details of the application are unimportant. Suffice to say that excluding the concept of a wait statement, this is quite similar to how a VHDL specification might be defined.

The function qsort and the auxiliary functions lside and rside define a quicksort algorithm over sequences. The definition follows the classic recursive style. As with other function definitions in these examples, these functions require some further study to understand completely. At this point it is important only to understand that parameters to the function are specified as param-var::param-type, separated by the ";" token and enclosed within parantheses. The final expression defines the return value. In the case of lside, all values less than or equal to the pivot value are found recursively and returned.

A potentially cleaner specification might have the form:

```
package work(T::univ) is
begin logic

   qsort(i::sequence(T))::sequence(T) is
      let (pivot::T is t(0); t::sequence(T) is tail(i)) in
          if i=nil
             then i
             else qsort(lside(pivot,t));;pivot;;qsort(rside(pivot,t))
          endif;

   lside(pivot::T; sequence(T))::sequence(T) is
      if i=nil
         then nil
         else if i(0) =< pivot
                 then i(0);;lside(pivot, tail(i))
                 else lside(pivot, tail(i))
              endif
      endif;

   rside(pivot::T; i::sequence(T)):: sequence(T) is
      if i=nil
         then nil
         else if i(0) > pivot
                 then i(0);;lside(pivot, tail(i))
                 else lside(pivot, tail(i))
```

Example 3 (Sort Constraints) *An alternative view of a component models performance constraints. The following definition models the power consumption constraints of a sorting component.*

```
    facet sort_constr
      power::real;
    begin constraints
      p1: power =< 5mW;
    end sort_constr;
```

The variable `power` is a real number representing power consumed by the component. The facet body defines a single term that limits power consumption to be less than or equal to 5mW. Both the semantics of constraints and the unit constructors required to define 5mW are defined in the constraints facet.

Example 4 (Timed Sort) *The facet `sort_timed` is an alternative definition of sort that places timing constraints on the definition. Here, instead of modeling what is true in an abstract next state, the sort is specified with respect to its behavior over time.*

```
facet sort_timed(T::type; i::in array(T); o::out array(T)) is
  use array_utils(T);
begin continuous
  l2: permutation(o@(t+5ms),i);
  l1: ordered(o@(t+5ms));
end sort_timed;
```

This definition uses the `continuous` domain rather than the `state_based` domain. The notation `x@t` refers to the value of x at time t. The term `l2` states that the output, o, 5ms in the future must be a permutation of the current input, i. The term `l1` states that the output must be ordered 5ms in the future.

No notion of next state as used previously is defined. However, this specification provides more detail in the form of hard timing constraints. Using the `continuous` domain, the user is allowed to define values of variables at specific times with respect to the current time t.

Example 5 (Operational Sort) *The facet `sort_op` provides an operation definition for a sorting algorithm by "implementing" a quicksort algorithm that will sort the input. Specifically:*

```
facet sort_op(i::in T; o::out T) is

  qsort(i::sequence(T)):: sequence(T) is
    let (pivot::T is t(0); t::sequence(T) is tail(i)) in
        if i=nil
          then i
          else qsort(lside(pivot,t));;pivot;;qsort(rside(pivot,t))
        endif;

  lside(pivot::T; i::sequence(T)):: sequence(T) is
    if i=nil
      then nil
      else if i(0) =< pivot
              then i(0);;lside(pivot, tail(i))
              else lside(pivot, tail(i))
           endif
    endif;
```

Example 2 (array_utils Package) *Packages are a parameterized mechanism for grouping together definitions. They are defined using the semantics of facets and will be discussed fully in a later section. Here, the definition of the array_utils package used by the sort_req facet is defined:*

```
package array_utils(T::univ) is
begin logic

  // numin - return the number of occurrences of x in i
  numin(x::T; i::array(T)):: natural is
    if i=[] then 0
            else if x=i(0) then 1+numin(tail(i))
                           else numin(tail(i))
                endif
    endif;

  // permutation - determine if a1 is a permutation of a2
  permutation(a1::array(T); a2::array(T)):: boolean is
    forall(x::T | numin(x, a1) = numin(x, a2));

  // ordered - determine if a1 is ordered.  =< must be defined on T
  ordered(a1::array(T)):: boolean is
    forall(i :: sel(x::natural| x =< #a1-1) | a(i) =< a(i+1));

  // tail - return the tail of an array.  based on sequence tail.
  tail(a1::array(T)):: array(T) is [head(~ a1)];

end array_utils;
```

The array_utils package defines four general purpose functions for arrays: (i) numin; (ii) permutation; (iii) ordered; and (iv) tail. It is difficult to explain these definitions fully without deeper understanding of Rosetta function definition. However, some exploration will aid in understanding and writing more complex specifications.

As an example, examine the definition of permutation:

```
  permutation(a1::array(T); a2::array(T)):: boolean is
    forall(x::T | numin(x, a1) = numin(x, a2));
```

This definition can be divided into two parts. First, the signature of permutation is given as

```
  permutation(a1::array(T); a2::array(T)):: boolean
```

The function name is permutation, (a1::array(T); a2::array(T)) are the domain parameters, and boolean is the return type.

The second part of the definition, following the keyword is, denotes the value of the return expression. The expression specifies the permutation. It is true when every element of T occurs in a1 and a2 the same number of times. It is false otherwise. The syntax of function declaration and the semantics of forall and other constructs are defined later.

Other functions are similarly defined. numin determines the number of occurrences of a value in an array using a simple recursive definition. ordered defines a predicate that is true when every element of its argument array is greater than or equal to the preceding element. Finally, tail for arrays is defined by extracting the elements into a sequence, finding the tail, and recreating an array. Remember, to fully understand these definitions requires further knowledge of Rosetta type and function semantics that will be presented later.

9

2.1.1 Examples

Examples are included here to provide motivation for the facet syntax and to provide context for the following sections. It is intended that these examples provide an overview, not a detailed language description. It is suggested that these be referred to while reading subsequent chapters as a means for understanding the utility of Rosetta definition capabilities.

Example 1 (Sort Definition) *A declarative specification for requirements and constraints associated with a* sort *function has the following form:*

```
facet sort_req(i::in array(integer); o::out array(integer)) is
  use array_utils(integer);
begin state_based
  l2: permutation(o',i);
  l1: ordered(o');
end sort_req;
```

The facet sort_req *defines a view of a component that accepts an array of integers as input and outputs the array sorted. This simple specification demonstrates several aspects of Rosetta specification using the* state_based, *axiomatic style.*

Parameters for sort_req *are simply an input and output arrays of type* integer. *The facet uses the* state_based *domain allowing the use of* o' *to represent the output in the state following execution. The package* array_utils *(defined later) is included to provide definitions necessary for defining sort. Specifically,* permutation *and* ordered. *These functions could be defined in the declaration section of the package, however this definition is cleaner and allows reuse of the array utilities in other specifications. Note that the* array_utils *package is parameterized over a type. This parameterization is used to specialize the* array_utils *for any appropriate type.*

It is possible to write a sort_req *definition that is parameterized over the contents of the input and output array. This implementation sorts arrays of integers. Although this may be interesting from a pedagogical perspective, it is not particularly useful or reusable. The following definition parameterizes the facet definition over an arbitrary type, T:*

```
facet sort_req(T::type; i::in array(T); o::out array(T)) is
  use array_utils(T);
begin state_based
  l2: permutation(o',i);
  l1: ordered(o');
end sort_req;
```

In this new sort_req *facet, the type T associated with the contents of the input and output arrays is a parameter. This allows specialization of the* sort_req *facet for various array contents. The only restriction being that an ordering relationship must be defined on the array elements.*

The following instantiation of the parameterized sort_req *is equivalent to the original* sort_req *facet:*

```
  sort_req(integer=>T);
```

This usage replaces all instances of T in the facet with the type integer. *The resulting facet is semantically identical to the original* sort_req *definition.*

where `l` is the label associated with the term and `term` is the definition itself. The label is used to reference the term in other definitions as well as when the term is exported. All terms defined in scope of the `begin-end` pair are considered top level terms.

The `register` uses two terms to define behavior. The first, labeled `l1`, defines the register's next state in terms of its current state, input and control inputs. The `if` statement implements the various cases for hold, shift right, shift left, and load. It should be noted that the shift operations are implemented using the built in bitvector shift functions `lshr` and `lshl` that provide logical shifts over bitvector types. The second term, labeled `l2` defines the next output. This simple expression states that the next output is the same as the next state as defined by term `l1`.

It should be noted that both terms defined in `register` hold simultaneously. Thus, both the next state and output definitions must hold for the component to behave correctly. The structure of the specification is much like the structure of a VHDL specification. Each state variable and output parameter is handled individually. The distinction here is the variability of definition semantics. In this case, the Rosetta function semantics is used to calculate next values for each variable.

The domain extends the semantics available for specification by adding new definitions and potentially syntax specific to the domain being specified. In the case of `state_based`, the basic addition is the concept of current and next state. Specifically in the register definition, `state` refers to the register contents in the current state while `state'` refers to the register contents in the following state. The `state_based` domain defines the semantics of "`x'`".

In declarative requirements facets, domains and associated terms are characterized by logical expressions like those used in `register`. In operational descriptions, terms may be imperative or functional program fragments. Regardless, the syntax and semantics of terms are determined by the semantics defined by the facets domain.

Parameter instantiation is achieved by traditional universal quantifier elimination. An object of the specified type is selected and the parameter replaced by that object. When formal parameters are instantiated with objects, those objects replace instances of parameters throughout the facet specification. When `A` is an actual parameter and `F` is a formal parameter, the notation `A=>F` allows direct assignment of actual parameters to formal parameters. This notation allows partial instantiation and is sometimes necessary when parameter ordering in constructed facets is ambiguous.

Consider the following modified `register` specification:

```
facet register(i::in bitvector; o::out bitvector; s0::in bit; s1::in bit) is
  state::bitvector;
  export state;
begin state_based
  l1: if s0=0 then
          if s1=0 then state'=state
             else state'=lshr(state) endif
          else if s1=0 then state'= lshl(state)
             else state'=i endif
       endif;
  l2: o'=state';
end register;
```

This specification is identical to the previous definition except that only the `state` variable is visible outside the facet scope. The terms `l1` and `l2` are no longer visible as they are not listed in the export clause. The variable `state` is accessed using the name `register.state` because `register` is the label assigned to the facet.

```
facet register(i::in bitvector; o::out bitvector; s0::in bit; s1::in bit) is
  state::bitvector;
begin state_based
  l1: if s0=0 then
        if s1=0 then state'=state
          else state'=lshr(state) endif
      else
        if s1=0 then state'= lshl(state)
          else state'=i endif
      endif;
  l2: o'=state';
end register;
```

This definition describes a facet `register` with data parameters `i`, and `o` of types `bitvector` and two control parameters of type `bit`. The variable `state` is defined to hold the internal state of the register and is of type `bitvector`. As can be deduced from examination of the specification, this register performs hold, logical shift right and left, and load operations given inputs of 00, 01, 10 and 11 on parameters `s0` and `s1` respectively.

All Rosetta variables and parameters are declared using the notation `v::T` where `v` is a variable name and `T` is a type. The "::" notation is used for declarations as it represents bunch membership and all types are in fact bunches. The notation `x::T` creates an item labeled `x` whose values are associated with type `T`. The declaration can be viewed as declaring an item `v` whose possible values are selected from the bunch associated with `T`. In the `register` specification, `state::bitvector` defines a variable labeled `state` whose values can be selected from the bunch type `bitvector`.

Parameters are universally quantified variables visible over the scope of the facet. Parameter definitions are similar to traditional declarations with the addition of a descriptor predicate. Specifically, `i::in bitvector` defines a parameter `i` of type `bitvector` and declares the predicate `in(i)` to be true. The semantics of `in(i)` are defined by the semantic domain currently being used. In the `register` example, the `state_based` domain defines `in` to explicitly disallow reference an input parameter in the next state. Specifically, `i'` is disallowed.

The declaration section following the facet interface includes declarations local to the facet. Items defined in this manner are visible throughout the facet. Such declarations may be made visible outside the facet using an `export` statement. In this case, the exclusion of an `export` clause makes all labels visible outside the specification.

When referenced in the facet body, a term, variable or parameter is referenced by its label without decoration. When referenced outside the facet, labels are referenced using the facet name as a qualifier. In the register example, `register.l1` refers to the first term in `register` while `register.state` refers to the variable `state`.

The `begin-end` pair delimits the domain specific terms within the facet and declares the facet's domain. The `begin` statement opens the set of terms and identifies the semantic domain those terms will use. In the `register` specification, the semantic domain is `state_based` providing the basic semantics for state and change in the traditional axiomatic style. Specifying a semantic domain indicates what domain theory the facet uses for its definition. Every facet must have an associated domain even if that domain is the `logic` domain common to all facets.

Terms in the term list define the behavior modeled by the facet. Each term is a labeled, well formed formula (wff) with respect to the facet's semantic domain.

The general form associated with any term is:

```
l: term;
```

Chapter 2

Facet Basics

The basic unit of specification in Rosetta is termed a `facet`. Each facet defines a single aspect of a component or system from a particular perspective. To define facets completely, it is necessary to understand the basics of Rosetta declarations, functions and expressions. This chapter intends only to introduce the concept and simple examples of facet definition to motivate the descriptions in following chapters. If concepts are not fully presented here, assume they will be in chapters dealing with the specifics of facet definition.

A facet is a parameterized construct used to encapsulate Rosetta definitions. Facets form the basic semantic unit of any Rosetta specification and are used to define everything from basic unit specifications through components and systems. Facets consist of three major parts: (i) a parameter list; (ii) a collection of declarations; (iii) a domain; and (iv) a collection of labeled terms. This section introduces the facet syntax, an *ad hoc* facet semantics, and provides structure for the remainder of the document. For a formal definition of facet semantics, please see the *Rosetta Language Semantics Guide*.

2.1 Facet Definition

Facets are defined using two mechanisms: (i) direct definition; and (ii) composition from other facets. In this section we will deal only with direct definition and defer facet composition to Section 2.3. Direct definition is achieved using a traditional syntax similar to a procedure in a traditional programming language or a theory in an algebraic specification language. The general formal for a facet definition is as follows:

```
facet <facet-label>(<parameters>) is
    <declarations>
begin <domain>
  <terms>
end <facet-label>;
```

The facet definition is delineated by the `facet` keyword immediately followed by a $< facet - label >$ providing the facet with a unique name. The facet label is immediately followed by a comma separated parameter list denoted above by $< parameters >$ and the keyword `is`. The declarations section, denoted by $< declarations >$, is used to declare labeled items and define visibility of locally defined labels. The keyword `begin` starts the definition section and is immediately followed by domain theory, denoted $< domain >$, used to provide a vocabulary for the definition. Declarations follow in the form of labeled terms, denoted $< terms >$ that provide a definition for facet. The definition concludes with the keyword `end` and the facet label.

As an example, a specification for a `find` component follows:

Chapter 1

Introduction

This document serves as a usage guide for the Rosetta specification language. It defines most of the base Rosetta semantics in an *ad hoc* fashion and provides general usage guidelines through examples.

The basic unit of Rosetta specification is a *facet*. Each facet is a parameterized collection of declarations and definitions specified using a domain theory. Facets are used to define: (i) system models; (ii) system components; (iii) architectures; (iv) libraries; and (v) semantic domains.

Although definitions within facets use many different semantic representations, the semantics of facet composition, inclusion and data types are shared among all facets. Collections of facets are composed using a collection of common operations that operate regardless of semantic domain. The basic facet definition provides an encapsulation, parameterization and naming convention for Rosetta systems.

This document describes the semantics of facets in an ad hoc fashion. Its intent is to provide an introduction to facets and their various uses without discussing any specific domain theory. In addition to facets themselves, this document also defines a type system shared among facet definitions. Basic types and operations available to Rosetta specifiers are identified and primitive definitions provided. Finally, the system construct used to describe hegerogeneous systems using facets is presented. The system construct supports definition of facets, assumptions, declarations, and verifications in support of a systems level design activity.

Contents

Rosetta Usage Guide

Perry Alexander and *Cindy Kong*
Information & Telecommunication Technology Center
The University of Kansas
2291 Irving Hill Rd.
Lawrence, KS 66044-7541
{alex,kong}@ittc.ukans.edu

David Barton
Averstar, Inc.
Vienna, VA
dlb@wash.inmet.com

January 16, 2002